

Building trust into light-handed regulations for cognitive radio

Kristen Ann Woyach

Electrical Engineering and Computer Sciences
University of California at Berkeley

Technical Report No. UCB/EECS-2013-236

<http://www.eecs.berkeley.edu/Pubs/TechRpts/2013/EECS-2013-236.html>

December 20, 2013



Copyright © 2013, by the author(s).
All rights reserved.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission.

Building trust into light-handed regulations for cognitive radio

By

Kristen Ann Woyach

A dissertation submitted in partial satisfaction of the
requirements for the degree of

Doctor of Philosophy

in

Engineering – Electrical Engineering and Computer Sciences

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Professor Anant Sahai, Chair
Professor Brian Carver
Professor Vern Paxson
Professor Jean Walrand

Fall 2013

Building trust into light-handed regulations for cognitive radio

Copyright 2013
by
Kristen Ann Woyach

Abstract

Building trust into light-handed regulations for cognitive radio

by

Kristen Ann Woyach

Doctor of Philosophy in Engineering – Electrical Engineering and Computer Sciences

University of California, Berkeley

Professor Anant Sahai, Chair

This thesis introduces an incentive-based trust model to let wireless spectrum regulation embrace diverse current and future means of implementing cognitive radio.

Cognitive radio has emerged as a way to combat inefficient spectrum use by allowing independently designed networks to share the same frequency band. This philosophy has been embraced by the FCC, which has already allowed cognitive use in the TV bands, and plans to make spectrum sharing the norm in other bands as well. To enact spectrum sharing, regulatory decisions, like band assignment, are made at runtime so that they can reflect local context. From a regulatory perspective, the most important question is how to trust that these decisions will be made and carried out correctly.

Right now, the FCC guarantees correct decisions by directly testing that any deployed technologies are incapable of making bad decisions. This process of testing is called certification. But certification has limitations. For example, a network of nodes could sense for a TV signal and decide as a group that the TV tower is far enough away that their interference to TV receivers would be negligible. However, this network will never pass a certification test. There is no way to prove that the network will stay silent if all the nodes are blocked by the same building so they cannot sense the tower but can cause interference.

This thesis provides a new model for trust that would allow networked sensing and any other novel spectrum sharing solution through light-handed regulations. The idea is to build a system that allows regulators to trust secondaries to follow sharing rules *regardless of whether they are technically capable of finding spectrum holes*. This is accomplished by an incentive mechanism, a *spectrum jail*, that will punish secondaries caught causing interference by degrading their quality of service. This thesis shows that for such a mechanism to work, cognitive radio must be thought of as a *band-expander*. If the same mechanism must apply to all radios, regardless of technology, there must be pretty good unlicensed or licensed bands that secondaries can use if they cannot share spectrum appropriately.

The mechanism explored here is inspired by the ideas in the law and economics literature as well as the spectrum policy literature. This thesis takes these mostly rhetorical arguments and develops the first mathematical model for incentive-based trust in spectrum regulation.

This model allows identification of the most difficult to enforce cases: the regulator must decide whether a primary will be protected even if it hardly ever uses its band. The regulator

must also decide what constitutes harmful interference. Some interference is unavoidable when bands are shared; the regulator must decide how much interference the primary must accept in a shared environment. When these decisions are made, this thesis shows that trust can be guaranteed with a sanction set *at certification time* and which is applied to *all* cognitive devices regardless of technology.

The model also gives quantitative performance metrics, measuring the ability for secondaries to reclaim spectrum holes, which illustrate the dependence on the regulator’s ability to catch wrongdoers. In particular, this thesis shows that while trust depends on the ability to catch those causing interference, runtime performance depends on the wrongful conviction rate. So, even applying the same sanction, as spectrum sharing technology and catching technology improves, performance will improve as well.

This model is extended to understand what role the primary can or must play in its own protection as new primary devices are developed to operate in a shared band. By controlling the cost of reporting, the regulator can trust a primary to report interference correctly. This also means that if a secondary is difficult to coexist with, the primary will not use the jail system to try to get rid of the secondary. It will instead hire a “band-sitter,” which is a preferred secondary system that coexists more easily with the primary.

This thesis also addresses multiple secondaries and aggregate interference by giving a basic framework of results to guide research in this direction. The distribution of aggregate interference from randomly placed nodes is explored to understand *placement risk*: the threat of too much interference caused by clusters of secondaries too close to the primary. Then, the thesis develops strategies to use the secondary location information that TV whitespace databases already have to address the problem of placement risk. Finally, a basic queuing model is suggested as a future direction to extend spectrum jails to deal with multiple secondaries.

Finally, this thesis answers the question of why jails? The original motivation is two-fold. First, jails lend themselves to simple modeling because the utility and the sanction are both measured in quality of service terms. Second, jails can actually be reasonably implemented. The FCC has allowed TV whitespace devices using databases to coordinate spectrum access. In order to actually secure this operation, databases will need to be able to identify malfunctioning devices and turn them off. These same identity and kill-switch technologies will also enable spectrum jails. Jails can even be implemented through the databases themselves as a denial of operating tokens.

At a more philosophical level, in-kind and monetary sanctions are fundamentally different things. Which one is actually better suited to the spectrum sharing enforcement problem? The last chapter will apply the same performance-based understanding from the the rest of the thesis to understand when fines or in-kind punishments should be preferred. It shows that in cases of high uncertainty, or when primary protection is the most important consideration, in-kind sanctions are the right approach.

The Thesis in a Nutshell

Cognitive radio aims to enable efficient spectrum sharing between different wireless systems that are not jointly designed. This new regulation paradigm cannot come about all at once. So, the initial stages of the transition include protecting the current wireless devices (primaries) from the potential interference caused by new devices (secondaries). There are therefore two sides to this problem: creating the technology that allows sharing to occur, and creating the regulations that guide technology to guarantee peaceful coexistence.

Technology has come a long way to make spectrum sharing possible, but regulations have not yet caught up to make these technology solutions legal. For example, take TV towers and receivers as the primary. Secondary devices could individually sense the surrounding spectrum, determine as a network that TV towers are far away, and begin to transmit. This is not legal. Perhaps a secondary network could determine that there were no primaries in a particular direction and beamform its transmission only in that direction. This is also not legal. A secondary network operator could build their entire network in subterranean Faraday cages, but could want to use TV frequencies for their superior transmission capabilities. If the network is underneath a TV tower, the operator would be breaking the law.

If these peaceful coexistence scenarios are technically feasible, why are they illegal? It is fundamentally a question of trust. Right now, the FCC certifies every deployed wireless device. The purpose is to guarantee that any deployed device will do exactly what it claims to do. Unfortunately, certification cannot provide such a guarantee for the examples above: the subterranean network, if deployed above ground, could cause significant interference. If the beam-forming solution did not correctly determine the relative position of primaries, it too would interfere. The network nodes could be deployed too close together within the shadow of a building. They cause interference, but cannot sense the TV themselves. With certification alone, *there is no way to trust devices that collectively change their operation based on local context.*

If enabling trust is the actual problem, are there other ways to do this that would allow more technical freedom? This thesis presents the first mathematical model for extending trust in cognitive radio *through vulnerability instead of strict control*. The effect of this new model is to allow *trust first and performance later*, the basis of light-handed regulation. Essentially, the certification procedure includes guaranteeing only that the secondary network is subject to a punishment-based enforcement mechanism (spectrum jails). Then, regardless of the secondary's method of finding spectrum holes, it will be in the best interest of the secondary to follow sharing rules.

This summary goes through the main results of this thesis. At the end is a table of results, broken down by chapter.

The idea of spectrum jails

The incentive mechanism explored in this thesis is spectrum jails. The idea is an in-kind punishment that will decrease the quality of service (QOS) that a secondary can achieve if it insists on causing harmful interference. The punishment must be aligned with the QOS desires of the secondary – this thesis covers secondaries who try to maximize their throughput (Chapter 2), energy usage (Chapter 4), or a combination of both (Chapter 4). The jail includes a time-out during which the secondary is not allowed to transmit in any of its available bands, which impacts the throughput-sensitive users. The jail also includes a requirement to burn energy during the time-out, which impacts the energy-sensitive users.

Spectrum jails are envisioned to be implemented through the databases that are already operational in the TV whitespaces. Right now, these databases accept GPS coordinates from secondaries, and return a list of available channels. The intent of the FCC is to port these databases into other bands, like the federal spectrum as in the PCAST report (see Chapter 1). It may be very difficult at first to roll out databases that can coordinate spectrum access for spectrum holes that only last for minutes or seconds. On the other hand, secondaries *may* be able to find these spectrum holes if given the freedom to do so, either through sensing technology or other kinds of sensing infrastructure.

This thesis proposes that the database instead act as a *spectrum manager* and issue “operation tokens” which the secondary needs in order to operate in any of its bands. This token does not tell the secondary where it is free to transmit. The token tells the secondary it is free to look for a band to transmit in. These tokens could be issued on a slower time scale and revoked if harmful interference is detected. The lack of an operation token is equivalent to being in spectrum jail. The database implementation is analyzed in Chapter 3.

General themes

The strength of this mathematical model of spectrum jails is that it makes explicit the difficult-to-enforce cases and the performance tradeoffs. The model also highlights the different architectural requirements and regulatory decisions that must be made for such an enforcement system to work.

The architectural requirements include an identification system with certain performance specifications on how well it can connect interference to the culprit. To guarantee primary protection, the identification system must be able to catch those causing interference. To reduce the negative impact of spectrum jails on secondaries following the rules, the identification system must have a low rate of wrongful conviction. Note that this identification system (as argued in Chapter 1) *is already needed to actually trust that secondaries in the TV bands will follow sharing rules*. Spectrum jails do not have an extra requirement, they are leveraging the solution to a problem that already exists.

It is also required that cognitive radio be treated as a *band-expanding* technology instead of the only means of spectrum access. If certification does not depend on the secondary’s technology, then secondaries with poor technology need somewhere to operate other than the primary bands. Otherwise, these secondaries’ only choice is to cause harmful interference.

The mathematical model presented here also makes explicit the decisions regulators must make about harmful interference and when primaries will be protected. There will always be some interference when different wireless systems share the same frequency band. The model of trust

here allows regulators to choose exactly how much harm is too much harm and then design the enforcement system to protect primaries from any interference over this level. These harm decisions have always been made; the only difference here is that they are made directly instead of implicitly through no-talk radii or the size of guard bands.

But primaries cannot be protected in all situations – if the primary uses the band only rarely, only a very high punishment will guarantee deterrence of harmful interference because rational secondaries are choosing between checking for the primary all the time and suffering punishment only rarely. So, regulators must also adopt a “use-it-or-lose-it” rule for shared bands that requires the primary to actually use its allocation.

So, primaries that are active enough of the time can be protected from too much interference. What about the secondaries? How much of an impact does spectrum jail make on their performance? It depends on the quality of the identity system and how well/easily the secondary can correctly find spectrum holes. As these technologies become better, the impact of spectrum jail on secondaries will become negligible.

Extending the role of the primary

Current legacy users are designed to operate alone in an exclusive band. So the rules (and enforcement mechanisms) must be designed to protect these legacy users. In the future, however, primaries will be designed to operate in a shared band. Chapter 5 explores whether the primary should be required to be designed to take part in the enforcement system.

In general, the primary is in the best position to report harmful interference, and it does have incentive to do so. However, it also has incentive to report correctly received packets as dropped in order to wrongfully send the secondary to jail. The regulator has some amount of control over these actions by choosing how much it costs to report interference. This could be done, for example, through a mini jail sentence applied to the primary for every accusation. But the cost of reporting must be chosen carefully – if the cost of reporting is too high, there is less incentive to report actual interference. If it is too low, the primary may want to report all (received or dropped) packets as dropped.

The primary has another option to influence the actions of the secondary. The primary has some number of packets it needs to transmit, but there is no mechanism forcing it to transmit only the packets it needs to. The primary can either transmit gibberish to make the band seem more occupied, or it can hire a band-sitter. This is a preferred secondary whose transmissions the primary will protect to either make other secondaries more honest or to kick out other secondaries. Chapter 5 explores when this option will be used, and how the cost of reporting must change to make hiring a band-sitter preferred over false accusations for protecting the primary band.

Multiple secondaries and aggregate interference: protecting the primary against “placement risk”

Extending spectrum jails to multiple secondaries requires understanding the equilibrium spectrum jails are trying to achieve. To this end, Chapter 6 explores “placement risk,” or the chance of too much aggregate interference because of randomly-located secondaries clustering around the primary receiver.

The distribution of aggregate interference from randomly placed secondaries changes its form depending on how dense the secondaries are compared to the distance of the closest interferer

to the primary receiver. We call the distribution coming from sparse interferers (rural environments) “near field.” The distribution has a heavy tail and is dominated by the interference from the closest node. We call the distribution of interference coming from dense interferers (urban environment) “far field;” it has a Gaussian-like form. Understanding the shift from near-field to far-field is important because the rule for a no-talk radius, for example, may have to change in rural and urban environments to adequately protect the primary.

With this distribution information, Chapter 6 explores different sharing rules a database may be able to enact in the TV bands. The database has access to information about the secondary density, and even the locations of individual secondary devices. This information can be used to tailor the no-talk radius or the per-node secondary power to the current locations of secondary interferers. The current rules use no local information, instead choosing a no-talk radius and the node power that will apply regardless of secondary location realization. The more local information is used, the better the database can protect the primaries and the more spectrum holes are available for use by secondaries.

A note on jails and fines

This thesis is built on an in-kind punishment scheme using spectrum jails. But the traditional sanction approach considered in the literature is a fine. Chapter 7 uses an abstract mathematical model to highlight the difference between the two and argue that in-kind sanctions are the right choice for spectrum.

Two major differences are explored: levels of uncertainty and the ability to treat a fine as a price.

In spectrum, because all radios are trying to transmit information across a wireless medium, there are a bounded set of quality of service parameters that can either be estimated or at least bounded. As is shown through the rest of the thesis, a sanction can then be built that deters harmful interference for a wide range of QOS cost functions. On the other hand, fines require understanding the monetary value of a transmission, which may depend on the QOS, the price the users are willing to pay, the business practices and pricing model of the company, etc. It is much harder to estimate the correct fine, and any estimate will have a high impact on any secondary wrongfully convicted.

Further, rational secondaries cause harmful interference when they are better able to achieve their QOS goals through interference instead of correct operation. An in-kind punishment removes the QOS benefit of causing harmful interference. A fine just makes that higher QOS more expensive. If interference is considered a certainly undesirable act that needs to be deterred, an in-kind punishment is best. Otherwise, the higher cost is just letting those secondaries who value their transmissions more highly to break the rules.

Table 0.1: Results of this thesis

Chapter	Results
1: A new model of trust for spectrum access	Presents the background and argument for spectrum jails

Table 0.1 – continued from previous page

Chapter	Results
2: Basic spectrum jails with a single secondary	<p data-bbox="537 394 1149 430">Shows the basic results of using a spectrum jail:</p> <ul data-bbox="586 430 1403 1003" style="list-style-type: none"> <li data-bbox="586 430 1403 569">• Trust can be guaranteed at certification time, regardless of the secondary technology. Certification is based on respecting the jail commands, not on the ability to correctly find spectrum holes. <li data-bbox="586 569 1403 716">• Performance will improve as technology gets better – the regulatory overhead of jails will go down and the protection for the primary will get better as spectrum holes become easier to find and the catching system improves its selectivity. <li data-bbox="586 716 1403 894">• New architectural requirements: each secondary must have access to a home band or unlicensed band where it does not have to search for a primary. The ability to deny the secondary access to its home band is vital to extend trust to any secondary device. <li data-bbox="586 894 1403 1003">• Regulator decisions: if there must be a fixed jail sentence, rarely active primaries cannot be guaranteed protection. The regulator must choose the lowest protected activity level.
3: Spectrum jails implemented through databases	<p data-bbox="537 1041 1403 1142">Extends the basic model to cover the database implementation proposed in the introductory chapter. This multiband model also covers wideband secondaries that encounter multiple primaries.</p> <ul data-bbox="586 1142 1403 1612" style="list-style-type: none"> <li data-bbox="586 1142 1403 1213">• The sanction must be adapted to account for the greater temptation of multiple bands. <li data-bbox="586 1213 1403 1325">• Once the adaptation has been done, the same properties of the original model still hold: trust first and performance later is still possible! <li data-bbox="586 1325 1403 1612">• If using a database to directly reallocate spectrum holes, the database cannot effectively recover spectrum holes if it cannot operate at the same time scale as the spectrum hole. If instead the database distributes conditional operation tokens and implements spectrum jails, spectrum holes can be recovered by any secondary technically capable of finding them. Even a slower time scale database can enable recovery of small spectrum holes.

Table 0.1 – continued from previous page

Chapter	Results
4: Universal jails	<p>Extends the model to a wider range of secondary cost functions and refines the model of harm.</p> <ul style="list-style-type: none"> • General utility functions: <ul style="list-style-type: none"> • The jail sentence must affect the QOS of the secondary device to be effective. The case shown here is a secondary that cares only about energy usage. Once an energy-wasting sanction is added, all of the same results from the previous two chapters still hold. • Because all radios are transmitting information in a wireless medium, there are only a limited range of QOS desires. This thesis covers energy and delay constraints, but does not handle sensitivity to timing of packets. There is hope of creating a sanction that works for <i>all</i> types of secondary devices. • A more refined model of harm: <ul style="list-style-type: none"> • If the secondary does not cause much harm when cheating, it is very difficult to catch. Therefore, a higher sanction must be in place to deter such cheating. • The secondary may cause harm when it thinks it is being honest (think of the secondary’s sensor mis-detecting the primary). The sanction must also be set such that honest but harmful secondaries choose to operate in their home bands instead of in the primary’s band. • Once the sanctions have been adapted, the trust first and performance later properties still hold. • Regulator decision: the regulator must decide the maximum “acceptable level of harm.” Think of this like an interference temperature. Once this decision has been made, a primary is guaranteed to be protected from any harm above this level as long as it is active often enough.

Table 0.1 – continued from previous page

Chapter	Results
5: The role of the primary	<p data-bbox="537 394 1403 464">If the next generation of primary can be designed to work with the spectrum jail system, what role should it take in its own protection?</p> <ul style="list-style-type: none"> <li data-bbox="586 468 1403 537">• Primaries will report harmful interference to protect themselves as long as the cost of reporting is low enough. <li data-bbox="586 541 1403 1073">• Can primaries be trusted to be responsible? <ul style="list-style-type: none"> <li data-bbox="651 575 1403 716">• If the cost of reporting is too low, the primary will report interference even when its packet was correctly received. The hope is to use these reports to kick the secondary out of its band. <li data-bbox="651 720 1403 821">• The primary will also transmit more than it needs in order to keep the secondary honest and then kick the secondary out of the band. <li data-bbox="651 825 1403 926">• The desired primary behavior is to use extra transmissions because these transmissions can be sub-leased to “band-sitting” secondaries that coexist well with the primary. <li data-bbox="651 930 1403 1073">• The cost of reporting can be used to control whether the primary uses false reports or extra transmissions to kick out the secondary. The exploration of this is begun in this thesis, but largely left to future work. <li data-bbox="586 1077 1403 1218">• Performance: as secondary technology improves, the primary has less incentive to lie or proactively find a band-sitter. As technology becomes perfect, the secondary and primary will peacefully coexist. <li data-bbox="586 1222 1403 1398">• Regulator decisions: The regulator must choose the cost of reporting to direct primary behavior. It must be low enough to allow correct reports, but high enough to deter false accusations. This choice also implicitly defines what constitutes harmful interference. <li data-bbox="586 1402 1403 1503">• Architectural requirements: the cost of reporting could possibly be controlled by having mini-jails for primaries. This is left to future work.

Table 0.1 – continued from previous page

Chapter	Results
6: Sharing risk amongst multiple secondaries	<p>Analyzes the distribution of aggregate interference from randomly-placed secondary nodes to understand how to build rules to account for multiple secondaries. Also introduces a way forward in adapting spectrum jails to handle multiple secondaries.</p> <ul style="list-style-type: none"> • Distribution of interference: near vs far field <ul style="list-style-type: none"> • As the density of secondary interferers increases (relative to the distance of the interferer located closest to the primary), the distribution of aggregate interference changes. • With very sparse interferers, the distribution is dominated by the interference of the closest, and has a heavy tail. • With dense interferers, the distribution resembles a Gaussian. • Designing rules that require the secondaries to be a certain distance from the primary must take into account the type of distribution it is facing. Urban and rural environments behave very differently. • Designing rules using database information – if the database can use its information about secondary locations, it can better protect the primary and better recover spectrum holes. <ul style="list-style-type: none"> • If the no-talk radius must be set at certification time with a fixed maximum per-node power (the current FCC rules), the rule will be very conservative and push secondaries much farther away than they need to be, especially for rural environments. Even then, the protection offered is probabilistic, promising only a small probability of too much interference because of poor secondary placement. • The database has access to secondary density information. It is possible to get better secondary performance, while offering the same probabilistic primary protection, by changing the power based on the density of interferers. • The database also has access to secondary location information. With this information, the database can guarantee too much interference will never result from poor secondary placement. At the same time, the database can give the secondary as much access to spectrum holes as possible. This section explores this through adapting either the no-talk radius or the power density. Optimizing over both is left to future work. • A first toy model for jails with multiple secondaries <ul style="list-style-type: none"> • The model from previous chapters could be extended using queuing theory. A simple example is shown here, but most of this adaptation is left to future work.

Table 0.1 – continued from previous page

Chapter	Results
7: Jails vs fines	<p>The thesis is built on in-kind punishment, but fines are the traditional tool for imposing sanctions. This chapter develops simple mathematical models to highlight the operational differences between the two approaches.</p> <ul style="list-style-type: none"> • With wireless, QOS can be bounded. However, the translation from the QOS value of a cheating transmission to a monetary value is very uncertain. In such a case, in-kind punishments will have much smaller overhead. • Jails and fines are fundamentally different things – one affects the achievable QOS while the other affects only the price at which that QOS can be obtained. When the QOS can be directly affected (like with spectrum jails), for most situations, especially given uncertainty over the QOS or price constraints, in-kind punishments are the better option.
8: Conclusion and Future Work	

This thesis is dedicated to my family

Contents

List of Figures	xiv
-----------------	-----

List of Tables	xxv
----------------	-----

1 A new model of trust for spectrum access	1
1.1 The cognitive radio problem	1
1.1.1 Different perspectives on “more dynamic allocation”	3
1.1.2 The TV whitespace compromise	7
1.1.3 The near future of spectrum regulation	8
1.2 The need for, and properties of, a technical regulatory layer	9
1.2.1 The goals of spectrum regulation	11
1.2.2 Moving regulation activities closer to runtime	12
1.3 The case for spectrum jails	18
1.3.1 In the future, trust through certification alone will not be good enough . . .	18
1.3.2 Goals and design philosophy of the certification/enforcement system	20
1.3.3 Assumed traits of secondary devices	21
1.3.4 Sanctions for radios	22
1.3.5 Putting it all together: spectrum jails implemented through databases	25
1.3.6 A note on evolvability	25
2 Basic spectrum jails with a single secondary	27
2.1 The model	28
2.2 Basic operation	32
2.2.1 Guaranteeing trust with a home band	36
2.3 Performance improvement with better technology	41
2.3.1 Performance metric definitions	41
2.3.2 Exploring the metrics	44
2.3.3 Improving performance in the future	49
2.4 Concluding remarks	52
2.4.1 What we have learned so far	52
2.4.2 Future work	52
3 Spectrum jails implemented through databases	56
3.1 The model and necessary sanction	57
3.2 Performance with improving technology	63

3.2.1	Recovering spectrum holes with fixed N, M	63
3.2.2	Optimal expansion	67
3.3	Concluding remarks	70
3.3.1	What we have learned so far	70
3.3.2	Future work – adapting jails to expansion	72
4	Universal jails	73
4.1	The need for different sanctions	73
4.2	Energy and mixed users	77
4.2.1	The Model	77
4.2.2	Extending trust to energy and mixed users	78
4.2.3	Better performance with better technology	80
4.2.4	Database implementation for energy and mixed users	86
4.3	General delay, energy, and mixed cost functions	88
4.3.1	Different kinds of mixed users	88
4.3.2	Margins for other unmodeled phenomena	89
4.4	A more refined model of harm	90
4.5	Concluding remarks	94
4.5.1	What we have learned so far	94
4.5.2	Future work – other kinds of secondaries	95
5	The role of the primary	97
5.1	The Model	99
5.2	Guaranteeing trust	103
5.2.1	Trusted protection: primary can only raise red flags	103
5.2.2	Trusting the primary to act responsibly	109
5.2.3	Setting C_{rf} to deter crying wolf	117
5.3	Improving Performance	119
5.4	Choosing to split the band	119
5.5	Improvement with technology	121
5.6	Concluding remarks	123
5.6.1	What we have learned so far	123
5.6.2	Future work	124
6	Sharing risk amongst multiple secondaries	126
6.1	The models	127
6.1.1	Per-node interference power model	128
6.1.2	Node placement models	129
6.2	Near vs far field: understanding the distribution of aggregate interference	130
6.2.1	Calculating the interference power	131
6.2.2	Near vs far field	140
6.3	Dealing with placement risk	143
6.3.1	r_n fixed at certification time	144
6.3.2	r_n set at certification time; power set at runtime	147
6.3.3	r_n set at runtime; power density set at certification time	148

6.3.4	Performance evaluation	150
6.4	Concluding remarks	151
6.4.1	What we have learned so far	151
6.4.2	Future work	153
6.4.3	A first toy model for enforcement	154
6.4.4	Observations and the way forward	157
7	Jails vs fines	159
7.1	The general problem	160
7.2	A data-processing toy	163
7.3	QOS vs prices toy	164
7.3.1	The model	165
7.3.2	Without uncertainty	167
7.3.3	With uncertainty	171
7.4	Concluding remarks	178
7.4.1	What we have learned so far	178
7.4.2	Future work	178
8	Concluding remarks	179
8.1	What has been accomplished	180
8.2	Where to go from here: moving spectrum jails closer to reality	183
	Bibliography	186

List of Figures

1.1	The regulatory “control system.” Engineers tend to think of only the control system of the wireless device trying to access local spectrum with help from a database or other entities. Regulators, however, are really trying to use certification and enforcement to control any <i>temptation</i> the engineers may have to cause interference for the benefit of their own networks.	10
1.2	The command-and-control system. Although the network is running at the network operation timescale, all regulatory operations happen at much slower human time scales. In particular, any enforcement events require a network administrator to complain, which triggers an FCC investigation and perhaps a cease-and-desist letter. If all networks have their own private frequency band, this slow intervention time scale is sufficient to control interference.	13
1.3	With the TV Whitespace ruling, cognitive devices are able to opportunistically use spectrum holes, with help from a database. However, intervention to turn off malfunctioning devices, for example, still goes through the same complaint, investigation, cease-and-desist cycle as with command-and-control. This cycle is far too slow to effectively deal with problems.	15
1.4	If the database has access to identity and a kill-switch, the new “spectrum manager” will actually be able to effectively protect primaries. These same technical primitives can also be used to enact a run-time enforcement scheme.	17
1.5	How much spectrum a database can recover depends on its time scale relative to the primary. Even if the primary is only using the band 25% of the time, the database cannot recover any spectrum holes if the primary usage is spread out over the database time slots.	18
1.6	Even if the database can predict perfectly which spectrum holes a primary will be using, it can only recover the time slot if the <i>whole</i> time slot is empty. So, if the primary uses 30% of the band, and the database token includes several primary time steps, this mismatch will decrease the recoverable spectrum. A single order of magnitude difference in time scales will effectively eliminate recoverable spectrum.	19
1.7	The utility a secondary can derive from cheating is a boost to the quality of service it can provide. Because the service is transmitting data, this utility really only lies along three dimensions: throughput, energy usage, and “sensitivity to burstiness.” This last one is capturing the sensitivity to the timing of packets, so voice communication will have a different utility along this axis than transmitting a large file. This thesis covers secondaries with throughput and energy-based cost functions.	23

2.1	Flowchart of events governing the operation and punishment of the secondary. . . .	29
2.2	Markov chain illustration for the jail system. The primary follows a Markov chain to turn on and off. When the primary is not transmitting, the secondary legally transmits. When the primary is transmitting, the secondary chooses whether to wait, or cheat by transmitting anyway. The secondary is sent to jail with probability P_{catch} if cheating and P_{wrong} if not cheating. Once in jail, the secondary stays there for an amount of time determined by P_{pen}	30
2.3	If the primary is <i>iid</i> , then the Markov chain can be simplified to two states: when the secondary is out of jail and when it is in jail. The movement to jail is determined by whether the primary is on and the secondary is cheating. Moving out of jail still depends only on P_{pen}	31
2.4	Example costs for the secondary given the level of cheating, with two different values of the sanction parameter P_{pen} . The cost is always minimized by $P_{cheat} = 0$ or 1. This is true in general: the secondary will never choose to cheat only part of the time. 33	33
2.5	The required average sanction (amount of time in jail per offense) to deter cheating for different values of the cost of sensing. Different lines indicate different probabilities that the primary will be active. As the cost of sensing rises, the necessary sanction rises without bound, indicating there does not exist a sanction that will work for all costs of sensing.	34
2.6	The required average jail time per offense to deter cheating vs the probability of the primary being active. Different lines represent different costs of sensing. When the primary is mostly active, there is little opportunity to transmit, so the secondary is more tempted to cheat. When the primary is rarely active, the secondary would rather go to jail every once in a while than sense at every time step.	34
2.7	With a home band, the operation in the primary band is the same as before. In the home band, the secondary does not have to respect a primary's priority, and so its only option is legal transmission. The secondary chooses its operation mode: operate just in the primary band, just in the home band, or in both. The jail is a global jail – the secondary loses transmission ability in all bands, including the home band. It is assumed that if operating only in the home band, the secondary is not subject to jail.	36
2.8	Required sanctions vs the cost of sensing. The bottom black line is the sanction required to deter cheating. The red line is the sanction required to make the secondary move to the home band. To deter cheating, the sanction must be the minimum of these two lines (dotted green line). Notice that it is possible to choose a level of sanction that will work for all values of the cost of sensing, and it does not go to an infinite sanction as the cost of sensing goes to 1. Even if the secondary is operating in the primary band, its behavior is modified by being able to transmit in the home band as well. When the cost of sensing is very high, the secondary still has incentive to not cheat in order to protect its operation in the home band.	38

- 2.9 The sanction required vs the probability the primary is active. The bottom black line is the sanction required to deter cheating. The red line is the sanction to move the secondary to the home band. The required sanction is the minimum of these two lines (green dotted line). Notice that there is no universal sanction possible. Regulators must define a minimum protected P_{tx} . Then, there exists a sanction that protects all primaries operational for more than P_{txMin} of the time. In this plot, there are two points to compare when choosing the sanction: the sanction needed at the minimum P_{tx} , and the sanction needed when the secondary will switch to the home band at high P_{tx} . This is a result of the chosen cost of sensing. Taking the maximum of these points over all costs of sensing leaves only the sanction needed at P_{txMin} 40
- 2.10 The primary transmission probability below which cheating actually occurs, given P_{txMin} and cost of sensing. If it is possible to improve sensing technology, the primary will be better protected. 41
- 2.11 The bounding curve from Fig. 2.10 for different values of P_{wrong} . Better wrongful conviction rates will also produce better realized protection for the primary. 42
- 2.12 Overhead metric values for different values of the parameters. The different plots use the indicated P_{txMin} and P_{tx} . Overhead is shown as a function of C_{DS} . Black solid lines are secondary overhead, green dotted lines are primary interference experienced, red dashed lines are holes due to jail, magenta solid lines are holes due to retreat to the home band. Notice that the primary is protected for only some values of C_{DS} when $P_{tx} < P_{txMin}$. Also, the lower the P_{txMin} , the higher the sanction. So, the secondary will move to the home band for smaller levels of C_{DS} 44
- 2.13 Overhead metrics vs P_{tx} . Different plots are for the indicated values of P_{txMin} , C_{DS} , and all plots are against P_{tx} . Black solid lines are secondary overhead, green dotted lines are primary interference experienced, red dashed lines are holes due to jail, magenta solid lines are holes due to retreat to the home band. Again here, the secondary moves to the home band for lower values of P_{tx} when P_{txMin} and therefore the sanction are higher. The primary is not protected when $P_{tx} < P_{txMin}$, but has fewer collisions for higher P_{tx} because the secondary is spending more time in jail. Secondary overhead gets worse for higher P_{tx} for the same reason – it is spending more time wrongfully in jail. 45
- 2.14 Summary overhead metrics against P_{txMin} . The values are averaged over the possible realizations of primary and secondary, so over P_{tx} and C_{DS} . The two plots have different values of P_{wrong} , .1, .01. All measures of overhead get better with smaller P_{wrong} 47
- 2.15 If the regulator chooses to change the sanction to balance the needs of the different parties, it can choose it based on the Pareto-like curve here. This is the value of the probability of primary collision vs the total spectrum holes. The curve represents different values of P_{txMin} . The circles are different choices of balancing – red is the min sum solution, green fixes the total collisions, and blue fixes the total spectrum holes. Different lines are different values of P_{wrong} 48

2.16	Given the three possible fairness balances (min sum, fix collisions, fix total holes), this plot shows the chosen P_{txMin} , average collisions, and average spectrum holes against P_{wrong} . Notice the dependence of the optimal P_{txMin} on P_{wrong} . If fixing the total holes metric, the best P_{txMin} is highly dependent on P_{wrong} , and so using this approach would require the ability to change the sanction as the technology for catching cheaters changed.	48
2.17	Band usage for different performance parameters. Blue is used by the primary, green is honest use by the secondary, black shows cheating use by the secondary, and red shows collisions between primary and secondary. The black area is curved because of time spent in jail. The green is hindered by wrongful convictions and sensing time. When there is no green filling in holes above the blue line, the secondary has retreated to the home band. As P_{wrong} and C_{DS} improve (get smaller), there are fewer collisions, the honest use more effectively fills the band, and the secondary retreats to the home band only for larger values of P_{tx}	50
2.18	The limit on the wrongful conviction rate in the home band or deterrence to still be possible, as a function of the size of the home band and P_{tx}	54
3.1	Assume a database can only issue one operation token every M time steps, and the primary is modeled as having an <i>iid</i> Bernoulli $P_{tx} = 0.3$ probability of transmitting in any time step. This is the band utilization (the primary uses 30% of the slots, the secondary recovers as much as possible). If the mismatch is high, even though spectrum is idle 70% of the time, very few holes can actually be recovered.	57
3.2	Flowchart of events governing the operation of the secondary with a database-token implementation of spectrum jails. The secondary decides whether it will operate in the primary bands; "In band" is the state variable indicating this choice. The database issues operation tokens when the secondary has no jail sentences to serve and jail tokens during jail sentences. The secondary can only transmit when it has an active operation token, which expires after M time steps. Sanctions are issued to secondaries caught causing interference during their operation periods.	58
3.3	A jail system run through a database with operation and jail tokens can be modeled as the multiband version of the Markov Chain model, as illustrated here.	59
3.4	Spectrum holes recovered by spectrum jails vs a database alone. If both sensing cost and wrongful convictions rates are high, spectrum jails will recover about the same amount of spectrum holes as a centralized database solution that knows ahead of time when the primary will transmit. However, as technology improves, spectrum jails can recover all spectrum holes, regardless of the speed of the database.	65
3.5	The necessary sanction goes up as the expansion rises because there is more opportunity and therefore more temptation to cheat as the secondary operates in more primary bands.	68
3.6	Overhead metrics as a function of different levels of expansion. As the expansion grows, so do the total spectrum holes because the wrongful conviction rate and the length of sanctions are scaling with the expansion. The collisions are decreasing because the secondary is spending more time in jail.	69

3.7	Although the <i>overhead</i> rises with expansion, the secondary cost tells a different story. The secondary will expand depending on the wrongful conviction rate, expanding more when the wrongful conviction is lowered.	70
3.8	This is the optimal expansion for a delay-constrained user over different values of the cost of sensing and the probability of primary transmission. If there is a lot of opportunity (P_{tx} small), the realized number of wrongful convictions will be small, so the desired expansion is very high.	71
3.9	Averaging over the range of sensing cost and primary transmission probabilities, this is the average expansion vs average spectrum holes per slot as a function of the P_{wrong} . Notice that the optimal average amount of spectrum holes, from the secondary perspective, is around 25%. The regulator can decrease this by restricting the potential secondary expansion.	71
4.1	This is the flowchart of spectrum jails without the infrastructure of the databases. This is identical to Fig. 2.1 in Chapter 2 in all but the inclusion of the home band — the secondary decides at the beginning whether it will operate only in its home band or in the primary bands as well.	74
4.2	The sanction and resulting percentage of time spent in jail for honest users when only delay-constrained secondaries are present.	75
4.3	If energy-constrained users are operational as well, using just a delay-oriented jail to deter cheating requires a potentially much higher sanction depending on the natural energy cost of sitting in jail. This causes much more wrongful jail time for honest users.	76
4.4	The sanction required to deter cheating in energy-constrained users as a function of the natural energy cost of jail relative to the energy cost of sensing.	76
4.5	Illustration of the Markov model used in this section. The secondary has the option to transmit in a home band along with or instead of the primary band. The transition probabilities and costs for different states are shown.	77
4.6	Summary metrics for different levels of wrongful conviction. Energy and delay-constrained secondaries have different lines for their overheads; mixed users will have the appropriate mix of the two overheads. The primary and spectrum hole lines are all averaged between energy and delay-constrained users. All metrics get better as wrongful conviction rates go down.	82
4.7	Overhead metrics as a function of the size of the margin for unmodeled phenomena Δ . The secondary overhead and spectrum hole metrics only rise sharply for high Δ , suggesting some robustness to additional margin. The primary collision metric gets better because higher Δ implies a higher sanction and therefore a larger range of protection P_{tx}	90
4.8	The change in P_{wrong} required to keep the secondary overhead equivalent as Δ gets larger. If the regulator wants to increase protection against different kinds of utility functions while maintaining the same performance for the honest secondaries, it can pour money and effort into reducing the wrongful conviction rate.	91
4.9	Spectrum jails guarantee that the primary will be protected for all $P_{tx} > P_{txMin}$ and $\theta > \theta_{max}$ where the θ includes any honest or cheating harm from the secondary. . . .	95

5.1	Spectrum jails guarantee that the primary will be protected for all $P_{tx} > P_{txMin}$ and $\theta > \theta_{max}$ where the θ includes any honest or cheating harm from the secondary. . . .	98
5.2	Cost functions for a primary that can raise red flags, but cannot raise their P_{tx} or cry wolf. The bottom plot is cost by P_{rf} ; the top plot is cost by \tilde{P}_{tx} optimized over P_{rf} . The red flag is correctly used to stop the secondary from cheating and to kick the secondary out of the band when \tilde{P}_{tx} is high enough.	104
5.3	This is a guide for reading Figs. 5.4 and 5.9. Above both dashed lines in the pair, the secondary is transmitting only in its home band. Between the two lines, the secondary is honestly using the cognitive band. Below the two lines, the secondary is cheating in the cognitive band.	104
5.4	Comparing the regions from Fig. 5.3 for primaries with different action spaces, using Parameter Set (A) from Table 5.2. The honest secondary (black dashed line) always raises a red flag, and is honest with the other parameters. The primary that can choose its probability of raising a red flag will follow the honest primary as long as the cost of red flag is low enough. The primary that can choose both P_{rf} and P_{cw} will use these to make the secondary not cheat and leave the band as soon as possible. Likewise, a primary that can choose P_{tx} as well will use this tool to kick out the secondary early as long as the cost of red flag is low. Notice that if the primary can choose only P_{rf}, P_{tx} , it coincides with the full game lines for high C_{rf} . This indicates that it is possible to choose C_{rf} such that the primary will prefer using P_{tx} over P_{cw} to control the secondary.	105
5.5	Maximum acceptable P_{rf} the primary is willing to use to get a γ decrease in dropped packets. If the cost of reporting is very low, the primary is willing to report more often to get the secondary to stop cheating or to exit the band.	107
5.6	Illustration of the choices the primary makes when it can report both dropped and received packets as interference. In the bottom plot, the primary can use P_{cw} to kick out the secondary, but it cannot use this tool to make the secondary stop cheating. The middle plot shows the choice over P_{rf} with the optimal P_{cw} . With higher P_{rf} , the secondary stops cheating and then in conjunction with higher P_{cw} , it will leave the band. The top plot shows the overall cost, and the optimal decision given an initial required value \tilde{P}_{tx}	110
5.7	The maximum P_{cw} a primary will use to kick the secondary out of the band given a cost of reporting, and a fixed probability of reporting interference. As the cost of reporting falls, the acceptable P_{cw} rises. However, even if the cost of reporting is high, a big enough change in dropped packets (γ) will still provoke the primary to cry wolf.	111
5.8	The costs for a primary that can operate in the full action space. P_{cw} is chosen in the bottom plot, noting that higher P_{cw} will cause a secondary to start cheating before it leaves the band. The effect of different P_{tx} , optimized over P_{cw} , is in the second plot from the bottom. P_{rf} , optimized over P_{tx} and P_{cw} , is in the plot second from the top. Finally, the summary plot is at the top. Again, because of the poor coexistence parameters, the primary uses all of its available tools to stop the secondary from cheating and to kick it out of the band as soon as it is possible (and cost effective). .	114

- 5.9 When better coexistence parameters are used (Parameter Set (B) from Table 5.2), the secondary is easily deterred from cheating and the primary does not try as hard to kick the secondary out of the band. 115
- 5.10 The maximum factor increase in P_{tx} the primary is willing to use to produce a γ drop in packet drop rate. This assumes C_{rf} is high enough that it is better to use P_{tx} to control the secondary, rather than P_{cw} 116
- 5.11 For an example secondary, and different values of $\tilde{P}_{wrong,other}$, the solid lines show the tradeoff between P_{tx} and P_{cw} above which the secondary will move to operating only in its home band. The dashed line shows the minimum P_{tx} required achieve a switch. Note that the original \tilde{P}_{tx} does not affect the line – to kick out the secondary, the primary must raise its transmission probability to this level, regardless of what activity level is actually desired. 117
- 5.12 These are example primary costs for different levels of C_{rf} . The plot is done against P_{cw} , but P_{tx} is also changing to trace out the line in Fig 5.11 that will actually make the secondary move to the home band. For this example, the $\tilde{P}_{wrong,other} = 0.01$ line is used. The cost of reporting must be high enough that the lowest cost is at $P_{cw} = 0$. In this example, that level is $C_{rf} > 0.359$ 118
- 5.13 The C_{rf} needed to make the primary want to sub-lease its band instead of crying wolf is dependent on the value of $\tilde{P}_{wrong,other}$. This means that it is not possible to define a C_{rf} that will deter crying wolf for all parameter combinations. The best the regulator can do is set C_{rf} at the highest value that does not violate the trust requirement in Thm. 22. Then, when the identity system is good enough, the correct primary behavior will happen. 119
- 5.14 If the primary is allowed to split its band and give up half the time to unlicensed use, it is left with a piece of spectrum that is much more easily defensible than its original allocation. So, it more easily stops the secondary from cheating. It also is able to share its band more easily with the secondary, so the primary does not try to kick out the secondary. 120
- 5.15 When the primary is allowed to split its band, it does not have to work as hard to protect itself. So, it has lower cost. 120
- 5.16 Performance metrics against P_{pen} assuming the secondary is purely delay-constrained ($k = 1$). The plots use the Monte Carlo method to average the metrics over the types of secondaries (parameterized by C_{DS}) and the types of primaries (parameterized by P_{tx}). In general, the most important factor for improving performance is to reduce the level of honest harm, θ_{honest} 122
- 6.1 Model setup: there is a TV transmitter at the origin, receivers within the protected radius r_p , no transmitters between r_p and r_n , and whitespace devices outside r_n placed according to a specified placement model. 128
- 6.2 2D deterministic placement model: the first ring of interferers is located $(m + \delta)d$ away from the center, and each subsequent ring is d away from the previous ring. The number of nodes in each ring is chosen so that there is one node per d^2 area. . . 129

6.3	Keeping the power density and the no-talk radius constant, this is the interference power for the deterministic placement models as m grows (nodes get closer together). The closest grid has $\delta = 0$ and the farthest has $\delta = 1$. As the nodes are packed more closely, the interference from all deterministic models converges to the sea model.	132
6.4	Example aggregate interference distributions for the 2D dithered grid model. If d is small (nodes are packed tightly), the distribution resembles a Gaussian. If d is large (nodes are spread out), then the distribution is heavy-tailed.	133
6.5	Approximate distributions for the dithered grid placement model. If the nodes are placed close together, the right approximation is a Gaussian distribution. If the nodes are spread out, the right approximation is the interference caused by a single randomly-placed interferer.	135
6.6	Example aggregate interference distributions for the Poisson placement model with no fading. As λ gets larger (nodes are more dense), this distribution also closely resembles a Gaussian.	136
6.7	Approximate distributions for the interference caused with the Poisson placement model. For high λ (nodes are dense), the approximation is a Gaussian. If λ is small, the right approximation is the distribution of the interference coming from a single interferer.	137
6.8	Distribution of the interference coming from nodes placed according to the Poisson model with fading. Even with fading, the distribution is more Gaussian-like for higher λ , and more heavy-tailed for smaller λ	139
6.9	The variational distance between the interference distribution and the appropriate Gaussian without fading, $\alpha = 2.5$. Notice that there are lines of equal variational distance to a Gaussian.	141
6.10	Variational distance by $r_n^2 \lambda$ with and without fading and with $\alpha = 2.5$. Notice that fading does not change the speed at which the distribution converges to Gaussian. This is because fading does not change <i>how many</i> nodes are contributing significantly to the interference.	142
6.11	Variational distance to Gaussian for different values of α . As with fading, the path loss exponent does not change how many more nodes are contributing significantly to the interference as $r_n^2 \lambda$ gets large, so it does not change how quickly the distribution converges to Gaussian.	143
6.12	The required r_n to ensure that the probability of total interference being too high is less than γ for different γ , and both the Poisson distribution and Gaussian approximation. This is for fixed per node power, and no-talk radius fixed at certification time. As nodes get more dense, the required no-talk radius must get larger. Therefore, with a fixed per-node power limit, and no control on the node density, there is no way to guarantee protection for primary users.	145
6.13	If the no-talk radius and the maximum <i>power density</i> is fixed at certification time, the database can allow secondaries to have higher per node power at runtime. However, as the node density gets smaller, the per node allowed power goes up, which means the required no-talk radius must be larger for smaller node densities. Again, no guarantee can be offered for primary users.	146

6.14	The required no-talk radius from Fig. 6.13 for different levels of protection γ . As γ gets smaller, the margin over the no-talk radius required for the sea gets larger. The difference from the Gaussian distribution also gets larger, for the same node density.	146
6.15	The per-node power cannot be increased indefinitely. If there is a fixed maximum per-node power, as well as a maximum power density and no-talk radius fixed at certification time, this is the required no-talk radius by node density. In this case, a protection guarantee can be offered to the primary user: set the no-talk radius at the highest required point, where the maximum per-node power and the per-node power determined by the power density are equivalent.	147
6.16	If the per-node power can be set at runtime based on the realized node placement (information that the database has from secondary access requests), the CDF of the achieved power density is shown here. The sea would achieve a power density of 5. Sparse secondary nodes (low λ) would see the most gain from favorable placements allowing much higher power densities.	148
6.17	If the power density is set at certification time, but the no-talk radius can be set at runtime based on the realization of the node placement, this is the CDF of the achieved no-talk radius. The sea would have a no-talk radius of 10. Sparse nodes see a higher no-talk radius because the power density rule allows these nodes to transmit at very high powers.	149
6.18	The performance as a function of the node density. Different plots have a different weighting, w , for the desire to transmit closer to the primary. The rules that allow modification based on the realization perform significantly better than those that have most parameters fixed at certification time. The rule with per-node power and no-talk radius fixed at certification time does well at high λ only because it is no longer offering sufficient protection to the primary. If it is important to have secondaries close to the primary, it is better to change the no-talk radius at runtime. If distance is less of an issue, it is better to change the power density at runtime. The optimal rule will be some combination of these two.	152
6.19	Placement model for the queuing example to extend spectrum jails. The primary receiver is in the center, and all secondaries are located a fixed distance away so that they all cause the same amount of interference.	154
6.20	The toy model to extend spectrum jails using queues. Secondaries enter the first buffer when they have something to transmit. As a group, they decide how many interference servers they will use, where N is the tolerable amount of interference at the primary. When finished transmitting, a secondary will either leave the system or go to jail depending on how much total interference is being caused. The goal of the regulator is to choose the length of the jail sentence (or the speed of the jail server μ_{jail} to make it in the secondaries' best interest to not cause too much interference.	155

7.1	An abstract idea of the considerations of choosing in-kind vs monetary sanctions. There exists an actual valuation on the crime, but there is no way to estimate or attack that actual valuation. So, sanctions could be applied at different layers of abstraction. For spectrum sharing, the in-kind abstraction includes all technical inputs that determine the quality of service available, and the sanction looks like spectrum jails. The monetary abstraction layer includes all business and personal inputs that affect the monetary value of cheating. It also includes the available quality of service. With more inputs, the uncertainty increases, so while it may be possible to understand the range of in-kind valuations, it may not be possible to reliably quantify the range of monetary valuations.	161
7.2	A model of how the type of sanction should be chosen. The sanction must be chosen at the regulation time scale with some noisy understanding of the monetary and in-kind inputs that are determining the utility function. These sanctions are then applied at runtime to any rightfully or wrongfully convicted agent. The agent decides whether to cheat based on the applied sanction, and its base utility while cheating or being honest.	162
7.3	A general model to understand operational differences between jails and fines. A provider can offer a service of a particular quality for a particular price based on whether or not it cheats. The user will buy the service as long as it is below the user's QOS-price tradoff. The sanction will either reduce the QOS available or raise the price at which that service is offered in order to make cheating not worthwhile. .	165
7.4	A simple version of the general model in Fig. 7.3 in which all relationships are assumed to be linear.	165
7.5	Case 1: This shows the options for sanctions given the depicted relationship between lines. The in-kind sanction can push the cheating line back to be worse than the honest line in QOS for the same price. Alternatively, the fine sanction can push the cheating line up so that any given QOS costs more to provide through cheating than honestly.	168
7.6	Case 2, version 1: The options for the sanction given this configuration either uses jails to push the cheating line so that it cannot provide a better QOS than the honest line. Or, fines can make the price of that high level of service so large that the user is unwilling to pay it.	169
7.7	Case 2, version 2: Given this configuration of lines, the only choice is to use either sanction to push the cheating line to be so expensive for a given QOS that the user is unwilling to pay it.	169
7.8	Case 3: Given this configuration, the in-kind sanction can push the cheating line back to the honest line, or the fine sanction can make cheating to expensive for the user to willing pay.	170
7.9	A guide to the parameter regions when jails or fines have better overhead. The line represents the points where the overheads are equivalent. To the left of the line, jails have better overhead. To the right of the line, fines have better overhead. δ is capturing the wrongful conviction, with higher δ implying a higher wrongful conviction rate.	171

- 7.10 Regions where jails or fines have better overhead, for different δ . Higher δ implies higher wrongful conviction. The regions are as in Fig. 7.9. As δ gets smaller, jails are the better choice for a wider range of parameters. 172

List of Tables

0.1	Results of this thesis	iv
1.1	Timescales of regulation activities under different regulatory schemes. C&C is command-and-control.	13
2.1	Parameters definitions for the basic spectrum jails problem. The jail sentence is defined based only on those parameters known at certification time, and performance depends on those parameters determined at runtime.	32
4.1	Parameters definitions for the extended spectrum jails problem. The jail sentence is defined based on those parameters used at certification time, and the performance is determined by those parameters used at runtime.	79
5.1	Parameters for the game between primary and secondary: cert. is set at certification time and needed for trust, run. is set at runtime and needed for performance, reg. is known to the regulator, pri. is known to primary, sec. is known to the secondary.	102
5.2	Values of the parameters used for different examples of the game between primary and secondary.	103
6.1	Parameter definitions for aggregate interference from multiple secondaries.	128

Acknowledgments

During my stay at Berkeley there have been many many people who have offered help, support, and guidance along the way. Thank you to all of those people, and my apologies to anyone I miss.

First and foremost, I would like to thank my advisor, Prof. Anant Sahai. It is rare, I think, to be able to spend your graduate school years working on a problem which is at the same time important and yet completely undefined. Anant taught me how to tackle such a problem and gave me significant support and guidance through the years to keep this (somewhat unbounded) research problem manageable. I very much appreciate this opportunity to learn and apply so many different tool sets from both policy and technology to try to understand trust in cognitive radio. I also very much appreciate being able to spend conferences trying to convince people that there are good reasons to send radios to jail :)

More seriously, I would like to thank Anant for teaching me what it meant to be a researcher. Also, thank you to Anant for teaching me to be open minded about finding interesting problems and considering crazy solutions that are instructive if not yet practical. It has been a pleasure and an honor being his graduate student.

Thank you to my committee, Prof. Jean Walrand, Prof. Vern Paxson, and Prof. Brian Carver, for all their help and input along the way.

Thank you to Pulkit Grover, my friend, collaborator, and now husband, who has kept me grounded throughout this process. I would not have made it through grad school without his encouragement and his help to keep everything in perspective. Whether exploring caves, talking about life the universe and everything, or learning new mathematical tools for theoretical research, he helped make these years fun and rewarding. I look forward to many more years of exploring life together.

Thank you to Gireeja Ranade who has always been there for me, whether for laughter, commiseration, or a deep technical discussion. She taught me new ways of looking at the world and new ways to approach research. I will miss our cookie-breaks that were often the highlight of my afternoon.

Thank you to the senior grad students who were particularly helpful in teaching me how to survive grad school and continue to offer helpful advice on life post-grad school, including Krish Eswaran, Rahul Tandra, and Mubaraq Mishra.

Thank you to my collaborators, all of whom made late nights of frantic paper-writing actually manageable (and often actually fun): Anant Sahai, Pulkit Grover, Gireeja Ranade, Rahul Tandra, Mubaraq Mishra, Hari Palaiyanur, and Kate Harrison.

Thank you to Martin Weiss and all of the policy experts at DySpAN who taught me about the policy issues related to my thesis.

Thank you to Prof. Francois Baccelli for the inspiring discussions and getting me started with stochastic geometry.

Thank you to Sheila Humphreys and WISCE for their support and friendship.

Thank you to Prof. Martin Haenggi, my undergraduate research advisor at Notre Dame, who taught me what research was and started me on this path to a PhD.

Thank you to all of the Wireless Foundations students and faculty, past and present, for helping to make my grad school experience as rewarding as it was.

A big thank you to my family and friends for all their love and support through this long (and often arduous) process. I could not have finished without you!

Finally, this thesis could not have happened without the NSF which has supported this research through a Graduate Research Fellowship and grants CNS-0403427, CCF-0729122, CCF-0917212, CNS-0932410, CNS-1321155, and ECCS-1343398.

Chapter 1

A new model of trust for spectrum access

This thesis explores the use of *spectrum jails* to enable light-handed spectrum regulation. The goal is to build an enforcement mechanism that will allow the FCC¹ to trust that new spectrum access technologies will respect the priority of current users *without* directly certifying the operation of these new technologies. The path from conception to deployment of new technologies will therefore be easier and more transparent than the current path, enabling innovation.

Spectrum sharing has become a reality with TV whitespace regulation. However, this introductory chapter will argue that the current rules do not fully exploit the potential of cognitive radio. The problem is that although dynamic allocation has been embraced, the FCC retains the same trust model that it has used since the beginning of spectrum regulation. Moving forward, a new trust model will be required.

To understand the nature and requirements of the new trust model, and the enforcement and certification systems that enable it, this chapter casts the spectrum regulation problem as an abstract hierarchical control problem. This serves to show trust as the main regulatory problem and indicates the technical primitives that will be required to enable trust. These include identity for cognitive radios and a run-time kill switch.

The chapter closes by arguing that spectrum jails are a good way to use these technical primitives to achieve trust for new technologies.

1.1 The cognitive radio problem

Cognitive radio, or spectrum sharing, is a blanket term for the response to the problem of apparent spectrum scarcity.² Obviously, some kind of scarcity exists: the last available spectrum allocation was in the 700MHz band and was sold at auction in 2008 for about \$20 billion [2]. The

¹This thesis is written from the perspective of U.S. spectrum regulation. But the ideas will apply for national regulators in general.

²The term itself was coined by Mitola [1], and originally was used to mean a radio that could use its current context to determine its operation. The term could include everything from sensing the local spectrum for opportunity to transmit to recognizing a device's location based on auditory cues. Eventually, the term was limited to include just spectrum access because this was the part of intelligent action that required significant changes to hardware, software, law, and policy. All other kinds of intelligent customization could be implemented at will in software.

demand for wireless data continues to grow exponentially [3]. At the same time, the Government has issued the National Broadband Plan [4], which calls for 500MHz of spectrum in the next 10 years to supply broadband Internet access to the entire country. If no new spectrum allocations are available, how can these new demands for data be met?

The key is that the scarcity is in available *allocations*, not in available spectrum. On average only 5% of the usable spectrum is actually utilized [5] at any given location and time.

The real problem is a mismatch between the scale of use and the scale of allocation. Intuitively, TV station allocations should not extend into regions that the signals cannot reach [6]. Cell towers do not need as much bandwidth at night when the usage is sporadic [7].

If the inefficiency is so obvious, why was the scale mismatch created in the first place? When command-and-control (the original spectrum regulation scheme) was first devised, and spectrum was abundant, inefficiency stemmed from the natural response to deal with the fact that the wireless medium is messy. In space, signal propagation is affected by bouncing off of buildings [8], how easily different frequencies pass through different materials [9], and even the current state of the ionosphere [10]. In frequency, signals naturally spill out of their defined bands and cannot be sharply contained [11]. Only time can really be completely controlled by turning a switch on and off. But before the Global Positioning System [12], there was no global reference for time, and therefore no way to synchronize devices.

In the face of such uncertainty, and few services demanding spectrum, the sensible response was to define boxes in time, frequency, and space. The boxes were separated in frequency with enough guard room to account for any spill-over. Space was simple because the problem was to fit in TV and radio towers with fixed locations and power footprints. The boxes just needed to allow enough guard space between towers to separate footprints. Time could be split into day and night operation, but generally allocations were given to a single company for a period of several years. Different license holders needed to figure out how to deal with uncertainty like fading and multipath within their own networks, but they could ignore any uncertainty coming from other license holders.³

Separating different networks like this has a corresponding benefit of being extremely easy to enforce. The FCC requires certification before any wireless device can be deployed. If a network's usage right is contained within its own private block of spectrum, certification means checking that the device will respect the limits of its block. Inter-network interference events are rare, consisting mainly of individual infractions and malfunctioning devices. The FCC handles these by investigating complaints and issuing cease-and-desist letters or fines as appropriate [17, 18].

With command-and-control, inefficiency is unavoidable – companies require stable usage rights to deploy systems. For the technology at the time, this meant that a company's allocation box could not change over the lifetime of the allocation. So, the box was defined to account for the *peak* system usage. Off-peak usage would certainly lead to wasted spectrum. But at the time, the waste was acceptable.

Now, however, wireless usage and the economic potential of new systems is skyrocketing [3]. Wasting spectrum through command-and-control-style regulation is no longer acceptable.

The solution to the problem is obvious: make allocation faster and more dynamic in order to correctly match usage needs with available spectrum opportunities.

³For information on the history of this separation method, see [13–15]. The resulting FCC allocation chart is here [16].

Unfortunately, *how* to architect that dynamic spectrum allocation system is not obvious and is hotly debated. Even worse, agreeing on which problems are most important, or most difficult, changes depending on one’s background.

The next part illustrates the different perspectives and problems, paying close attention to how the different groups think about the enforcement problem.

1.1.1 Different perspectives on “more dynamic allocation”

In general, there are three groups working to solve the problem of dynamic allocation: engineers, policy experts, and regulators. There is extensive literature about the future of spectrum regulation from all three groups. This thesis cannot possibly go through the nuances of all of the important issues discussed in that literature. So, this section pulls out a few issues that are of particular importance here.

Technical perspective

⁴Wireless engineers have a complicated relationship with spectrum regulation. They are at the same time stuck within the world-view defined by current regulation, but are also the disrupting force causing that regulation to change. Right now, for more dynamic regulation, technical advances are enabling it to happen, and the engineering world-view desires more light-handed regulation. But, until this thesis, the technical perspective on spectrum enforcement is dictated by the current regulatory paradigm.

Cellular technology is a good example of how engineers and regulations have interacted in the past. Cellular systems evolved out of the regulatory structure of the time: they assumed they would have a particular spectral mask, tower restrictions, and no outside interference. The innovation was figuring out how to deal with their own self-interference, and dynamically assign different handsets different sub-bands for operation [8]. However, cellular technology was disruptive to regulation: the regulatory regime at the time was not liberal enough to immediately accept a new kind of architecture. So, new rules were developed, new certification procedures enacted, and a new allocation for cellular bands was found. From the first cellular design to the first deployment took about 25 years [19].

Now, again, technical advances are bringing about spectrum sharing. There are two important enabling technologies: frequency agility and methods of finding spectrum holes.

Frequency agility is key because it allows cognitive users to switch from one empty band to another [20]. The technology was developed to allow dynamic assignment within networks, and so was not immediately disruptive to regulation itself. The change to a shared spectrum regulatory regime comes with the realization that frequency agility can change the meaning of a stable spectrum usage right. Networks do not need to always use the same set of frequencies; they can rely on having a band for transmission without requiring *the same* band for transmission.

But having agile radios cannot allow shared spectrum without also having the means to find spectrum holes. Currently, spectrum holes can be found using a distributed sensing approach [21], or a centralized database approach [21]. Sensing includes listening to the local activity to discover whether a frequency band is already being used. The database solution includes registering

⁴The author of this thesis has a technical background, so the perspective of this thesis grows from this base even as it is trying to address the problems raised by other perspectives on the problem.

priority users in a database, so that opportunistic users can contact a centralized agent to find out which local bands are free.

Now that regulation is changing, the engineer’s goal for future regulation is to enable innovation. Cellular adoption took on the order of a decade [19]. Spectrum sharing was determined to be a good idea in 2002 with the FCC Spectrum Task Force Report [22], was made legal in 2008 [21], and the first commercial device by Adaptrum Inc. was approved in 2012 [23]. But the rules are still evolving [24]. Can the new rules be defined to shorten this time from conception to roll-out?

Simplifying and shortening the process is crucial because innovation is still happening along two fronts: how to build networks and how to find spectrum holes.

In terms of building networks, *no one yet knows the best way to use spectrum*. Information theory gives limits for how much information can be passed through a channel [25]. But only point-to-point links are well understood. Networks [26], multiple antennas [27], and interference cancellation techniques [28] are open areas of research that could produce big gains in the future. Even the limits that are known are not necessarily the right limits when trying to optimize total energy usage [29].

More importantly for this thesis, the process of finding spectrum holes is also still evolving. No one has found the best way to organize a network of nodes to sense for spectrum usage, although significant advances have been made [30]. As this thesis will discuss later, databases are still in their infancy in terms of exploiting their real potential.⁵ Combination approaches are also important possibilities [34]. In fact, this thesis will propose an architecture in which databases ultimately control access, even while individual devices may be sensing to find spectrum holes.

When building regulation that relies on changing technology, there are two choices. Either there must be a significant change to regulation every time these technologies change, causing major inefficiencies (this is the current path). Or, the rules must enable an evolutionary path for these underlying technologies to improve naturally. This thesis aims to enable the second, light-handed choice.

The difference between these two choices hinges on the trust model.

Right now, the FCC relies almost entirely on certification to produce trust – they guarantee that new devices will follow access rules by testing at certification time that devices cannot do anything but follow access rules. Under this model, *any* change to technology will require at least a new certification procedure. Changing the technology to find spectrum holes will require a change to the rules themselves as well.

When regulators require trust, the technical response has been policy engines.⁶ These are essentially easily-certifiable decision trees that guarantee a device will make certain decisions based on the sensory input. There are many problems with this approach. First, the policy engine cannot guarantee that all contexts will produce appropriate answers. Perhaps the network of sensing nodes in the TV bands are all located behind the same building. No policy engine can guarantee these nodes will be able to realize their predicament and not transmit.

More importantly, what happens when a new sensing technology is introduced? The first option is to build a policy language rich enough to cover most of the future technical possibilities. This approach will quickly run into halting [45] or self-reference [46] problems. The other option is

⁵See [31–33] for some of the work on how databases might be improved.

⁶The policy engine began with the Darpa XG program [35–37], but the community has jumped on this idea as the key to enabling trust. A significant amount of literature followed, including [38–44]

to build a unique policy engine for every new technology. Each new engine will need new regulatory approval. This can be expected to take a decade or more. There must be a better solution.

For new technology to be easily certified and deployed, a paradigm shift is needed in the trust model. The certification requirements cannot depend directly on the spectrum access technology being certified. This thesis is the first technical work providing a trust model with this property.

Policy perspective

⁷The policy perspective comprises three main subjects: what should usage rights look like in the future, what is the right way to get to those future rights, and how are the rights enforced once established.⁸

There are essentially three competing views of what more dynamic spectrum regulation should look like in the future: secondary (real time) markets, managed commons, and open access.

With secondary markets run by auctions (eventually in real-time), either the government [52,53] or private interests [14] hold wide swaths of spectrum and manage the distribution of access rights to actual users. In an extreme version, micro-transactions occur in which packets transmit their payment as they are flying through the air [54]. Managed commons are the decentralized version of secondary markets. Different networks have usage rules but the networks are trusted to make many local decisions themselves [55,56]. Open access is essentially transferring all of spectrum into unlicensed bands with power limits [57]. Individual networks must figure out how to coexist, and those that have better coexistence properties will be more successful, thereby encouraging innovation [57].

Related to this debate is how property rights (or more generally, usage rights) in spectrum should actually be defined [50]. Competing views argue that rights should be very specifically defined [49], perhaps along the nine dimensions of spectral radiation [58], or just through (maybe stochastic) power limits at the boundaries [10,49]. Alternatively, usage rights could be very loosely defined, relying on dispute resolution through courts to iron out details later [59,60].

The goals of the parties in this debate are similar. Each wants to provide the simplest rights to control interference but allow the most freedom for innovation. At the same time, the desire is to allocate spectrum to the most valued use. The discussion is about the best way to get to this goal.

Underlying all of these views and issues are Coasian bargains and transaction costs [13]. The idea is that disputes over interference will be solved through bargains between the affected parties to come to a mutually-desirable solution. Transaction costs are the cost of finding and implementing the solution. Obviously, if usage rights are not well defined, there may be more opportunity to converge to good rights in the future, but the transition period will be more difficult for the parties involved.

As with the future of spectrum access technology, it is not yet clear which version of the future of spectrum regulation is actually the best. It is also not clear how usage rights should be

⁷For a more extensive treatment of the policy perspective and literature, see [15]

⁸This section does not cover the significant policy effort to make spectrum usage rights more liberal. This means separating the statement of the rules from the service they are meant to contain. Originally, spectrum access rules defined even the protocol to be used for TV and radio [13]. Now, however, regulation uses spectral masks, maximum power limits, etc, that can apply to any service. This paradigm shift took many years and much work. For reference, see [47–51].

defined, or how extensive the transaction costs will be to change rights under any of these regimes. What is clear, however, is that none of these futures are possible without being able to trust that wireless devices will correctly follow their usage rights [61]. Further, actually being able to compare these futures requires being able to quantify what these transaction costs will be. This cannot be simply argued in rhetoric. This thesis provides the first quantitative results that can speak to trust and transaction costs.

Once the future of spectrum usage is chosen, what is the evolutionary path to it? Beyond a “big bang auction” [62], the choices are limited to deciding what will happen to the next available piece of spectrum. Should it be given to unlicensed or licensed use [63]? Should market mechanisms be used to make this choice [64–66]? Unlike the visions of the future, which are generally strategies for spectrum as a whole instead of individual pieces⁹, the movement toward any of these futures is generally considered on a band-by-band basis, adopting a generational evolution of spectrum regulation. This thesis questions this paradigm as well, giving the possibility of a smoothly-transitioning regulatory regime that does not require large changes between generations.

Finally, who should resolve disputes, and what should enforcement look like? These are philosophically distinct questions. The first is a path for evolving usage rights to better fit the needs and values of the involved parties. Enforcement assumes that the current usage rights are correct and should remain fixed, but need some mechanism to ensure they are followed. Both are likely required – the first as a mechanism to update rights, the second to keep the peace between transition periods.

For both, the debate in the policy literature is between traditional courts [55, 59] and spectrum tribunals [67]. Traditional courts have the benefit of already being in place and already being well versed in how to handle disputes. But, the current legal system is not filled with experts in spectrum. It is not even clear how interference disputes should be handled. For example, should they fall under nuisance law, or trespass law, or something entirely different [70]? Spectrum tribunals, on the other hand, are thought of like patent courts [67]. They are envisioned as courts that handle only spectrum disputes and so would be able to expertly address problems. However, they would also require entirely new infrastructure to create.

Regulator’s perspective

In spirit, regulators care about all of the above problems. They want to enable innovation and technical freedom. They want to understand the best version of future regulation and find the right way to get there. But regulators have a hard constraint on anything they choose to implement: *regulators must protect current spectrum users*.

Whether because of the strength of the lobby [47, 71] or the importance of the service (like public safety), regulators cannot allow new sharing technology to be deployed that will harm legacy services. This has two implications: sharing rules must be designed to limit interference to networks that were designed to operate alone in their bands. Also, any new device that is deployed must be completely trusted to follow that sharing rule. The FCC must prioritize trusted, safe operation over efficiency.

Right now, even as sharing rules are evolving, the trust model is not. The FCC trusts new devices because it uses certification to rigidly control the device’s operation. Section 1.1.2 will

⁹Some authors do suggest a zoning approach is appropriate where different solutions work for different bands and services [67]. For a technical response to this idea, see [68, 69].

show how this led to an interesting compromise on the rules in the TV whitespaces.

A unified view of the different perspectives

The views of the three parties can be unified by thinking of what a device should, can, and will do. The *should* relates to what usage rights should be now and in the future. *Can* relates to what a device is technically capable of doing, and *will* relates to trusting that a device will act appropriately once deployed.

Engineers tend to think of *can* and *will* being the same thing – the perspective right now is to design technology that can follow sharing rules, and then use the technical capability to prove that the devices will follow rules through policy engines. Policy experts tend to focus on the *should* and the *will*. They care about what the future should look like, and try to figure out how to use the current justice system to ensure the *will*. Regulators are most focused on the *will* – they need to make sure that new devices will not overwhelm the current spectrum users.

This thesis focuses on the *can* and *will*, arguing that these can be designed together, but are fundamentally separate things. Regardless of the future of spectrum regulation, there must be trust. But, the technical abilities that guarantee trust can be separated from the technical abilities that are capable of correctly finding spectrum holes. So, trust can be guaranteed, and then performance can improve later.

1.1.2 The TV whitespace compromise

The television bands are the natural place to start in experimenting with more dynamic regulations because TV towers are easy to share with, and the service is not safety-critical if something does go wrong. The problem here is relatively simple because the TV towers are at known, constant locations and are always transmitting. The TV receivers (after the DTV transition) have strict signal-to-noise ratio (SNR) requirements with a sharp decline in quality below the limit [72]. Whitespaces, or blocks of spectrum with no active TV channel, are abundant [73]. So, there is a huge opportunity, if whitespace devices can correctly figure out their location relative to the TV tower and receivers.

Technology has offered two solutions to finding these whitespaces: sensing and databases. Originally, sensing was the preferred option for two reasons: it was tempting to implement rules that required no new infrastructure. The technology could be integrated in a decentralized way in the devices, pushing the implementation burden to the private developers. Further, there was hope that this was sufficient because sensing had already been successful for the 802.11a WiFi standard to avoid interfering with radar [22].

However, the TV band problem is harder than that faced by 802.11a because the sensing must be accurate at low SNR to adequately protect TV receivers. Because of the need for robustness against uncertainty inherent in real noise [74], reliably distinguishing a TV signal from noise at low SNR requires a network of sensors working together [30]. As mentioned before, no policy engine can guarantee correct placement of sensors, so certification alone cannot allow trust of networks of sensing nodes. Sensing, therefore, could only be allowed with a *per-node* sensitivity of -114dBm [75]. With this rule, almost none of the whitespaces can be recovered [6].

Databases, which at first seemed too costly because of their required infrastructure, became *the* solution to finding whitespaces in the TV bands. Whitespace devices are required to submit GPS coordinates to a database, to receive a list of available channels, and to transmit only

in the indicated channels [75]. Databases are programmed to consider a channel available if the secondary device (with antenna height less than 30m above average terrain) is further than 11.1km from the closest TV tower with an active signal on the same channel¹⁰ [24].

The database approach fits nicely into the FCC's current trust model. The databases themselves are trusted both because they are certified to give known answers to simple requests and because there are only a few database providers in a long-term relationship with the FCC. Although not required yet by the current rules, the database message could be secured through cryptographic means. The whitespace devices themselves can be trusted because they are certified to only transmit on channels explicitly provided by the database.

There is a further limit on total power, and a power spectral density limit [24] to complete the basic access rules for TV whitespaces. Under these rules, the database solution is good, but not perfect, at recovering TV whitespaces [73].

Note that whitespace use is at this point a kind of experiment. There is not yet a compelling business model that fits into the TV whitespaces [76], although there are now commercial deployments of whitespace devices [23]. So, the current usage model is essentially an unlicensed one – the devices follow specific guidelines, but can be used for any purpose.

1.1.3 The near future of spectrum regulation

After the engineering details are worked out, the TV bands can be considered a success in building a framework to allow recovery of spectrum holes. But the TV bands are only the first step in transitioning to more efficient spectrum use.

The National Broadband Plan [4], endorsed by the Whitehouse [77], requires 500MHz of spectrum to be made available in the next 10 years. TV whitespaces alone cannot meet this demand. So the near future of spectrum regulation includes two projects aimed to free large swaths of spectrum for new use.

The first, the incentive auctions in the TV bands [78], aims to clear parts of the TV bands for uses other than whitespace devices. The auction will include three phases. A backwards auction will buy spectrum from TV broadcasting companies that want to give up their channel allocation. The remaining channels will then be reorganized to pack them efficiently into a smaller band. The newly freed spectrum will then be sold at auction. It is unclear at this point how much spectrum will become available and how much whitespace will be left. But the resulting organization of the TV bands will certainly use those bands more efficiently than they are used now.

The second project, outlined in the PCAST report [79] and endorsed by the Whitehouse, aims to extend spectrum sharing to recover unused federal spectrum. The premise is that it is not an economically viable option to directly transition federal users from their current allocation to a more efficient allocation. The better solution is to share federal bands, let commercial entities design new devices, and transition federal users to commercial products in more efficient bands when the transition is natural.

The PCAST report plans to enable sharing by building infrastructure to bootstrap the database solution into the federal bands [79]. Given the FCC's current trust model, this is potentially very hard. If spectrum holes in federal bands are primarily geographical and change slowly

¹⁰The rules are more complicated than this statement makes them seem. The availability of a band depends on distance from a TV tower with an active channel or an active adjacent channel. It also depends on the secondary's height above average terrain. See Paragraph 15 of [24] for the full current distance rules.

in time, databases will port naturally. These are the same characteristics that make databases relatively simple to implement in the TV bands. If, however, the spectrum holes in federal bands change quickly over time, the databases will need to be able to catalog and distribute large amounts of information very quickly. This issue will be discussed in more detail later.

The PCAST report also opens the door for two important new ideas. First, it proposes that in the future, new generations of primaries must be designed to enable secondary sharing. It has been noted [80, 81] that the primary has control over the usability of spectrum holes because the primary controls its own receiver design. With a poor receiver filter, the primary will be more easily hurt by interference. Regulators must therefore decide what can be reasonably called “harmful interference” and what is actually just an overly-sensitive primary.

The PCAST report makes this decision explicit by proposing that standards be created for primary receiver design [79]. The FCC Technological Advisory Council has followed this advice with a whitepaper on receiver standard design [82]. This proposal is similar in spirit to the trust-through-control model embraced by the FCC. If primary receivers need to be able to handle some interference, then the only available option is to directly specify the expected receiver performance.

The PCAST report also introduces the idea of licensed secondaries – these devices will have to respect primary priority, but have more guaranteed usage rights than the current unlicensed-type best effort available in the TV whitespaces. This is important because it is the first sign of really embracing the reality that new licenses can have stable usage rights without long-term band assignments.

These two ideas come together to indicate a real policy shift. In the future, spectrum sharing will be the *norm* instead of the exception. New networks built by current or new license holders will need to be designed with sharing in mind.¹¹

As a final note, this thesis adopts the stance that a shared spectrum future must come with a caveat – it is very tempting to think that building networks that share means building networks running the same underlying protocol. Essentially, this means the future will have every network running some version of LTE, which already includes many technical provisions for sharing [84]. While this may actually be the best future, it is not necessarily so. If it is important to enable networks that are *not* running LTE, it is important to build regulation that does not accidentally force this future to happen.

1.2 The need for, and properties of, a technical regulatory layer

The core cognitive radio insight is that having more dynamic, flexible regulation requires moving regulatory decisions closer to runtime, where they can be made based on local context. To allow decisions at runtime, there must exist a technical entity trusted to ensure those decisions are made correctly. The required properties of this entity can be seen more easily by thinking of regulation as an abstract hierarchical control problem, as illustrated in Fig. 1.1.

Engineers tend to think of cognitive radio as including only the bottom half of this system. The cognitive device is trying to access spectrum, with input from local context, a database, or both. The control problem here is trying to control noise and the uncertainty of shifting spectrum holes as other networks’ usage patterns change. The regulation problem, however, is different. The regulatory problem is to make sure wireless devices coexist peacefully without causing too much

¹¹For more on this conceptual shift, see [83].

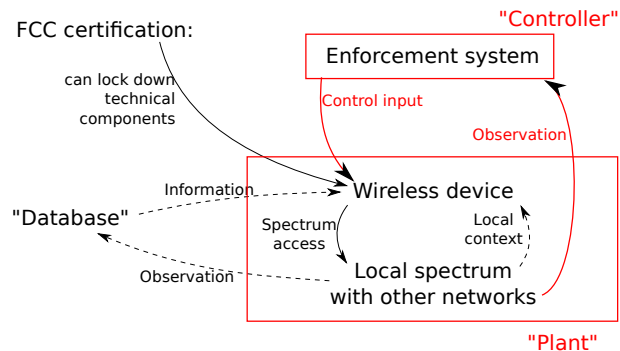


Figure 1.1: The regulatory “control system.” Engineers tend to think of only the control system of the wireless device trying to access local spectrum with help from a database or other entities. Regulators, however, are really trying to use certification and enforcement to control any *temptation* the engineers may have to cause interference for the benefit of their own networks.

interference to each other. So, regulation is trying to control the *temptation* to build devices that will cause interference in exchange for higher utility.

So, for the regulatory problem, the “plant” can be thought of as the entire system of spectrum and the wireless devices that access it. There are two regulator inputs to this system. The FCC certification is done once, before devices are deployed. Then, there is an enforcement mechanism that acts as the “controller” to provide ongoing input.

This abstract control model clarifies the primary goal of spectrum regulation and the technical meaning of “trust.” In control theory, the first order purpose of the controller is to keep the plant stable. In spectrum regulation, a stable plant means that none of the diverse deployed networks are running amok and causing too much interference to any of the others. Trusting the deployed systems means that the spectrum plant is guaranteed to remain stable.

This section will argue that moving regulation decisions closer to runtime will place requirements on the certification and enforcement systems that ensure trust. Because runtime decisions naturally introduce instability to the plant, one of these requirements will be that the enforcement mechanism has an element that also operates close to runtime.

To actually allow dynamic regulation, therefore, there must exist a *technical regulatory layer* that actively manages spectrum access. What does this spectrum manager need to do, and what should it not be allowed to do?

We believe the answer to this question must be guided by the ideas presented by Lessig in *Code* [85]. He focuses on copyright, where digital rights/restrictions management (DRM) software implements the restrictions on usage rights defined in copyright law. The problem is that the software developers have the ability to redefine the enforced law without any required consultation with Congress. Once that new “law” has been deployed, it is very difficult to use traditional courts to change it.

The same kind of problem exists on the Internet to an even more extensive degree: those who code virtual worlds can define the rights their “citizens” enjoy without any Constitutional consultation. There are no courts that can enforce real-world laws and rights in digital communities.

There are two lessons to learn: when implementing code that will act as law, the code must be carefully designed to ensure it upholds the desired values. Further, if the implementation

cannot uphold the value, either the value must be reconsidered or a new kind of implementation must be found.

Spectrum regulation must take these lessons to heart. For example, if regulators care about innovation in the technical means to find spectrum holes, they *cannot* use a trust model that precludes such innovation if a better model is available.

This section therefore begins by crystallizing the goals of spectrum regulation and the tasks of the regulator to achieve those goals. The transition to a sharing paradigm is then understood as moving regulation activities closer to runtime and into the set of operations done by the technical regulatory layer. Finally, by using the control analogy, this section shows the primitive technical capabilities that the technical regulatory layer must possess.

1.2.1 The goals of spectrum regulation

In essence, spectrum regulation is about facilitating spectrum access by different devices, using different technologies, as part of different networks that are not jointly designed. There are many views of what regulation should look like in the future – but which one is right? Ideally, spectrum regulation will fulfill the following goals.¹²

- Political goals:
 - Politically neutral and transparent
 - Rationally rationing resource when scarce
 - Maintain access for services whose value is not monetary (like public safety)
 - Low barriers to entry
- Technical goals:
 - Encourages technical innovation
 - Supports diversity of uses and technology
 - Easy to introduce new and better technology
 - Regulation does not depend too much on the technology
 - Allows high spectral utilization
 - Evolves to protect changing primary needs

The political goals are designed to reduce the opportunity for corruption, monopoly, and loss of valued services. Auctions and markets are the economists' traditional tools to ensure political neutrality [52] and rationally rationing in the case of congestion [52]. However, auctions cannot be the only mechanism, as they do not easily preserve valued but low-revenue services like public safety networks or other military and government uses.¹³

Low barriers to entry protect both political and technical values. They help prevent monopolies [86], and they also help encourage technical innovation. Think of a start-up trying to get venture capital funding. If their device stands little chance of surviving the certification process, or if that process could take anywhere from one month to ten years, few VCs would want to invest regardless of the technology. The success and support for development of a product should depend on its ability to survive the market, not its luck in traversing the certification process.

¹²These are the goals the author can identify as important guidelines. This should certainly not be taken as an exhaustive list of everything spectrum regulation needs to accomplish.

¹³The PCAST report suggested that this problem could be overcome by having fake auctions that allow government services to bid without requiring actual money to change hands [79].

Many of the technical goals are also aimed at encouraging innovation through simpler certification requirements. The ideal way to make certification easy and transparent would be to create a process that will work *even for as yet undiscovered systems*. This is not possible if the process depends too strongly on the particular technology. TV whitespace rules are the proof that this ideal is not crazy – certification depends on the mechanism to communicate with the database and the spectral mask. Beyond that, *any* network architecture is allowed, even if it looks nothing like what has been come before.¹⁴

Next, the *regulatory overhead* must be kept low. This is any inefficiency in spectrum usage caused by the regulations themselves. For example, the collapse of whitespaces using a per-node sensing requirement is one type of regulatory overhead. Even with databases, the spectral mask for the current TV whitespace devices is extremely difficult to achieve [24, 88]. The extra processing power and any resulting communication inefficiency incurred because of the spectral mask is another kind of regulatory overhead.

Regulatory overhead is not currently well understood, partly because the overhead appears in so many different forms. Sources of inefficiency are often qualitatively argued [13, 89], but there are few metrics for quantitative comparison. The first metrics for regulatory overhead are in [90, 91] which measure recovered transmission area recovered in the TV bands. One of the goals of this thesis is to identify the different types of regulatory overhead associated with enforcement so that different kinds of mechanisms can be quantitatively compared in the future.

Finally, regulation must evolve to protect changing primary needs. Legacy systems are not static – they evolve as new devices are built and come online. Although the current legacy systems assume they are isolated from other networks, *new primary systems could be designed with different assumptions*. The PCAST recommendation for primary receiver standards even suggests that new primary systems should be required to be sharing-friendly [79]. One such provision may be requiring primaries to transmit a beacon signal so they are easier to sense [74].

Sharing rules must change along with newly deployed primary systems so that secondaries are not forever constrained to avoid helpless primaries. The path to future regulation requires all of these systems to evolve together.

1.2.2 Moving regulation activities closer to runtime

This part takes a closer look at the different regulatory tasks under three different regimes: command and control, TV whitespaces, and a proposed scheme with a full technical regulatory layer. The purpose is to understand how the requirements of different tasks change as regulation decisions are moved closer to runtime. The timescales at which regulation tasks take place can be compared for the three schemes in Table 1.2.2

Table 1.1: Timescales of regulation activities under different regulatory schemes. C&C is command-and-control.

Timescale	C&C	TV whitespaces	Proposed
Long-term regulatory	Define rule paradigm Allocation Assignment	Define rule paradigm Allocation	Define rule paradigm
Device lifetime	Certification	Certification	Certification
Human intervention	Enforcement	Enforcement	
Network operation	Data Transfer	Assignment Data Transfer	Allocation Assignment Data Transfer Enforcement

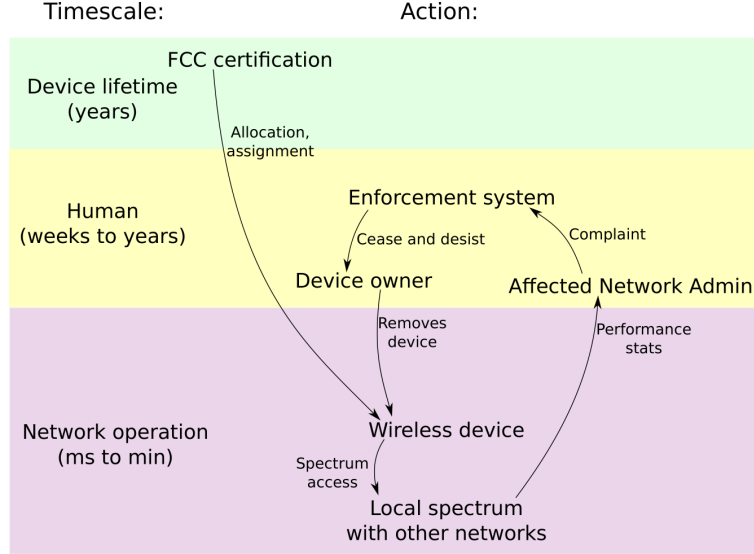


Figure 1.2: The command-and-control system. Although the network is running at the network operation timescale, all regulatory operations happen at much slower human time scales. In particular, any enforcement events require a network administrator to complain, which triggers an FCC investigation and perhaps a cease-and-desist letter. If all networks have their own private frequency band, this slow intervention time scale is sufficient to control interference.

Command-and-control

The evolution of command-and-control established the set of tasks that regulators use to manage spectrum. In its most technically-neutral form¹⁵ these tasks are the following:¹⁶

1. Define rule paradigm type – Command and control regulation separates different spectrum licenses by spectral masks, power limits, tower heights, geographic areas, and time spans.
2. Allocation – define the type of use for any particular frequency band – TV, cell, unlicensed use, etc. – and set the parameters of the rules accordingly.
3. Assignment – choose the particular companies that can access subdivisions of the allocated frequency band. Verizon and AT&T operate in different parts of the cellular band.
4. Certification – check that assigned companies’ proposed devices follow their respective rules.
5. Data Transfer – This is a placeholder for the actual operation of the deployed network. Regulators have no defined tasks here.
6. Complaints, check-in, and enforcement – If a company notices someone else causing interference, they can complain to the FCC, who will investigate the claim and issue cease-and-desist orders.

Under command-and-control, all regulatory actions occur in this order and at a human time-scale.¹⁷ People must work to define the spectrum rule type and choose appropriate parameters for them when allocations are defined.

Allocations (or setting the specific parameters of the rules) last for decades. TV band allocations have changed only twice in their history: when UHF frequencies were added [92] and when the digital TV switch happened [93]. They are changing again now with the incentive auction [78]. Allocation changes only happen when prompted by changing technology. When cellular architectures were created, the FCC then had to find a new spectrum allocation that would fit the new technology [19]. These decisions were naturally made at the scale of the architecture design lifetime.

Assignments, which determine the specific company that will use a band, also last for years. Again, this is natural. In order for a company to invest the high startup cost for wireless infrastructure, it needs a stable guarantee that its access rules will not change [49]. So, if a device is frequency-locked (vs the agile devices that exist now), assignments must last for a cycle that appropriately ensures stability, on the order of ten years.

Certification processes are created for every new device and are applied before any device is deployed. So, it naturally happens on the time scale of device lifetime, usually years.

Enforcement, even today, relies on a complaint by a human, so it too operates at a human time scale. When a complaint is made [17], the FCC sends a van with a directional antenna to investigate, and then issues cease-and-desist orders or fines when necessary [18]. This process can take days to years to complete.

¹⁴There has been much work in the policy literature along these same lines. Liberalized spectrum allocations [50,51], or investigating whether spectrum is fungible [87] is about streamlining regulatory decisions to fit more than one technical standard or architecture.

¹⁵Policy experts have worked hard to liberalize spectrum regulation [51], so there are actually many different forms of command-and-control whose rules depend to greater or lesser degrees on the technology they are governing. From a technical perspective, all of these rule paradigms look the same because they all have the same operational list of tasks that are performed in the same order.

¹⁶A similar, but slightly different break-down of tasks can be found in [61].

¹⁷To see this list in the order of the more precise timescale at which tasks occur, and to compare to the other versions of regulations, see Table 1.2.2.

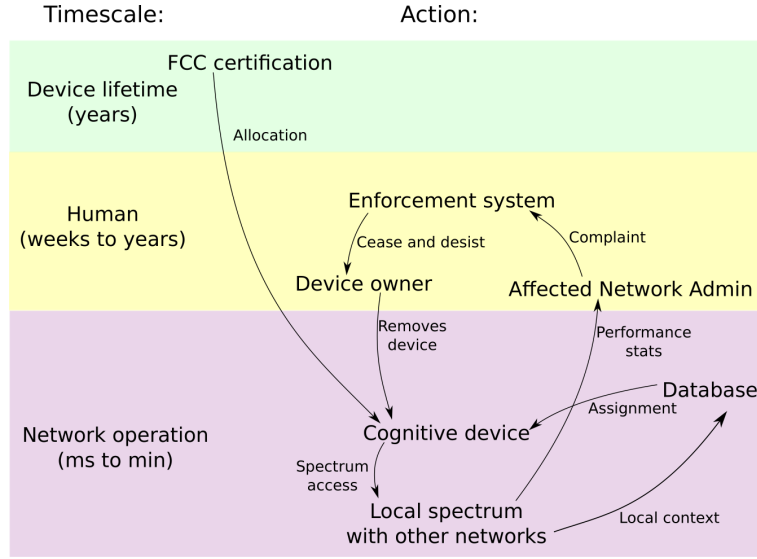


Figure 1.3: With the TV Whitespace ruling, cognitive devices are able to opportunistically use spectrum holes, with help from a database. However, intervention to turn off malfunctioning devices, for example, still goes through the same complaint, investigation, cease-and-desist cycle as with command-and-control. This cycle is far too slow to effectively deal with problems.

Fig. 1.2 shows these tasks in terms of the control metaphor with the time scale included. This model helps to understand why the FCC has the trust model it does. If enforcement occurs at the human time scale and the data transfer occurs at the network operation time scale, *enforcement cannot stabilize the spectrum access plant*. It is too slow [94]. So, the plant must be designed to always be stable. This means only certifying devices that can prove they will not cause interference. This implies giving all networks their own, fixed spectrum assignment that is sufficiently isolated from all other spectrum assignments.

When spectrum is abundant, command-and-control does well in meeting the goals for spectrum regulation. It can meet neutrality goals through auctions on assignments and even allocations [52]. It can protect public safety and other services with exceptions to auctions for these services. It can even provide the stability of rules required to encourage innovation. When spectrum is abundant, inefficiency does not even restrict new entrants.

The only real regulatory overhead is time and uncertainty. New allocation and certification procedures must be developed for every new technology. This process can take decades and is completely opaque [95], creating significant risk for new ventures.

Of course, the ability of command-and-control to meet regulation goals breaks down in the face of high demand and spectrum scarcity. The solution is to move some regulation tasks closer to runtime.

TV whitespace regulation

With the introduction of frequency-agile radios, a stable rule set does not require a fixed assignment for the lifetime of the device. Already in current cellular systems, individual handsets

are dynamically assigned sub-bands according to the local situation [8]. WiFi systems similarly can change their operation channel at runtime [96].

TV whitespaces are the first attempt to have network-level frequency assignment happen at runtime as well. This changes the operation of regulation to look like this:

1. Define rule paradigm (human time scale)
2. Allocation (human time scale)
3. Certification (human time scale)
4. Assignment (network operation time scale)
5. Data transfer (network operation time scale)
6. Complaints, check-in, and enforcement (human time scale)

Notice that certification and assignment have switched places, and assignment is now done at the network administration time scale. Essentially, the FCC has created a new rule paradigm that includes spectral masks, power limits, no-talk radii, and database consultation. The allocation is for a type of service called “TV whitespace device” that specifies the parameters of the rule paradigm. Certification involves proving that a device will consult a database for the runtime frequency assignment based on GPS location. Enforcement is assumed to work the same way it always has: complaints will prompt an investigation. In extreme situations, complaints may prompt a change to the assignment rules, which will then be propagated through the databases.

This version of the regulatory process patches some problems of the command-and-control model. It improves efficiency in the TV bands by allowing whitespace reuse. It also simplifies the certification process by streamlining many different services through the same process of checking the spectral mask, power limit, and database access protocol. It, finally, lowers barriers to entry by opening new, effectively unlicensed spectrum for new entrants.

The current rules are a good first step, but they are not sufficient to actually protect TV receivers. Consider the control view of TV whitespace regulation shown in Fig. 1.3. As with command-and-control, the enforcement controller is acting at a human timescale. This time, however, regulatory decisions are being made at runtime. The controller is too slow to actually control runtime behavior, so the question becomes whether certification alone is sufficient to ensure a naturally stable plant.

Unfortunately, the answer is no. What happens when a device malfunctions? Removing that device requires the full path of complaint, investigation, cease-and-desist, and removal. Other problems can arise as well. In urban areas, if there is enough market penetration by cognitive devices, the aggregate interference will overwhelm TV receivers under the current rules [31]. Further, the system is vulnerable to attack. The database transmissions are not currently required to be secured [24], so the database message to change the no-talk radius could be spoofed. Even a type of denial of service attack is possible in which all cognitive devices are instructed to turn on. This would certainly overwhelm TV receivers [31].

These examples suggest that the TV whitespace rules are missing some key technical primitives required to trust the deployed systems.

There must exist a way to correct instances of too much interference. This includes both the ability to recognize that interference is happening, and the ability to turn off cognitive devices. The controller needs feedback and a kill switch that operate *close to runtime*. If regulators want to be able to direct this kill switch at individual malfunctioning devices, individual identity will also be required.

These extra technical abilities will allow trust. Further enhancements can help improve

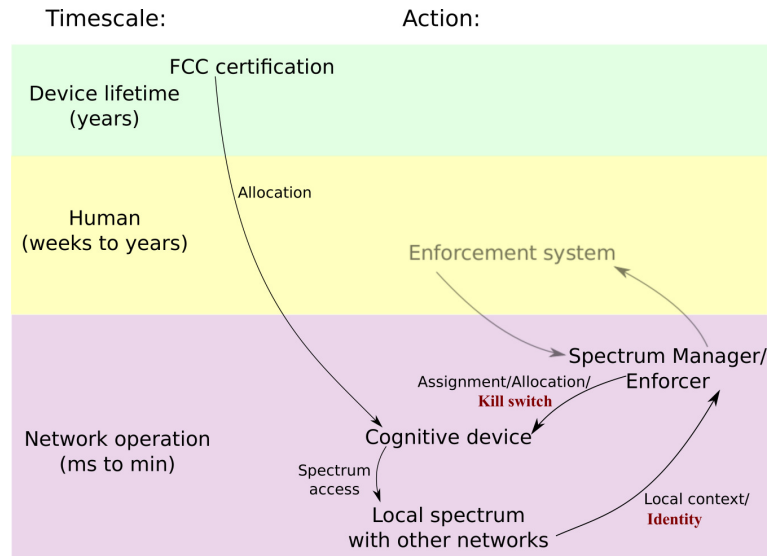


Figure 1.4: If the database has access to identity and a kill-switch, the new “spectrum manager” will actually be able to effectively protect primaries. These same technical primitives can also be used to enact a run-time enforcement scheme.

performance.

Right now, the databases give channel assignment information based on individual node locations. But they have much more information available. They can easily record node locations to make assignment decisions based on aggregate statistics. This would allow databases to restrict assignments when the node density is too high so they do not have to rely as much on the kill-switch.

But if the databases are making decisions based on aggregate node information, why are they only allowed to change channel assignments? From a technical point of view, no-talk radii and power limits are no different from channel assignments. They only require a slightly more complicated decision algorithm.

Fully dynamic regulation with a technical regulatory layer

If the extra technical primitives are included, and allocation parameter decisions can be made at runtime, the regulatory process changes to the following:

1. Define rule paradigm (human time)
2. Certification (human time)
3. Allocation (network operation time scale)
4. Assignment (network operation time scale)
5. Data transfer (network operation time scale)
6. Complaints, check-in, and enforcement (Potentially mixed time scale)

In the TV whitespaces, databases could let sparse secondaries transmit closer to the primary or let them use higher powers. The database could even change the no-talk radius or power limits to deal with congestion. These efficiency gains require no extra technical abilities to achieve.




Figure 1.5: How much spectrum a database can recover depends on its time scale relative to the primary. Even if the primary is only using the band 25% of the time, the database cannot recover any spectrum holes if the primary usage is spread out over the database time slots.

The resulting proposed system is illustrated in Fig. 1.4, and includes the technical requirement of enforcement action close to runtime.

The “spectrum manager and enforcer,” which replaces the database, is the technical regulatory layer proposed at the beginning of this section. The technical regulatory layer is responsible for determining operating parameters, distributing them to all managed devices, and ensuring that the parameters will be followed correctly.

The form of the regulatory layer is determined by how it is supported by certification and the enforcement mechanism. These two must work together to ensure the spectrum “plant” remains stable. For example, if the regulation layer includes a centralized database entity, secondaries must be certified to interact with the database or it cannot possibly control interference.

On the other hand, a network could generate its own access parameters through sensing, implementing part of the regulatory layer in a decentralized way. It is not possible to use certification alone to bound all the uncertainty inherent in networked sensing. So, a different kind of certification/enforcement system is required.

The next section explores different ways to leverage certification and enforcement to produce systems that are trusted to be stable. It concludes that spectrum jails are a good way to use the technical primitives that are already necessary to extend trust to even decentralized decisions.

1.3 The case for spectrum jails

1.3.1 In the future, trust through certification alone will not be good enough

Even enhanced with the extra technical primitives identified in the last section, trust through certification alone will not be sufficient to realize spectrum sharing in the future.

Assume, like in the PCAST report, that future spectrum sharing will use a database-centered mechanism to coordinate access to spectrum holes. There are two major sources of inefficiency. First, there may be mismatch between the allocation the database can provide and the technical capability of the device using that allocation. Ultra wide band systems [97] cannot use

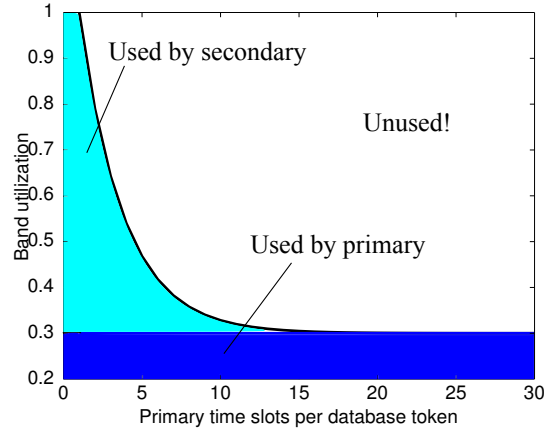


Figure 1.6: Even if the database can predict perfectly which spectrum holes a primary will be using, it can only recover the time slot if the *whole* time slot is empty. So, if the primary uses 30% of the band, and the database token includes several primary time steps, this mismatch will decrease the recoverable spectrum. A single order of magnitude difference in time scales will effectively eliminate recoverable spectrum.

the same allocation boxes as a narrow-band signal. The database must be complex enough and must be updated with every new type of secondary in order to fill requests correctly.

But there exists an even simpler source of inefficiency: a mismatch between the time scale at which the database can produce and deliver allocation information and the time scale of the spectrum hole it is trying to recover. In the TV bands, this is not a problem – secondaries must consult with the database only every 24 hours [97], and spectrum holes are indefinite from the point of view of the secondary device. The database can easily keep up with demand at this time scale.

What happens when the primary is not always on or always off? If the spectrum holes exist on the order of minutes or seconds, will the database be able to keep up? The problem is illustrated in Fig. 1.5. Although the primary is only active 25% of the time for both cases, the database can only recover the spectrum holes when its tokens are distributed at the same time scale as the holes. In the bottom example, the database cannot recover any spectrum.

Fig. 1.6 shows a back-of-the envelope calculation of the inefficiency resulting from a timescale mismatch. Assume the spectrum holes exists at some time scale, and the database can only send information at the rate of 1 per M primary time steps. Even if spectrum is idle 70% of the time, a time mismatch of only one order of magnitude collapses all of the spectrum holes. There must be a way to do better.

Using the control analogy from earlier, trust through certification with the extra technical primitives is like building a generally stable plant and using the controller only for infrequent problems. It is far more efficient to build an unstable plant and actively use the controller to stabilize it. For example, fighter planes are naturally unstable systems, and must have a computer actively work to keep them in the air. However, their efficiency is far higher than that possible with a commercial, stable airplane [98].

Applying this idea to spectrum regulation, regulators should let secondaries make spectrum access decisions. There is great potential for interference (a.k.a. instability), but that can be controlled by relying more heavily on a runtime enforcement system.

The control analogy is only an analogy. The problem of enforcing spectrum sharing rules does not fit naturally into traditional control theory. So, the next few parts will flesh out the enforcement problem, from goals to available tools. Section 1.3.5 will argue that spectrum jails are a good starting point for addressing runtime enforcement.

1.3.2 Goals and design philosophy of the certification/enforcement system

What should the certification/enforcement systems look like? What is the right model for trust? It must at least satisfy the following goals:

1. Works – if the certification/enforcement system does not adequately limit interference, it has obviously failed.
2. Fulfills as many of the regulatory goals as possible.
3. Keeps regulatory overhead low

In this thesis, these goals are distilled into a single goal:

- Guarantee trust regardless of secondary technology, and then ensure performance will improve as technology improves.

The purpose is to allow the same enforcement mechanism to work for any secondary device presented for certification. This is what this thesis means by “light-handed regulation.” But, performance will be tied to the technology, so that better secondary devices will work better. Likewise, as regulator technology improves, all secondary devices will perform better.

This thesis fulfills this goal by espousing the design philosophy of applying certification/enforcement activity at the bottleneck – that point in the operation of wireless networks where it can have the most effect for the least amount of pressure.

This concept is not new. Compliance theory is a field of research that investigates how to apply different kinds of enforcement action to guarantee compliance with laws by corporations or individuals [99]. The idea is to use the action (whether education, personal contact, sanctions, etc) that is most effective at inducing proper behavior instead of jumping straight to a punishment for illegal activity.

A similar philosophy was used by security researchers trying to understand how to combat email spam [100]. The usual approach is to build better spam detectors, but this simply leads to an arms race between the spam creators and the email filters. The actual bottleneck in fighting spam is the small handful of banks that process the payments for transactions initiated through spam emails.

In spectrum regulation, the problem is more complicated. We are trying to target the enforcement bottlenecks in a system that does not yet exist! There is lots of discussion on what future enforcement should look like, but rhetoric is not enough. Technical studies are needed to explore the enforcement possibilities, understand what the tradeoffs even are, and then choose the right bottleneck.

This thesis begins that process by exploring one reasonable direction. In spectrum, secondary misbehavior manifests itself as interference to primaries. So, this thesis takes harmful interference itself as the right bottleneck and tries to control that directly instead of specifying particular allowed secondary technologies. How should the control be implemented? The possibilities are discussed in the next section.

1.3.3 Assumed traits of secondary devices

What kind of enforcement mechanism is used naturally depends on what the regulator is willing to assume about the companies and secondary devices that they are trying to regulate. So, this section discusses different assumptions, their weaknesses, and how certification and enforcement can be designed to reflect those assumptions. This section will argue that sanctions are a good first approach for a technical study of enforcement.

Completely known technical capabilities

This is the assumption that the FCC is currently using. In order for certification to prove that a device will follow rules, its operation with respect to those rules must be completely known. Because the rules take the form of technical boundaries like spectral masks, power limits, etc, the technical capability of the device must be completely known.

As has already been argued, this is a very strong assumption when devices share spectrum dynamically, because all operating conditions cannot possibly be tested.

Reputation and long-term relationships

This is the assumption that allows databases to be trusted. Essentially, if the FCC is in a long-term relationship with a company that values the relationship and its own reputation, that company can be trusted to act correctly. There may be some certification requirements, and complaint investigation. The basis of trust, however, is derived from the relationship itself instead of these processes. By necessity, this would imply a vetting process for new relationships, and a small number of trusted companies.

This is a good assumption for those entities serving as databases and spectrum managers in the future. The FCC has to be able to deputize these companies to make regulatory decisions. With the needed kill-switch, this includes both allocation/assignment type decisions and enforcement decisions. A small number of long-term deputies seems like a good solution.

However, for secondary device manufacturers, relying on reputation and long-term relationships is less desirable particularly from the perspective of innovation. It would be very difficult for new entrants to establish their trusted relationship, so new technologies would come mainly from already established companies.

Network effects

Network effects allow trust through the need to inter-operate in order to succeed. For example, there is no law or regulation requiring Internet components to use TCP/IP. However, all Internet components do use TCP/IP because they cannot otherwise communicate with any other device on the Internet.

For spectrum regulation, trust based on network effects would be difficult to achieve. Although standards bodies like IEEE create protocols for spectrum sharing, like 802.22 [101], there is not a strong natural incentive to adopt the protocol. Paradoxically, TCP/IP may actually contribute to the weak network effects in spectrum. Because TCP/IP provides a common protocol to build upon, there is less need for the physical layer to adopt a common platform.

Ultimately, different networks and companies want to *avoid* direct coordination requiring joint design. Until the perfect sharing standard emerges, which attracts use because of its superior performance, trust cannot rely on network effects.

Rationality

Although generally a poor assumption for humans [102], rationality is a good starting point for radios. The purpose of a radio is to transmit information with some quality of service. Presumably, radios will be designed to optimize that quality of service.

A way forward for enforcement can be inferred from work by Etkin *et al* who use game theory to analyze what happens if networks must coexist without external influence [103]. The result shows that spectrum will only be shared equally with equal amounts of vulnerability. If two networks can cause similar amounts of interference to each other, they will share fairly because of the credible threat of too much interference. If, however, one network can cause interference while the other cannot, then the powerful network will transmit, while the weaker network will stay silent.

This idea will be taken further in this thesis to use certification to *engineer* the vulnerability required to make secondaries follow sharing rules. Essentially, some outside force must be applied so that a radio will have better utility when it is operating honestly and not causing interference. This outside force will be a sanction, or an in-kind punishment that is applied when interference is detected.

Human traits like altruism

Radios are not necessarily selfish agents, because they are designed by humans that will almost certainly design good, honest radios. Indeed, human traits like altruism have been engineered into systems in other contexts to bring large gains in performance [104, 105].

However, unless altruism can be certified, it cannot be assumed to be present in all deployed devices. Therefore, rationality is a worst case assumption that gives a good baseline for performance, even though reality may be able to do much better.

1.3.4 Sanctions for radios

The key to making a sanction-based enforcement system work is choosing the right sanction. But how would one effectively punish a radio? The answer follows the rationale of the law and economics literature for humans.¹⁸

Humans have natural vulnerability to punishment because humans have basic needs for a high quality of life. Humans suffer without appropriate leisure time and money. Therefore, they are vulnerable to jail and monetary sanctions [109]. Whether a sanction is effective then depends on whether jail or monetary sanction is applied correctly to the person [112].

What are radios vulnerable to? A radio wants to transmit information with the highest possible quality of service. A radio that will try to optimize its quality of service will therefore respect any threat to that quality of service. Etkin *et al* shows a good example of this: radios will share fairly to avoid harmful interference [103]. But, like people, not all radios are vulnerable to the same punishment.

¹⁸This literature grew from the seminal papers [106–108]. See [109–111] for a general introduction to this area.




Figure 1.7: The utility a secondary can derive from cheating is a boost to the quality of service it can provide. Because the service is transmitting data, this utility really only lies along three dimensions: throughput, energy usage, and “sensitivity to burstiness.” This last one is capturing the sensitivity to the timing of packets, so voice communication will have a different utility along this axis than transmitting a large file. This thesis covers secondaries with throughput and energy-based cost functions.

A radio’s utility is dependent on its quality of service matching the type of service it is trying to provide. If a laptop is trying to download a large file, it may require high average throughput, but bursty transmission may not matter. However, for any voice communication, even lags on the order of ten milliseconds are highly disruptive for the user [113].

In general, the quality of service a radio requires can be measured along the three axes in Fig. 1.7: throughput, energy used, and burstiness. Burstiness is the timing of packets and whether there are long delays in the middle of the transmission.

An effective sanction, then, reduces utility along these dimensions. Forcing a radio to turn off for a certain amount of time will be an effective sanction for those radios that want to maximize throughput. Forcing a radio to burn energy will work to deter an energy-sensitive user from causing harmful interference. Because the goal is to establish trust first, and then improve performance, this thesis tries to build a sanction that is effective for all secondary devices *regardless of their individual sensitivities*. Chapters 2 and 4 will cover secondary devices sensitive to throughput and energy-based sanctions. Radios that are sensitive to changes in packet timing are left for future work.

As presented in terms of vulnerability, focusing on in-kind sanctions makes intuitive sense. However, from a traditional economics perspective, this is not the natural choice. Traditionally, sanctions are assumed to be fines [109].¹⁹ But for cognitive radio enforcement, in-kind punishment is the right choice for three reasons: implementation, estimation uncertainty, and operational differences. The argument will be presented briefly here and explored quantitatively in Chapter 7.

¹⁹It is traditionally argued that fines are more efficient than in-kind punishments because they redistribute wealth instead of destroying a resource [109]. However, some have argued that in-kind sanctions can be more efficient in some cases [114, 115]

Implementation

Section 1.3.5 will explain how in-kind punishments can be implemented through the technical primitives from Section 1.2.2 that are already required for trust in cognitive radio. Fines cannot be implemented so simply because they require a billing system. This means that identities for cognitive devices must be individual, instead of groups, even in the first roll-out. They also must be persistent and secure enough to tie a radio to a financial account.

This is not true for in-kind punishments. Section 1.3.5 will show that they can be implemented in the device itself, so all punishment can be local. The identity can also change periodically, because it only needs to locally and temporarily tie an offense to a device.

Estimation Uncertainty

For rational agents, the decision to cause interference is based on maximizing a utility function. So, the punishment must affect the utility function such that it is less advantageous to act honestly than to cheat. This thesis takes the stance that the punishment must effectively deter cheating regardless of the characteristics of the secondary.

With an in-kind punishment, quality of service can be degraded along a limited set of dimensions. Also, because there are limits on the possible quality of service in any dimension, it is possible to bound the size of the sanction needed to deter secondaries from cheating.

With fines, it is less clear that it is even possible to set an appropriate fine. How much is a cheating transmission worth? It depends on the user, the quality of the cheating transmission, the business practices of the company, and many other factors. Where would one even begin in developing a monetary utility function for a secondary device?

Does finding the correct utility function even matter? Presumably a fine could be set very high, and then it would certainly cover all but the most highly valued secondary cheating transmissions. As the rest of this thesis will show, the sanction needs to be set to the lowest available effective sanction. This is because no identity system can be perfect. What happens when a device is wrongfully convicted? It will end up paying that very high fine, causing significant overhead. If the right fine cannot be estimated appropriately, in-kind punishments will be a much more efficient solution.

Operational Differences

The law and economics literature recognizes in-kind and fine based punishments to be qualitatively different things. The main difference comes from being able to treat a fine as a price [115, 116]. [116] has an illustrative example: if parents must pay a fine if they pick their kids up from school late, some parents will try to avoid the fine and arrive on time. Other parents, however, think of the fine as the price of a few hours of after school child-care. The fine would have to be higher than these parents are willing to pay in order to have any deterrent effect.

The problem is that fines and in-kind punishment operate very differently. Fines can be thought of as a payment for a higher level of service, or as a way to transition resources to higher-valued uses. In-kind punishments, however, serve only as a deterrent. So, regulators must decide the kind of protection offered to primaries: is it protection regardless of the secondary service? Or are primaries protected only from lower-valued secondaries?

1.3.5 Putting it all together: spectrum jails implemented through databases

This thesis proposes that implementing sanctions as a “spectrum jail,” using databases and the technical primitives already required in TV whitespaces, is a good way to extend trust to secondary devices, even with light-handed regulations.

As already established, to truly trust the database implementation in TV whitespaces, some kind of identity system and a kill-switch will be required. These can be leveraged to implement spectrum jails as well. The system would be as follows:

The extended database, operating as a spectrum manager, issues “operation tokens,” with a specific time to live, to secondary devices. The operation token gives secondaries permission to transmit in the primary band, but it does *not* guarantee that the primary is not also transmitting. The secondary is responsible for finding its own spectrum holes. This could be done through sensing as a service [117], networked sensing [118], some kind of database service, etc. However, if the secondary fails to correctly find spectrum holes and causes harmful interference to the primary, the next token request will result in the database issuing a “jail token.” While the jail token is active, the secondary will not be allowed to transmit. It must also perform any other sanction activity required, like burning energy.

Can this implementation be trusted? The spectrum managers can be trusted because of long-term relationships with the FCC. The secondaries can be certified to include the necessary hardware and software to implement the jail sanction and token compliance. The tokens could be secured with cryptographic protocols. As long as the sanctions are effective in deterring harmful interference, trust can be achieved. Most of the rest of this thesis is devoted to designing effective sanctions.

Does the jail system fulfill the enforcement goals from Section 1.3.2 and the regulatory goals from Section 1.2.1? This is the question that will be explored throughout the thesis and revisited in the Conclusion.

For now, notice two points that spectrum jails have in their favor: trust of secondary devices is based on certifying compliance with the jail system, and proving that the sanction will be sufficient to deter rational secondaries from cheating. This thesis will show that determining a sufficient sanction does not depend on secondary technology. So, one, transparent certification process can be applied to any secondary device, thereby lowering barriers to entry.

Second, spectrum jails offer an interesting solution to the policy debate between spectrum tribunals and traditional courts from Section 1.1.1. Spectrum tribunals would require significant new infrastructure, but would have the expertise to adequately handle spectrum disputes. Traditional courts are already in place, but might not be able to address spectrum concerns. Spectrum jails are the best of both worlds: they can be implemented through infrastructure that is already in place or already required. They are also able to effectively control harmful interference.

1.3.6 A note on evolvability

Evolvability is one of the regulatory goals from Section 1.2.1 that has not yet been addressed in detail. Thus far in the history of spectrum regulation, evolution from one regulatory scheme to another has been generational. Each new spectrum allocation and each new technology has tweaked the rules. The result is that every band has a slightly different set of rules.

Generational evolution of regulation cannot continue in a shared spectrum paradigm. No longer will networks be separated enough to allow different networks to follow different generations

of the rules. So, rules in shared bands can change in two ways. Either all of the affected networks will have to deploy new devices conforming to the new rules at the same time, or shared spectrum regulation can build in an evolutionary path.

Spectrum managers controlling operation and jail tokens present a very nice evolutionary path. What information is contained in the operation token? These can change depending on primary capability and requirements. Perhaps the primary is willing to locally register its activity in the local database. The operation token can include a basic version with just operation and a more advanced version with the extra primary information. Any secondary from an older generation, or different local area, will use the basic version of the token. But secondaries with a software update to take advantage of the new information will be able to better recover available spectrum holes.

As another example, perhaps secondary devices would like to deploy in an area with no primaries. There is no reason for implementing spectrum jails here. So, operation tokens could be broadcast by satellite to any secondaries in the area. If in the future primaries do come that location, the tokens can be moved to local broadcast and the spectrum jail capabilities can be added.

Because the trust is coming from the spectrum managers holding a leash on the operation of secondary devices, the actual implementation can be customized to the local situation. The jail capabilities can also be rolled out in stages as they are needed. Evolution of regulation can happen on a more natural cycle instead of along hard generational boundaries.

Chapter 2

Basic spectrum jails with a single secondary

This chapter quantitatively investigates whether spectrum jails can provide trust first and better performance as technology improves by considering the simplest possible case.¹ This basic model will be expanded in the next three chapters.

The problem includes three “players:” the single primary which must be protected, the single secondary which is capable of causing interference, and the regulator who designs the jail mechanism. The primary is assumed to be a legacy device that cannot react to its current situation. It is defined only by whether it is transmitting and whether its packets are being dropped. The secondary is assumed to have a utility function that measures total average throughput, which changes depending on how it responds to the primary and to the jail mechanism. The regulator chooses the jail parameters to regulate the interaction between the other two players.

This chapter assumes a worst-case interference relationship, chosen based on the insights from Etkin *et al* [103]. If the primary can cause interference harm to the secondary, the secondary will have incentive to respect its priority. Therefore, the worst case is when the primary causes no interference harm to the secondary, but the secondary can cause interference harm to the primary. The secondary is vulnerable only to punishment through jail.

The regulator designs the parameters of the spectrum jail to produce trust and performance. Trust means that the primary is protected from secondary interference; performance is measured by how well spectrum holes can be recovered. These goals are in direct conflict: the primary can be fully protected with an infinite jail sentence that essentially removes the secondary from the band. On the other hand, spectrum usage could be maximized by letting the secondary always cheat. Because this model assumes the secondary would not experience harm from interference, the band would be fully, productively utilized.

This chapter shows that trust can be guaranteed even while allowing spectrum holes to be recovered. Further, more spectrum holes can be recovered as technology improves.

The key to guaranteeing trust is understanding what are the most difficult situations and what needs to be done to address them.

This thesis assumes that jail parameters are set at a time scale slower than runtime. This can either be at certification time, or during firmware updates. The possibility of runtime-adaptive

¹Many of the ideas in this chapter are published in [15, 119–121].

jail sentences is left to future work because trust should not depend on jails being adaptive; this performance improvement can be added later. If jail sentences are set on a slow time scale, they must protect all runtime realizations of primaries (parameterized by probability of transmission) from all possible secondaries (parameterized by the cost of looking for the primary).²

With this assumption, this chapter shows that there are two worst cases: when there are very few spectrum holes available (because the primary is almost always transmitting), and when there are too many spectrum holes available (rare primary transmissions). If the primary is almost always transmitting, a frequency agile radio should and will naturally choose to operate elsewhere, like in an unlicensed band. This chapter models this choice as a home band, where the secondary does not have to respect primary priority.

When the primary is rarely around, the secondary is choosing between checking for the primary all the time, and going to jail only rarely. In this situation, the regulator must make a choice. It can either embrace a “use it or lose it” approach to primary protection, or it can implement a different mechanism for protecting rarely active primaries.

With a home band and a chosen limit on primary protection, this chapter will show that trust is not only possible, but it is prior to performance! Secondaries can be trusted to not cause interference even if it is not possible for them to recover any spectrum holes without cheating.

The chapter then quantitatively explores how spectrum hole recovery will improve in the future with improved technology, *using the same jail parameters*. As technology becomes perfect, all spectrum holes can be recovered. Furthermore, primaries experience better protection: secondaries will not cheat even if the primary does not meet the minimum usage requirement.

2.1 The model

The implementation of spectrum jails will follow the flowchart in Fig. 2.1. At each step, nature determines whether the primary is transmitting, and the secondary chooses whether to cheat or look for the primary. If the primary is transmitting, the secondary has a possibility of going to jail whether or not it cheats. The operations continue in a loop indefinitely, and all cost functions will involve average long term behavior.

The flowchart of events from Fig. 2.1 is modeled as the Markov chain illustrated in Fig. 2.2. There exists a primary that transitions between transmitting and not transmitting based on some two-state chain. In response to this, the secondary chooses between different actions. When the primary is not transmitting, the secondary can legally transmit.³ The primary and secondary are assumed to be co-located. So, when the primary is transmitting, the sharing rule dictates that the secondary stay silent. It can, however, choose to cheat and transmit anyway with probability P_{cheat} .

If the secondary cheats, it will be sent to jail with probability P_{catch} . If it waits, it may still be sent to jail wrongfully with probability P_{wrong} . Once in jail, it must wait there for an

²Another way to think about the need for universality over primary transmission probability is to look at the problem from the point of view of different runtime time scales. The secondary has some time scale at which it looks for the primary and transmits. This time scale may be different than the one used by the primary for transmission. The primary probability of transmitting will seem different depending on the time scale at which it is observed. The jail sentence must work for any time scale natural to the secondary.

³Fig. 2.2, shows the option of the sensing producing a false alarm. For simplicity, this chapter will assume that the probability of false alarm is zero. This assumption will be revisited in Chapter 4.

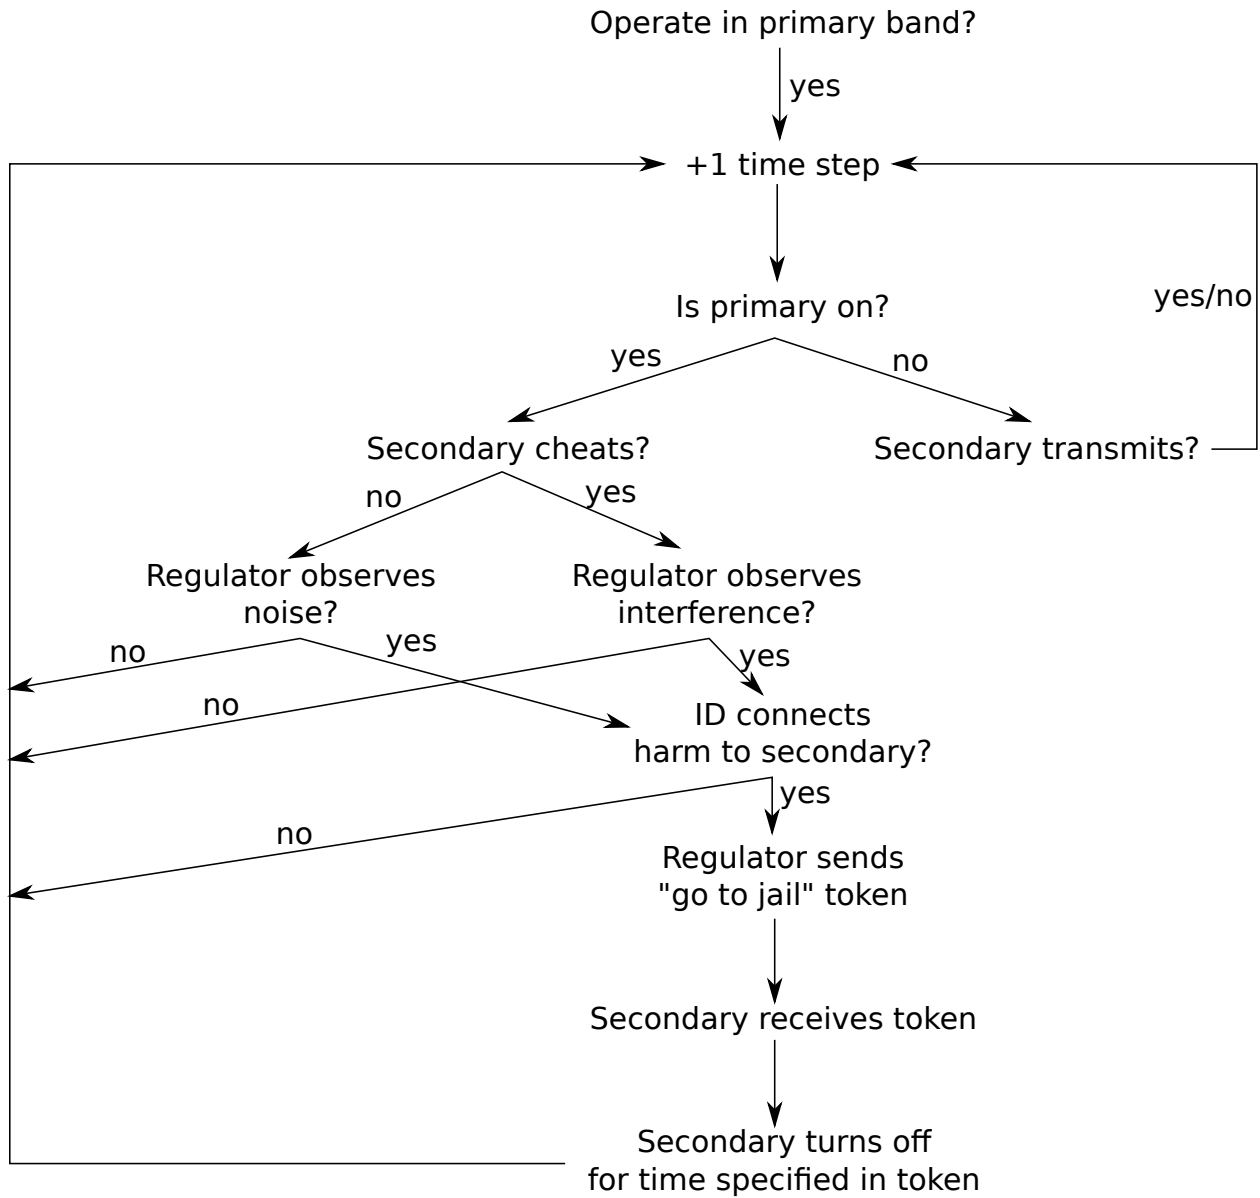


Figure 2.1: Flowchart of events governing the operation and punishment of the secondary.

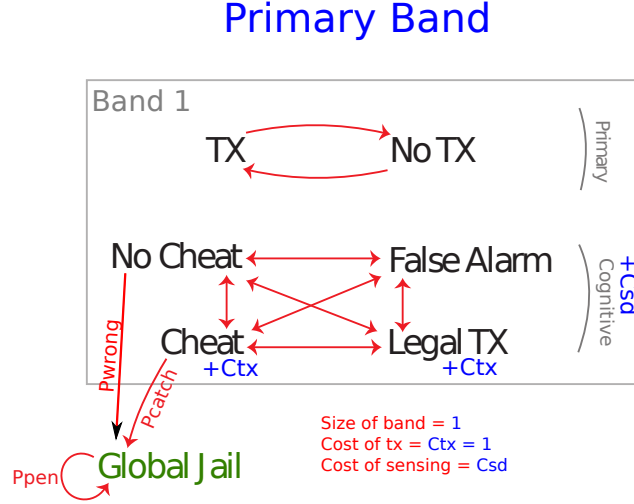


Figure 2.2: Markov chain illustration for the jail system. The primary follows a Markov chain to turn on and off. When the primary is not transmitting, the secondary legally transmits. When the primary is transmitting, the secondary chooses whether to wait, or cheat by transmitting anyway. The secondary is sent to jail with probability P_{catch} if cheating and P_{wrong} if not cheating. Once in jail, the secondary stays there for an amount of time determined by P_{pen} .

amount of time determined by P_{pen} . Once out, the secondary goes back to choosing how to operate in response to the primary.⁴

P_{catch} and P_{wrong} are determined by different things, but both are ultimately determined by the identity system.⁵ P_{catch} is how well the identity system connects an instance of interference to the offending secondary. Wrongful conviction happens when a primary dropped packet due to other sources is attributed to the secondary by the identity system. These dropped packets could be due to noise, or other secondaries causing harm. These parameters will be reconsidered in Chapters 4 and 5. This chapter assumes that because the identity system is known at certification time, these parameters can be known at certification time. They can also improve as technology gets better in the future.

This model can, without much loss in generality, be made even simpler. The primary two-state Markov chain can be thought of as having three parameters: the average probability of transmission and mixing parameters that determine how likely the primary is to stick in either state.

⁴This model treats the jail sentence as a random quantity for simplicity. In reality, the average length of the jail sentence, $1/(1 - P_{pen})$ would be fixed at certification time. Because the cost functions defined here include only average, long term costs, there is no difference between a fixed and probabilistic jail sentence.

⁵The particular form of the identity system is abstracted away in this thesis to performance specifications on catching and wrongful conviction. This is because although identity is known to be important [75], there is no consensus on the form this identity should take. It could be like that in [119, 120, 122], and stamp the identity as a code into the on-off characteristic of the secondary transmission. Identity could also be gleaned from unique features of the transmitting radio [123], or otherwise be introduced as a watermark into the signal [124]. An information-theoretic approach combining identity codes [125, 126] and transmission over unknown channels (see [127] and the references therein) could also be applied. This thesis assumes that such a system exists, and then shows what this system must achieve to be effective for the jail mechanism. With this performance guidance, a good identity system can be developed in future work.




Figure 2.3: If the primary is *iid*, then the Markov chain can be simplified to two states: when the secondary is out of jail and when it is in jail. The movement to jail is determined by whether the primary is on and the secondary is cheating. Moving out of jail still depends only on P_{pen} .

From the perspective of the secondary, a “sticky” primary has times when it is often transmitting, times when it is rarely transmitting, and transitions. The secondary can presumably learn the current state of the primary, so it can track the transitions. As long as the jail sentence works for a wide range of transmission probabilities, then, it is sufficient to assume the primary is *iid* with a transmission probability P_{tx} .

With the *iid* primary assumption, the model can be simplified to that in Fig. 2.3. The secondary now has two states: In Jail, and Not in Jail. It stays in jail according to P_{pen} , and goes to jail with probability P_{toJail} depending on the primary and the secondary’s cheating behavior as given later in Eq. (2.3).

This chapter assumes a secondary that cares only about maximizing overall throughput. So, define its cost function as the average number of steps required to send one message, assuming that the secondary can send one message if it transmits for the entire time step.⁶ The cost function is then:

$$C_D = \frac{1}{\pi_{NJ}(P_{cheat} + (1 - P_{cheat})(1 - P_{tx})(1 - C_{DS}))} \quad (2.1)$$

where the denominator is the number of messages the secondary is able to send per time step on average. Let π_{NJ} be the probability that the secondary is not in jail. C_{DS} is the amount of time per step that the secondary must spend looking for the primary, and can be interpreted as either the time spent sensing or the time spent in consultation with a database. Note that if the secondary chooses to cheat all the time ($P_{cheat} = 1$), then it does not have to spend any time sensing. Given the Markov model with *iid* assumption on P_{tx} :

$$\pi_{NJ} = \frac{1 - P_{pen}}{P_{toJail} + 1 - P_{pen}} \quad (2.2)$$

where the probability of going to jail is

$$P_{toJail} = P_{cheat}(P_{tx}P_{catch}) + (1 - P_{cheat})(P_{tx}P_{wrong}) \quad (2.3)$$

⁶The secondary message is assumed to be arbitrarily split into sections without the penalty of extra headers on subsequent pieces.

Table 2.1: Parameters definitions for the basic spectrum jails problem. The jail sentence is defined based only on those parameters known at certification time, and performance depends on those parameters determined at runtime.

Parameter	Definition	Timescale when used
P_{tx}	Probability the primary is transmitting	Runtime
P_{pen}	Probability of staying in jail	Certification
P_{catch}	Probability cheating secondary is caught	Certification
β	Size of the home band	Certification
P_{txMin}	Minimum protected probability of primary transmission	Certification
C_{DS}	Delay cost of checking for primary	Runtime
P_{wrong}	Probability honest secondary is sent to jail	Runtime

The objective of the secondary is

$$\min_{P_{cheat}} C_D \quad (2.4)$$

All of the parameters and definitions are included in Table 2.1.

2.2 Basic operation

Fig. 2.4 shows the secondary cost function for different values of the other parameters. In these examples, the best choice for the secondary is either to never cheat, or to always cheat. With this cost function, these are the only two rational choices, regardless of the parameters:

Theorem 1. *The rational strategy for the secondary with the cost function in Eq. (2.1) is to choose $P_{cheat} \in \{0, 1\}$*

Proof. Lemma 1 can be applied because with $x = P_{cheat}$, it has the correct form and the denominator is always positive. Therefore, C_D is monotonic in P_{cheat} . The result follows. \square

Lemma 1. *Any function $f(x)$ of the form*

$$f(x) = \frac{Ax + B}{Cx + D} \quad (2.5)$$

for any A, B, C, D , is monotonic in x over the range $x \in [x_1, x_2]$ if $-D/C \notin [x_1, x_2]$.

Proof. The condition $-D/C \notin [x_1, x_2]$ guarantees that there are no discontinuities in the range of interest. Then, taking the derivative of $f(x)$:

$$\frac{df}{dx} = \frac{ACx + DA - ACx - BC}{(Cx + D)^2} \quad (2.6)$$

$$= \frac{DA - BC}{(Cx + D)^2} \quad (2.7)$$

The sign of the derivative depends on $DA - BC$ and cannot change due to changes in x . Therefore, $f(x)$ is monotonic in x over the range $x \in [x_1, x_2]$. \square




Figure 2.4: Example costs for the secondary given the level of cheating, with two different values of the sanction parameter P_{pen} . The cost is always minimized by $P_{cheat} = 0$ or 1 . This is true in general: the secondary will never choose to cheat only part of the time.

Because the secondary will either always cheat or never cheat, Figs. 2.5 and 2.6 show the sanction required to guarantee that $\text{argmin}_{P_{cheat}} C_D = 0$ for different values of C_{DS} and P_{tx} . As more time is spent sensing, the necessary sanction rises as well. For a fixed jail sentence, sensing is getting expensive while the cost of cheating and going to jail is remaining the same. So, the jail sentence must rise accordingly.

Fig. 2.6 is likewise intuitive. When the primary is transmitting most of the time, there is little opportunity for honest use. If the jail sentence remains fixed as P_{tx} rises, the cost of jail is going down relative to the cost of waiting for the primary to turn off. When the primary is very rarely transmitting, the secondary is choosing between paying a sensing cost all the time and going to jail only very rarely when it actually causes harm.⁷ As the chance of getting sent to jail goes down, the sanction must go up for the overall cost of jail to remain higher than the cost of sensing.

These figures show that with the current model, trust is strongly tied to secondary technology. The next theorem makes the dependence explicit.

Theorem 2. *With C_D as written in Eq. (2.1), there does not exist a value of P_{pen} such that secondaries will be honest for all C_{DS} and P_{tx} .*

Proof. Using Thm. 1, $P_{cheat} \in \{0, 1\}$. Therefore, the cost function in (2.1) can be split into two

⁷The model in this chapter makes no distinction between when the primary is transmitting messages and when the band is being defended against secondary transmission. The secondary can therefore only be sent to jail when it causes a primary packet to be dropped. However, these are not necessarily the same thing. Chapter 5 will allow the primary to police the band even when it is not the one transmitting the protected message.

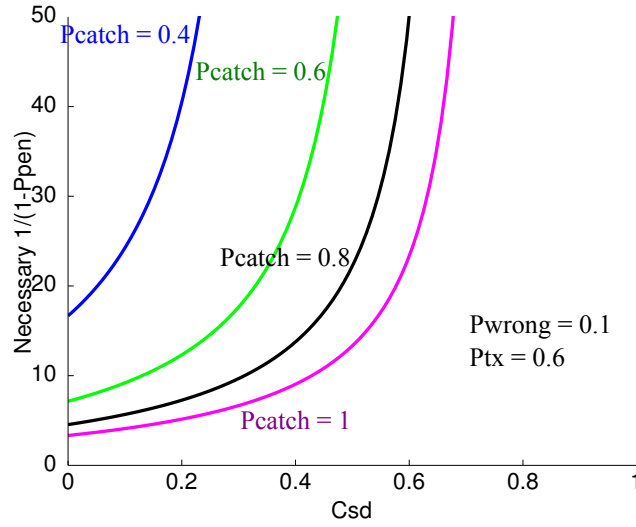


Figure 2.5: The required average sanction (amount of time in jail per offense) to deter cheating for different values of the cost of sensing. Different lines indicate different probabilities that the primary will be active. As the cost of sensing rises, the necessary sanction rises without bound, indicating there does not exist a sanction that will work for all costs of sensing.




Figure 2.6: The required average jail time per offense to deter cheating vs the probability of the primary being active. Different lines represent different costs of sensing. When the primary is mostly active, there is little opportunity to transmit, so the secondary is more tempted to cheat. When the primary is rarely active, the secondary would rather go to jail every once in a while than sense at every time step.

cases. When cheating:

$$C_{D,cheat} = C_D(P_{cheat} = 1) = \frac{1}{\pi_{NJ}} \quad (2.8)$$

$$= \frac{P_{tx}P_{catch} + 1 - P_{pen}}{(1 - P_{pen})}. \quad (2.9)$$

substituting in π_{NJ} from Eq. (2.2). When not cheating:

$$C_{D,honest} = C_D(P_{cheat} = 0) = \frac{1}{\pi_{NJ}(1 - P_{tx})(1 - C_{DS})} \quad (2.10)$$

$$= \frac{P_{tx}P_{wrong} + 1 - P_{pen}}{(1 - P_{pen})(1 - P_{tx})(1 - C_{DS})}. \quad (2.11)$$

Comparing the two, the required sanction P_{pen} to guarantee the secondary will not cheat is

$$C_{D,cheat} \leq C_{D,honest} \quad (2.12)$$

$$\Rightarrow P_{pen} \geq 1 - \frac{P_{tx}(P_{catch}(1 - P_{tx})(1 - C_{DS}) - P_{wrong})}{1 - (1 - P_{tx})(1 - C_{DS})} \quad (2.13)$$

P_{pen} cannot be greater than 1. So, there exist situations when no sanction is sufficient to deter cheating:

$$1 > 1 - \frac{P_{tx}(P_{catch}(1 - P_{tx})(1 - C_{DS}) - P_{wrong})}{1 - (1 - P_{tx})(1 - C_{DS})} \quad (2.14)$$

$$\Rightarrow \frac{P_{wrong}}{P_{catch}} < (1 - P_{tx})(1 - C_{DS}) \quad (2.15)$$

For every combination of $P_{wrong} > 0$, $P_{catch} > 0$, there exists a (P_{tx}, C_{DS}) such that it is not possible to find an appropriate sanction. \square

At first glance, this seems like a negative result – the goal of building spectrum jails was to remove any dependence between the jail sanction and the secondary technology. But what is the real problem here?

Jails are being used as a rationing mechanism. When the primary is almost always around, there are very few spectrum holes. These primaries should not have to share at all, and certainly should not have to deal with cheating secondaries. Likewise, secondaries that cannot effectively sense for a primary should not be trying to fill spectrum holes.

If jails are the only mechanism to distribute the scarce resource of easily recoverable spectrum holes, they will of course need to depend on the current availability of the resource. However, jails do not actually need to be a rationing mechanism. Secondaries are already frequency-agile devices with the capability to learn. They will be able to switch to a different band, like an unlicensed band, if it is too hard to share with a primary.

In the next part, this is modeled as a *home band* – a band where the secondary can transmit without having to check for a primary. The home band could be an unlicensed band. It could also be a band where that device has primary rights. The key is that the secondary is using the spectrum holes as a *band expander* to get more utility beyond its already-available band.

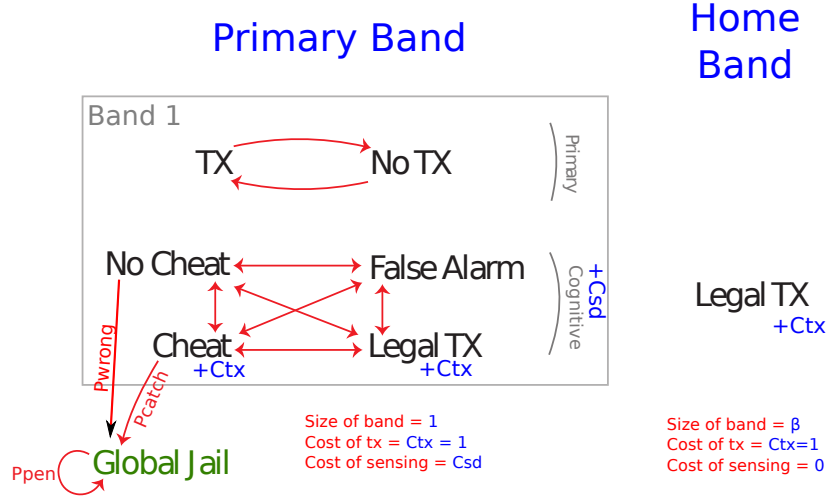


Figure 2.7: With a home band, the operation in the primary band is the same as before. In the home band, the secondary does not have to respect a primary's priority, and so its only option is legal transmission. The secondary chooses its operation mode: operate just in the primary band, just in the home band, or in both. The jail is a global jail – the secondary loses transmission ability in all bands, including the home band. It is assumed that if operating only in the home band, the secondary is not subject to jail.

The regulator, then, must either force all cognitive devices to have a paid allocation, or it must guarantee there are available unlicensed bands. These bands can then be rationed, if needed, through more traditional means, like pricing.⁸

2.2.1 Guaranteeing trust with a home band

The new model is illustrated in Fig. 2.7. The secondary can choose where to operate: just the primary band, just the home band, or both simultaneously. If the secondary is operating in the primary band or both, it can be sent to jail. While in jail, it must cease transmission in all bands it occupies.⁹

If the secondary commits to operating just in the home band, it is not subject to jail. Therefore, if jail is too expensive, because of high sanctions or high wrongful conviction rates, the secondary will have incentive to operate in the home band alone.¹⁰

The new secondary objective includes choosing the minimum over three cost functions and operating modes:

⁸Note that for a pricing mechanism to work, the secondaries must first be trusted to actually follow the results of the pricing mechanism. Adapting spectrum jails to work for rationing unlicensed bands is left for future work.

⁹If the secondary is a database provider or other secondary market facilitator, the home band could be a geographical region with the same frequency as the primary bands. So, the secondary has its base operation with a dedicated band in Berkeley, for example, but can expand its reach to other parts of the country on an opportunistic basis. Global jail would then require the database provider to shut off access to its bands in all geographical areas.

¹⁰Being able to trust a decision at runtime to operation just in the home band may be a hard problem. Section 2.4.2 will discuss this problem. It shows that even if the home band is subject to jail, keeping P_{wrong} low enough during just home band operation will still allow trust.

1. Use just the primary band:

$$C_{D,1} = \frac{1}{\pi_{NJ}(P_{cheat} + (1 - P_{cheat})(1 - P_{tx})(1 - C_{DS}))} \quad (2.16)$$

2. Use both home band and primary band:

$$C_{D,2} = \frac{1}{\pi_{NJ}(P_{cheat} + (1 - P_{cheat})(1 - P_{tx})(1 - C_{DS}) + \beta)} \quad (2.17)$$

3. Use just the home band:

$$C_{D,3} = \frac{1}{\beta} \quad (2.18)$$

Then, the overall cost for the secondary is:

$$C_D = \min_i C_{D,i} \quad (2.19)$$

and the objective remains $\min_{P_{cheat}} C_D$.

With this new model, the secondary will still choose to either always cheat or never cheat.

Theorem 3. *With a home band, $P_{cheat} \in \{0, 1\}$.*

Proof. $C_{D,2} \leq C_{D,1}$ for all values of all parameters, given $\beta \geq 0$, so it is enough to consider just $C_{D,2}$ and $C_{D,3}$. $C_{D,2}$ has the form required by Lemma 1, so it is monotonic in P_{cheat} . Therefore, the secondary will never cheat, always cheat, or always sit in the home band. \square

With the home band, the required sanction to keep a secondary from cheating is the minimum of two values: the sanction required to keep the secondary honest, and the sanction required to make the secondary choose to operate in the home band alone.

Fig. 2.8 shows the sanction required to deter cheating over the cost of sensing. With a home band and a high enough sanction, all secondaries can be deterred from cheating. Furthermore, the sanction for universal deterrence is determined by the value of P_{pen} required for the secondary to choose home band operation over cheating.

Also notice that the sanction required to not cheat is different than it was without the home band: the secondary will always choose to use the home band along with the primary band. When it is losing too much home band utility due to jail, it will stop cheating to preserve its home band. Therefore, even if the secondary cannot choose to move to the home band entirely, cheating can be controlled with a reasonable sanction.

Theorem 4. *The necessary sanction to guarantee that no secondary will cheat given a particular P_{tx} is:*

$$P_{pen} > 1 - P_{tx}P_{catch}\beta \quad (2.20)$$

Proof. It is sufficient to guarantee that it is always less expensive to go to the home band than to always cheat. Therefore, the condition is:

$$C_{D,3} \leq C_{D,2}(P_{cheat} = 1) \quad (2.21)$$

The result follows from solving this equation for P_{pen} . \square




Figure 2.8: Required sanctions vs the cost of sensing. The bottom black line is the sanction required to deter cheating. The red line is the sanction required to make the secondary move to the home band. To deter cheating, the sanction must be the minimum of these two lines (dotted green line). Notice that it is possible to choose a level of sanction that will work for all values of the cost of sensing, and it does not go to an infinite sanction as the cost of sensing goes to 1. Even if the secondary is operating in the primary band, its behavior is modified by being able to transmit in the home band as well. When the cost of sensing is very high, the secondary still has incentive to not cheat in order to protect its operation in the home band.

Fig. 2.9 shows the sanction required vs P_{tx} for a particular value of the cost of sensing. Notice that at high P_{tx} , the sanction is no longer unbounded: the secondary will choose to go to the home band if there is not enough transmit opportunity in the primary band. However, at low values of P_{tx} , the required jail time still goes to infinity. This happens for the same reason it did before – the secondary is choosing between sensing all the time and getting sent to jail only very rarely when the primary is present. This is a problem of the temptation to cheat being too great, not a problem of the legal options being too scarce. The home band is not the right patch to compensate.

At this point, regulators must question what uses of spectrum are most valued. Perhaps there should be a “use it or lose it” rule which states that a primary will be protected from cheating secondaries only if it uses the band a sufficient percentage of the time.

However, “use it or lose it” is certainly not the right rule for all situations. For example, public safety users do not utilize the band most of the time. However, in an emergency, their transmissions must be protected.

There are a few options. If the primary can wait a few seconds to minutes before the packets must get through, the problem could take care of itself: presumably, the secondaries are learning P_{tx} to make their cheating decisions because this value will not be declared beforehand. Public safety users would be absent for large blocks of time but would have very high P_{tx} for the periods when they are active. If they can tolerate learning time, the secondaries will likely vacate the band as soon as they realize the primary is now transmitting constantly.

If the low P_{tx} use requires immediate uninterrupted access, the primary could be given a go-to-jail beacon that would ban all secondary devices to their home bands for the duration of the emergency. If beaconing is not an option, and the use is important enough, these primaries *must* have their own band. Regulators must pay to protect these important users with wasted spectrum.¹¹

Given such a minimum protected P_{tx} , or P_{txMin} , there is a sanction that works for all cases of interest:

Theorem 5. *Given a minimum protected P_{tx} , P_{txMin} , the required sanction to protect all other P_{tx} from any secondary realization is:*

$$P_{pen} > 1 - P_{txMin}P_{catch}\beta \quad (2.22)$$

Proof. Take Thm. 4 and maximize over $P_{tx} \in [P_{txMin}, 1]$. □

At this point, trust that the secondary will not cheat has been achieved. Notice Thm. 5; the jail sanction depends on three values: $P_{catch}, P_{txMin}, \beta$. The regulator knows the capability of its identity system, and it can choose a protection level for the primary. It also knows what kind of home band a secondary has access to, and can determine its size.¹² This sanction can therefore

¹¹This point will be revisited in Chapter 5. If it is possible for the primary to subcontract its band to a secondary it trusts, then the primary’s band can have higher utilization without sacrificing protection.

¹²If an unlicensed band is being used as the home band, knowing the size of the home band requires the regulator to either guarantee a fixed congestion level on the unlicensed band, or it requires the regulator to monitor the congestion. If it is simply monitoring the congestion, the jail sanction will need to change to reflect times of high unlicensed band congestion (translating to low home band size). This can presumably be done on a slower time scale. The change could be propagated through the database, as the suggested implementation for jails will include an enhanced database/spectrum manager.




Figure 2.9: The sanction required vs the probability the primary is active. The bottom black line is the sanction required to deter cheating. The red line is the sanction to move the secondary to the home band. The required sanction is the minimum of these two lines (green dotted line). Notice that there is no universal sanction possible. Regulators must define a minimum protected P_{tx} . Then, there exists a sanction that protects all primaries operational for more than P_{txMin} of the time. In this plot, there are two points to compare when choosing the sanction: the sanction needed at the minimum P_{tx} , and the sanction needed when the secondary will switch to the home band at high P_{tx} . This is a result of the chosen cost of sensing. Taking the maximum of these points over all costs of sensing leaves only the sanction needed at P_{txMin} .




Figure 2.10: The primary transmission probability below which cheating actually occurs, given P_{txMin} and cost of sensing. If it is possible to improve sensing technology, the primary will be better protected.

be set at certification time and does not depend on the secondary spectrum access technology or the primary usage characteristic. The only thing that needs to be certified is that the secondary will receive go-to-jail commands and follow its sanction appropriately.

Now, the question is performance. The next section will give metrics for performance and show how these improve as technology improves, even with the same sanction set at certification time.

2.3 Performance improvement with better technology

2.3.1 Performance metric definitions

Performance is measured for both primary and secondary: the primary is not guaranteed protection below P_{txMin} , so its performance metric is interference. The secondary will spend time in jail wrongfully, and will retreat to the home band if the primary band is inhospitable. So, it has metrics of overhead (or inefficiency) due to jail and spectrum holes not recovered due to moving to the home band.

For the primary, the appropriate metric is its probability of a dropped packet. This will be referred to as a collision:

$$\text{collision} = \pi_{NJ} P_{cheat} \mathbb{1}_{\text{in primary band}} \quad (2.23)$$

and will only happen when the secondary is in the primary band, not in jail, and cheating.

When will the secondary cheat? The sanction is set such that primaries will be protected as long as $P_{tx} > P_{txMin}$. However, secondaries are not guaranteed to cheat below P_{txMin} . They will start cheating depending on the sanction and their cost of sensing:

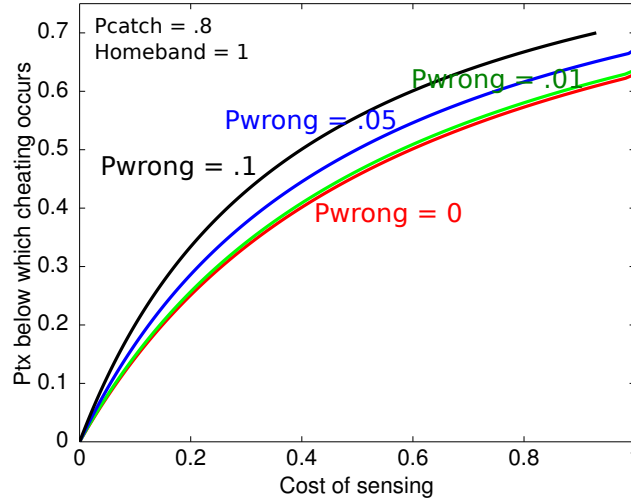


Figure 2.11: The bounding curve from Fig. 2.10 for different values of P_{wrong} . Better wrongful conviction rates will also produce better realized protection for the primary.

Theorem 6. *The P_{tx} below which the secondary will start cheating is*

$$P_{txcutoff} = \min(P_{txMin}, \tilde{P}_{tx}) \quad (2.24)$$

where \tilde{P}_{tx} is the solution to:

$$1 - P_{pen} = \frac{\tilde{P}_{tx} P_{catch} ((1 - \tilde{P}_{tx})(1 - C_{DS}) + \beta) - \tilde{P}_{tx} P_{wrong} (1 + \beta)}{1 - (1 - \tilde{P}_{tx})(1 - C_{DS})} \quad (2.25)$$

Proof. By construction, the secondary would rather operate only in the home band than cheat if and only if $P_{tx} > P_{txMin}$. Also, as established earlier, the secondary will always use the home band, with or without operating in the primary band as well. Therefore, only $P_{tx} < P_{txMin}$ is of concern. Compare operating honestly to cheating when operating in both cognitive and home band. Solving for $1 - P_{pen}$ in the following equation gives the result:

$$C_{D,2}(P_{cheat} = 1) > C_{D,2}(P_{cheat} = 0) \quad (2.26)$$

□

The actual cutoff is shown in Fig. 2.10 for different levels of P_{txMin} and sensing cost. Fig. 2.11 shows the bounding curve from Fig. 2.10 for different levels of P_{wrong} . This bounding curve is when $\tilde{P}_{tx} = P_{txMin}$ in Thm. 6. Improving the sensing cost or wrongful conviction will both improve the realized primary protection level, for any P_{txMin} .

This actual cutoff P_{tx} is a penalty of universality. To keep the overhead of the honest secondaries low, there must be a limit on the primary usage the sanction will defend. If the sanction could be tailored to the particular situation (and further to the particular secondary device), then even primaries at $P_{tx} < P_{txMin}$ could be protected. This issue will be discussed in more detail in Section 2.4.

For the secondary, the first metric is overhead due to jail, which is measured relative to a secondary that will always follow sharing rules but is not subject to jail sentences.

Without jail, an honest secondary has cost function:

$$C_{D,noJail} = \min_i C_{D,noJail,i} = C_{C,noJail,2} \quad (2.27)$$

where

$$C_{D,noJail,1} = \frac{1}{(1 - P_{tx})(1 - C_{DS})} \quad (2.28)$$

$$C_{D,noJail,2} = \frac{1}{(1 - P_{tx})(1 - C_{DS}) + \beta} \quad (2.29)$$

$$C_{D,noJail,3} = \frac{1}{\beta} \quad (2.30)$$

When not facing jail, the honest secondary will always use both bands because it wants as much transmit time as possible.

Defining the overhead due to jail is complicated by the choice between different operating modes. So to simplify, define the secondary overhead *only within the primary band*:

$$O_S = \frac{U_{D,noJail} - U_D}{U_{D,noJail}} \quad (2.31)$$

$$= 1 - \frac{C_{D,noJail,1}}{C_{D,1}} \quad (2.32)$$

where $U_D = 1/C_{D,1}$ is the utility of a secondary with the jail system and $U_{D,noJail} = 1/C_{D,noJail,1}$ is the utility of the secondary without jail.

Theorem 7. *The honest secondary's overhead relative to an honest user not subject to jail is*

$$O_S = 1 - \pi_{NJ} \quad (2.33)$$

or the time the secondary spends wrongfully in jail.

Proof. Plug in definitions of $C_{D,1}$ and $C_{D,noJail,1}$ into the definition of O_S . \square

An honest secondary spends time in jail only with wrongful conviction, so this overhead will be controlled by P_{wrong} .

From the regulator's perspective, the inefficiency due to jail is really the spectrum holes left unrecovered because of the jail system. This is captured in two metrics: spectrum holes remaining due to wrongful conviction, and spectrum holes remaining due to retreating to the home band.

$$\text{holes}_{jail} = (1 - \pi_{NJ})(1 - P_{tx})\mathbb{1}_{\text{in primary band}} \quad (2.34)$$

$$\text{holes}_{homeband} = (1 - P_{tx})\mathbb{1}_{\text{in home band}} \quad (2.35)$$

There will be a spectrum hole due to jail whenever the secondary is in the primary band, the primary is not transmitting, and the secondary is in jail. There will be a spectrum hole due to retreat to the home band whenever the secondary is in the home band and the primary is not transmitting.

These can be put together into a total remaining spectrum holes metric:

$$\text{holes}_{total} = \text{holes}_{jail} + \text{holes}_{homeband} \quad (2.36)$$




Figure 2.12: Overhead metric values for different values of the parameters. The different plots use the indicated P_{txMin} and P_{tx} . Overhead is shown as a function of C_{DS} . Black solid lines are secondary overhead, green dotted lines are primary interference experienced, red dashed lines are holes due to jail, magenta solid lines are holes due to retreat to the home band. Notice that the primary is protected for only some values of C_{DS} when $P_{tx} < P_{txMin}$. Also, the lower the P_{txMin} , the higher the sanction. So, the secondary will move to the home band for smaller levels of C_{DS} .

2.3.2 Exploring the metrics

Fig. 2.12 shows the overhead metrics. Each plot is a different value of P_{tx} and P_{txMin} : P_{txMin} grows in the horizontal positive direction and the realization of P_{tx} grows downward vertically. The individual plots show how the metrics change with C_{DS} . All of the metrics exhibit a threshold behavior against C_{DS} because the sensing cost determines only the operating mode of the secondary. It does not affect the amount of time in jail.

The threshold trends in different plots are as one would expect: for high P_{txMin} and low P_{tx} , the secondary will cheat for a larger range of sensing cost because the cost of jail is lower and the temptation is higher. On the other extreme, for low P_{txMin} and high P_{tx} , the secondary will retreat to the home band much sooner because the sanction, and therefore the cost of wrongful conviction, is high. At the same time, the opportunity in the primary band is small, so the cost of wrongful conviction is outweighing the benefit of extra transmit time.




Figure 2.13: Overhead metrics vs P_{tx} . Different plots are for the indicated values of P_{txMin} , C_{DS} , and all plots are against P_{tx} . Black solid lines are secondary overhead, green dotted lines are primary interference experienced, red dashed lines are holes due to jail, magenta solid lines are holes due to retreat to the home band. Again here, the secondary moves to the home band for lower values of P_{tx} when P_{txMin} and therefore the sanction are higher. The primary is not protected when $P_{tx} < P_{txMin}$, but has fewer collisions for higher P_{tx} because the secondary is spending more time in jail. Secondary overhead gets worse for higher P_{tx} for the same reason – it is spending more time wrongfully in jail.

Fig. 2.13 shows the same metrics against P_{tx} for different values of P_{txMin} and C_{DS} . P_{txMin} grows horizontally to the right, and C_{DS} grows vertically downward. These plots are also intuitive: for high P_{txMin} and high C_{DS} , the secondary cheats for a larger range of P_{tx} because the sanction is very low compared to the cost of sensing. For very low P_{txMin} , the secondary is driven out of the band even with very small P_{tx} since the cost of wrongful convictions is too high.

The holes metric follows $1 - P_{tx}$ because those are the holes that would have been filled by a secondary but are not because the secondary is in the home band. The best case here is high P_{txMin} and very low C_{DS} . When the cost of sensing is very small, even a small sanction is enough to deter cheating, but the secondary will not retreat to the home band unless P_{tx} is very high.

The lines for holes due to jail and secondary overhead follow each other closely when the secondary is not cheating – their only difference is the $1 - P_{tx}$ term in the holes metric. When the rational secondary is cheating, the holes due to jail metric goes up because a cheating secondary will spend much more time in jail than an honest one.

These figures give a good intuition for how these metrics behave, but cannot easily show dependence on different parameters. Fig. 2.14 gives a summary view given a choice of P_{txMin} .

Each figure is a different value of P_{wrong} , and the lines are calculated by averaging each metric over P_{tx} and C_{DS} . The holes metrics are active only for part of this range, so they are calculated as having a uniform distribution over the active range.

With the option of a home band, the sensing cost splits users into two camps: those that are sufficiently technically mature to take advantage of the spectrum holes in the primary band, and those that are not. The regulator may want to focus on the performance for a particular targeted set of secondaries. Therefore the holes due to transition to the home band includes several lines. Each is calculated only on a specified range $C_{DS} \in [0, C_{SD,max}]$.

The tradeoff between the metrics can be seen even more directly in Fig. 2.15, which shows the total holes metric vs the collisions metric for different levels of sanction and different P_{wrong} . The regulator can use this plot to inform its decision of what primaries to protect based on how easily the spectrum holes may be recovered.

The dots show different possible choices for a fair trade between primary collisions and filled spectrum holes. Red dots show the min sum solution. Blue dots fix the total holes at .01 and do the best possible in terms of primary collisions. Green dots are like the previous definition of P_{txMin} : the primary collisions are fixed and then the regulator does the best possible in terms of filling spectrum holes.

Finally, Fig. 2.16, shows the same solution points as a function of P_{wrong} . This figure also shows the P_{txMin} required to achieve these points. If the regulator can produce a better P_{wrong} , its choice becomes much simpler.

The tradeoff between different metrics is now very clear – smaller P_{txMin} means that the primary is more protected, but the secondary will have higher overhead and the regulator will see a worse usage efficiency. But throughout, the metrics depend very strongly on the wrongful conviction rate.

When spectrum jail systems are first rolled out, perhaps the best identity system will include sending all cognitive radios in the area to jail whenever any interference is detected. This will catch interferers well, but wrongful conviction rate will be very high. However, as technology improves, it may be possible to identify guilty individuals, and the wrongful conviction rate will fall. The plots so far have indicated that performance will improve as wrongful conviction rate improves. The next part makes this notion precise.




Figure 2.14: Summary overhead metrics against P_{txMin} . The values are averaged over the possible realizations of primary and secondary, so over P_{tx} and C_{DS} . The two plots have different values of P_{wrong} , .1, .01. All measures of overhead get better with smaller P_{wrong} .




Figure 2.15: If the regulator chooses to change the sanction to balance the needs of the different parties, it can choose it based on the Pareto-like curve here. This is the value of the probability of primary collision vs the total spectrum holes. The curve represents different values of P_{txMin} . The circles are different choices of balancing – red is the min sum solution, green fixes the total collisions, and blue fixes the total spectrum holes. Different lines are different values of P_{wrong} .




Figure 2.16: Given the three possible fairness balances (min sum, fix collisions, fix total holes), this plot shows the chosen P_{txMin} , average collisions, and average spectrum holes against P_{wrong} . Notice the dependence of the optimal P_{txMin} on P_{wrong} . If fixing the total holes metric, the best P_{txMin} is highly dependent on P_{wrong} , and so using this approach would require the ability to change the sanction as the technology for catching cheaters changed.

2.3.3 Improving performance in the future

Trust has already been guaranteed, and the last part hinted that improving technology would improve performance. This could happen in two ways: as the catching and identity systems get better, cheaters will be found with greater fidelity, and innocents will be less likely to be wrongfully convicted. Because it has a greater effect, this section will ignore the change in P_{catch} and focus on the wrongful convictions.

Technology will also improve how spectrum holes are discovered. Either with better databases or networked sensing solutions, the effect will be to reduce the sensing cost C_{DS} .

Fig. 2.17 shows the effect of reducing these two parameters. P_{wrong} gets smaller horizontally, while C_{DS} gets smaller vertically. The different colors show who is using the band what fraction of the time. Blue is primary use; green is honest secondary use; black is cheating secondary use; and red is collisions where the secondary is cheating and wiping out the primary transmission. The cheating use does not fill all spectrum holes because the secondary spends time in jail. Likewise, the honest use does not fill the band because of time spent on sensing and wrongful convictions. When there is no secondary use, the secondary has moved to the home band.

The dotted black line shows the guaranteed protection level for the primary, which is setting the sanction. As P_{wrong}, C_{DS} get smaller, performance improves! The secondary cheats only for smaller values of P_{tx} , and can honestly fill spectrum holes more efficiently. When both are zero, all spectrum holes are filled by honest secondary use. The following theorem makes this precise.

Theorem 8. *Performance improves with technology in the following ways:*

1. $\lim_{P_{wrong}, C_{DS} \rightarrow 0} \text{collision} = 0 \forall P_{tx}$ where collisions are defined in Eq. (2.23).
2. $\lim_{P_{wrong}, C_{DS} \rightarrow 0} \text{holes}_{jail} = 0 \forall P_{tx}$ where the metric is defined in Eq. (2.34)
3. $\lim_{P_{wrong}, C_{DS} \rightarrow 0} \text{holes}_{homeband} = 0 \forall P_{tx}$ where the metric is defined in Eq. (2.35)

Proof. This proof treats each metric separately:

1. The collisions metric is defined as $\text{collision} = \pi_{NJ} P_{cheat} \mathbb{1}_{\text{in primary band}}$, and the sanction is defined to guarantee no cheating for $P_{tx} > P_{txMin}$. So it is sufficient to show that the secondary will choose $P_{cheat} = 1$ for smaller values of P_{tx} as P_{wrong}, C_{DS} improve. As in Thm. 6, the secondary will choose to cheat for any $P_{tx} < P_{txMin}$ such that

$$1 - P_{pen} > \frac{P_{tx} P_{catch} ((1 - P_{tx})(1 - C_{DS}) + \beta) - P_{tx} P_{wrong} (1 + \beta)}{1 - (1 - P_{tx})(1 - C_{DS})} \quad (2.37)$$

Rearranging, and plugging in the fixed sanction $P_{pen} = 1 - P_{txMin} P_{catch} \beta$ gives the following:

$$1 - P_{pen} > \frac{P_{tx} P_{catch} ((1 - P_{tx})(1 - C_{DS}) + \beta) - P_{tx} P_{wrong} (1 + \beta)}{1 - (1 - P_{tx})(1 - C_{DS})} \quad (2.38)$$

$$0 > P_{tx} P_{catch} ((1 - P_{tx})(1 - C_{DS}) + \beta) - P_{tx} P_{wrong} (1 + \beta) \quad (2.39)$$

$$- (1 - P_{pen})(1 - (1 - P_{tx})(1 - C_{DS})) \quad (2.40)$$

$$= -P_{wrong} (1 + \beta) P_{tx} + -C_{DS} (1 - P_{tx})(P_{tx} P_{catch} + 1 - P_{pen}) \quad (2.41)$$

$$+ P_{tx} P_{catch} (1 - P_{tx} + \beta) + -P_{tx} (1 - P_{pen}) \quad (2.42)$$

$$= -P_{wrong} (1 + \beta) P_{tx} + -C_{DS} (1 - P_{tx}) P_{catch} (P_{tx} + P_{txMin} \beta) \quad (2.43)$$

$$+ P_{tx} P_{catch} (\beta(1 - P_{txMin}) + 1 - P_{tx}) \quad (2.44)$$




Figure 2.17: Band usage for different performance parameters. Blue is used by the primary, green is honest use by the secondary, black shows cheating use by the secondary, and red shows collisions between primary and secondary. The black area is curved because of time spent in jail. The green is hindered by wrongful convictions and sensing time. When there is no green filling in holes above the blue line, the secondary has retreated to the home band. As P_{wrong} and C_{DS} improve (get smaller), there are fewer collisions, the honest use more effectively fills the band, and the secondary retreats to the home band only for larger values of P_{tx} .

Notice that as P_{wrong} and C_{DS} get smaller, the right side negative terms get smaller, so the condition will not be satisfied even with smaller values of P_{tx} . This means that the secondary will choose not to cheat even for smaller values of P_{tx} . As $P_{wrong}, C_{DS} \rightarrow 0$, the right side becomes positive unless $P_{tx} = 0$, at which point there cannot be collisions. So, $\lim_{P_{wrong}, C_{DS} \rightarrow 0} \text{collision} = 0 \forall P_{tx}$.

2. The holes due to jail metric is defined as $\text{holes}_{jail} = (1 - \pi_{NJ})(1 - P_{tx})\mathbb{1}_{\text{in primary band}}$, so it is sufficient to show that $\pi_{NJ} \rightarrow 1$ as $P_{wrong}, C_{DS} \rightarrow 0$ when the secondary is being honest. Plugging in the definition of π_{NJ} from Eq. (2.2), and setting the sanction as $P_{pen} = 1 - P_{txMin}P_{catch}\beta$:

$$\pi_{NJ} = \frac{1 - P_{pen}}{P_{toJail} + 1 - P_{pen}} \quad (2.45)$$

$$= \frac{P_{txMin}P_{catch}\beta}{P_{tx}P_{wrong} + P_{txMin}P_{catch}\beta} \quad (2.46)$$

Clearly, π_{NJ} gets bigger as P_{wrong} gets smaller, and $\lim_{P_{wrong} \rightarrow 0} \pi_{NJ} = 1$. There is no dependence on C_{DS} .

3. The holes due to retreat to the home band metric is defined as $\text{holes}_{homeband} = (1 - P_{tx})\mathbb{1}_{\text{in home band}}$. It is sufficient to show that the secondary will not retreat to the home band for any P_{tx} . The secondary will operate only in the home band if the cost of the home band is less than the cost of honest use in both bands, where costs are defined in Eq. (2.19):

$$C_{D,3} < C_{D,2}(P_{cheat} = 0) \quad (2.47)$$

$$\frac{1}{\beta} < \frac{1}{\pi_{NJ}((1 - P_{tx})(1 - C_{DS}) + \beta)} \quad (2.48)$$

$$\frac{1}{\beta} < \frac{P_{tx}P_{wrong} + 1 - P_{pen}}{(1 - P_{pen})((1 - P_{tx})(1 - C_{DS}) + \beta)} \quad (2.49)$$

$$\frac{(1 - P_{pen})(1 - P_{tx})}{P_{tx}} < \frac{P_{wrong}}{1 - C_{DS}} \quad (2.50)$$

As P_{wrong} and C_{DS} get smaller, this condition will be violated even by larger values of P_{tx} . And in the limit as $P_{wrong}, C_{DS} \rightarrow 0$, this condition becomes

$$\frac{(1 - P_{pen})(1 - P_{tx})}{P_{tx}} < 0 \quad (2.51)$$

which is violated only when $P_{tx} = 1$ at which point there are no holes left in the primary band to fill. □

So, even with the sanction set conservatively at certification time, the improvement in runtime parameters for wrongful conviction and sensing cost will reduce and possibly eliminate regulatory overhead.

2.4 Concluding remarks

2.4.1 What we have learned so far

This chapter has used a simple model for spectrum jails to show that trust and performance can be addressed independently!

For trust to be established, the regulator must pick a lowest protected primary usage, P_{txMin} . It must also guarantee that all secondaries have a home band, either licensed or unlicensed. Finally, the regulator must know its catching ability, which it controls through the identity system. These parameters are known before certification time, and therefore the sanction to provide trust can be set before certification time. Moreover, these parameters *do not depend on the technical capability of the secondary*, so certification must check only those parts of the secondary system that affect catching and punishment.

Then, as technology improves, the system will work better. Technology improves with better identity producing lower wrongful conviction rates, and with better methods to find spectrum holes reducing the cost of sensing. In the limit of perfect technology, regardless of sanction, spectrum hole recovery and primary protection will be perfect.

Because this analysis has been quantitative, it has also allowed identification of the hardest-to-protect primaries. Although intuition says that a rarely active primary operates in a band ripe for cognitive colonization, this primary is actually very hard to protect because there is a relatively smaller punishment for interference. So, public safety users will need other means of protection.

The analysis also gave quantitative metrics for performance that can be applied to *any* enforcement system for cognitive radio. For spectrum jails, these metrics can be evaluated to quantitatively understand the tradeoff between the primary protection and the band utilization to really understand the effect of the regulator's choices.

2.4.2 Future work

Chapters 3 and 4 extend the basic model from this chapter to include the database implementation described in Chapter 1, and to include different kinds of secondaries. So, future work along these directions will be deferred.

For the ideas in this chapter, the future work should be focused in two directions: understanding what it means to commit to operating just in the home band, and what an adaptive jail would require.

Dedicated home band

Allowing the choice of just home band operation may be difficult. The secondary would have to be certified that once this decision is made at runtime it will actually only operate in the home band. It may be possible to implement a policy engine that secures this choice if the home band is a separate band. However, this thesis envisions the home band to potentially be a safe slot within a set of frequencies and time slots specified by a database. How could the FCC certify a home band waveform that is defined entirely in software? This requires future work.

This part will show that the situation may not be dire. What would happen if the secondary was subject to wrongful convictions even when operating just in the home band?

As will be discussed in Chapter 4, there is a possibility of harm from even an honest secondary. For example, there may be a mismatch between when the primary turns on and when the secondary notices its presence. This honest harm will be a source of wrongful conviction. Presumably, if the secondary is actually in the home band, it is not even causing any honest harm. So, the wrongful conviction rate when operating just in the home band should be lower than when operating in the primary band. Call this new wrongful conviction rate $P_{wrongHB}$, and let the probability of going to jail while in the home band alone depend only on this and not on the primary transmission probability.

Can cheating still be deterred with wrongful conviction in the home band? Because the cost functions are based on long term behavior, wrongful conviction while operating only in the home band will essentially just reduce the quality (in this case, the size) of the home band.

Lemma 2. *With a wrongful conviction rate of $P_{wrongHB}$ while operating just in the home band, the effective size of the home band is*

$$\beta_{eff} = \frac{1 - P_{pen}}{P_{wrongHB} + 1 - P_{pen}} \beta \quad (2.52)$$

where β is the original size of the home band.

Proof. β represents the number of messages that can be transmitted in the home band per unit time. With wrongful conviction, this number of messages can be sent only when the secondary is not in jail. So,

$$\beta_{eff} = \pi_{NJ} \beta \quad (2.53)$$

The secondary in the home band is sent to jail with probability $P_{wrongHB}$ at every time step, so the probability of not being in jail is $(1 - P_{pen}) / (P_{wrongHB} + 1 - P_{pen})$. The result follows. \square

The necessary sanction will also change because the home band is effectively smaller.

Lemma 3. *When there is wrongful conviction in the home band at a rate of $P_{wrongHB}$, the new required sanction is*

$$P_{pen} > 1 - P_{txMin} P_{catch} \beta + P_{wrongHB} (1 + \beta) \quad (2.54)$$

Proof. The cost of the reduced home band must be compared to operation in both bands, with the actual β because the overall wrongful conviction with operation in both bands already includes the conviction rate in the home band alone:

$$\frac{1}{\beta_{eff}} < \frac{1}{\pi_{NJ}(1 + \beta)} \quad (2.55)$$

$$(1 - P_{pen})(1 - P_{pen} + P_{wrongHB}(1 + \beta) - P_{tx} P_{catch} \beta) < 0 \quad (2.56)$$

$$P_{pen} > 1 - P_{tx} P_{catch} \beta + P_{wrongHB}(1 + \beta) \quad (2.57)$$

where the second equation uses Lemma 2. Maximizing over $P_{tx} > P_{txMin}$ gives the result. \square

This is not always possible to achieve, because $P_{pen} = 1$ implies an infinite jail sentence. Deterrence is only possible with the following condition:




Figure 2.18: The limit on the wrongful conviction rate in the home band or deterrence to still be possible, as a function of the size of the home band and P_{tx} .

Theorem 9. *With a $P_{wrongHB}$ wrongful conviction rate in the home band, deterring cheating is only possible if*

$$P_{wrongHB} < P_{tx}P_{catch}\frac{\beta}{1+\beta} \quad (2.58)$$

Proof. Take the sanction from Lemma 3. The result comes from solving $P_{pen} < 1$. \square

So, this patch will only work if the wrongful conviction rate in the home band is low enough. Luckily, looking at Fig. 2.18, even a surprisingly high wrongful conviction rate will be sufficient to allow trust.

Performance, on the other hand, will suffer. So, solving the problem of trusting the decision to operate only in the home band is an important problem.

Adaptive jail sentences

This chapter, and the rest of this thesis, assume that the sanction must be chosen at certification time and cannot be adapted at runtime. Part of the motivation is that the requirements for trust should be as simple as possible. Adaptation would certainly allow for better performance; it would even allow low P_{tx} primaries to be protected. But it cannot be required for trust to be established, because productive adaptation may be a very hard engineering problem.

If the jail sentence is fixed, but calculated depending on device-specific parameters, those parameters must be trusted to be communicated correctly. This would at least require a trusted certification process for each device.

If the jail sentence changes with feedback on interference, it can be done in two ways depending on the identity system employed. If the identity system cannot pick out individuals, the sentence must rise for all secondaries. So, the sentence would be determined by the highest cost of sensing, effectively reproducing the non-adaptive solution.

Even if the jail sentence is effectively universal over secondary devices, it would still adapt to P_{tx} , right? Not necessarily. If P_{tx} changes faster than the adaptation can track, the sentence will follow the worst P_{tx} . Further, this is not the worst *actual* P_{tx} , it is the worst *observed* P_{tx} . Because the primary and secondary are not synchronized and may not have the same transmit cycle time scale, different secondaries may observe different P_{tx} . The adaptive jail may learn a sanction that is much higher than the fixed approach would advise, because the fixed approach sets a particular P_{txMin} .

To be truly adaptive to individuals instead of the worst case over the group, the identity system must be able to distinguish interference from different individuals. The jail manager, which issues go-to-jail commands, would also have to keep track of the individual sentence for each individual. This would at least require significant coordination.

Adaptive jails are an important problem, and they may be the right solution to minimize the regulatory overhead of spectrum jails. So, they certainly should be studied. But adaptation is not a core problem in developing spectrum jails, and they should not be required to establish trust. They can be developed and implemented later to maximize performance.

Chapter 3

Spectrum jails implemented through databases

For modeling simplicity, it was assumed up to this point that spectrum jails were implemented through a stand-alone monitoring and jail-managing system. This system would watch for interference, would identify culprits, and would send go-to-jail messages. However, the major benefit of using spectrum jails, as argued in the Introduction Chapter, is that they can be implemented through technical solutions that are already required. These include databases, some kind of identity, and kill switches.

The Introduction Chapter also identified a major shortcoming of the database solution as spectrum sharing extends into federal spectrum bands. The databases for the TV whitespaces are designed to receive a GPS coordinate and return a list of available channels [24]. If this coordination service is extended to other bands, as the PCAST report suggests [79], what happens if the databases cannot coordinate access on the same time scale as the spectrum holes the database is trying to recover?

Even if the database could know ahead of time when the primary was going to transmit, it still cannot recover many spectrum holes if there is a large time scale mismatch. Assume that the primary has an *iid* 30% probability of transmitting at any time step, on the primary's time scale. If the database can only send coordination messages every M primary time steps, it can recover a spectrum hole as long as the primary is not active in any of the M steps. Therefore, there is a $(.7)^M$ probability of recovering a spectrum hole. This recovery probability is shown in Fig. 3.1.

But the database does not have to tell the secondaries exactly when to transmit. If the database instead acts as a spectrum manager, and issues operation tokens that will be revoked if the secondary causes interference, this chapter will show that the time scale mismatch is much less of an issue. Any spectrum hole that can be found can be accessed. The secondary can itself sense, a company may implement sensing as a service as a different kind of infrastructure, or the database can step in if it is fast enough.¹ Any spectrum hole that can be found, through any available means, can be used.

The chapter will first extend the model from the previous chapter to capture the database

¹The capability of the database is relative to the primary in the area and is dependent on the local infrastructure. It is possible that the database in a high-use area like New York could be built fast enough to handle secondary requests. But a very rural area may not have the economic incentive to build a highly capable database. Spectrum jails allow the database to step in with full control only when it is capable of doing so appropriately.

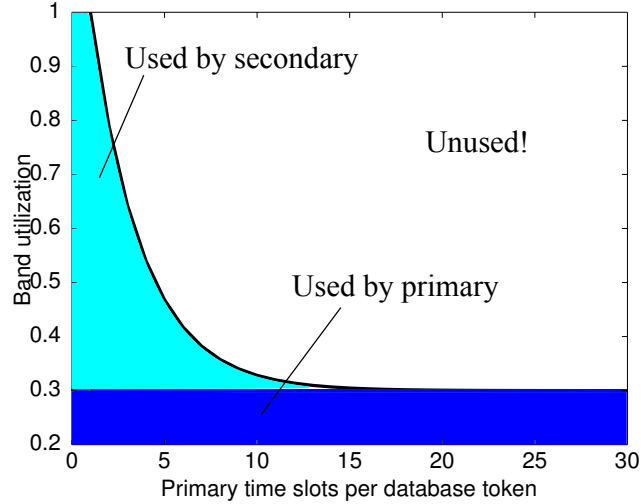


Figure 3.1: Assume a database can only issue one operation token every M time steps, and the primary is modeled as having an *iid* Bernoulli $P_{tx} = 0.3$ probability of transmitting in any time step. This is the band utilization (the primary uses 30% of the slots, the secondary recovers as much as possible). If the mismatch is high, even though spectrum is idle 70% of the time, very few holes can actually be recovered.

implementation. This implementation will be a multiband, multi-time-step version of the model from Chapter 2, so it actually also applies to any situation with wide-band secondaries and multiple primaries. This chapter will then discuss the necessary sanction to achieve trust. Finally, it will explore performance to understand how well spectrum holes can be recovered as technology improves.

3.1 The model and necessary sanction

Taking into account the need for a home band, the database implementation of spectrum jails is shown in flowchart form in Fig. 3.2. The database issues operation tokens every M time steps, and is good for operation in N primary bands. This token includes one “safe” slot, which is one time and frequency slot that is guaranteed to not have primary activity. This slot, which can be left open all the time or which can be a slot designated as empty by the primary, will operate as the home band.²

The secondary must do whatever it needs to in order to not transmit at the same time as the primary. If it causes interference and is caught, the next token it receives from the database will be a “jail” token. For the duration of this token, the secondary must turn off its radio.

To model this, notice that from the secondary’s perspective the only real difference between the token system and the dedicated monitoring/jail system is that the tokens last for more than one time step and frequency band. So, the extended model is the multiple band and multiple time-step version of the Markov jail model, shown in Fig. 3.3.

²Building a safe slot into the database implementation is not necessary. Unlicensed bands can again be used for the home band. This chapter models the home band as a safe slot to demonstrate the idea.

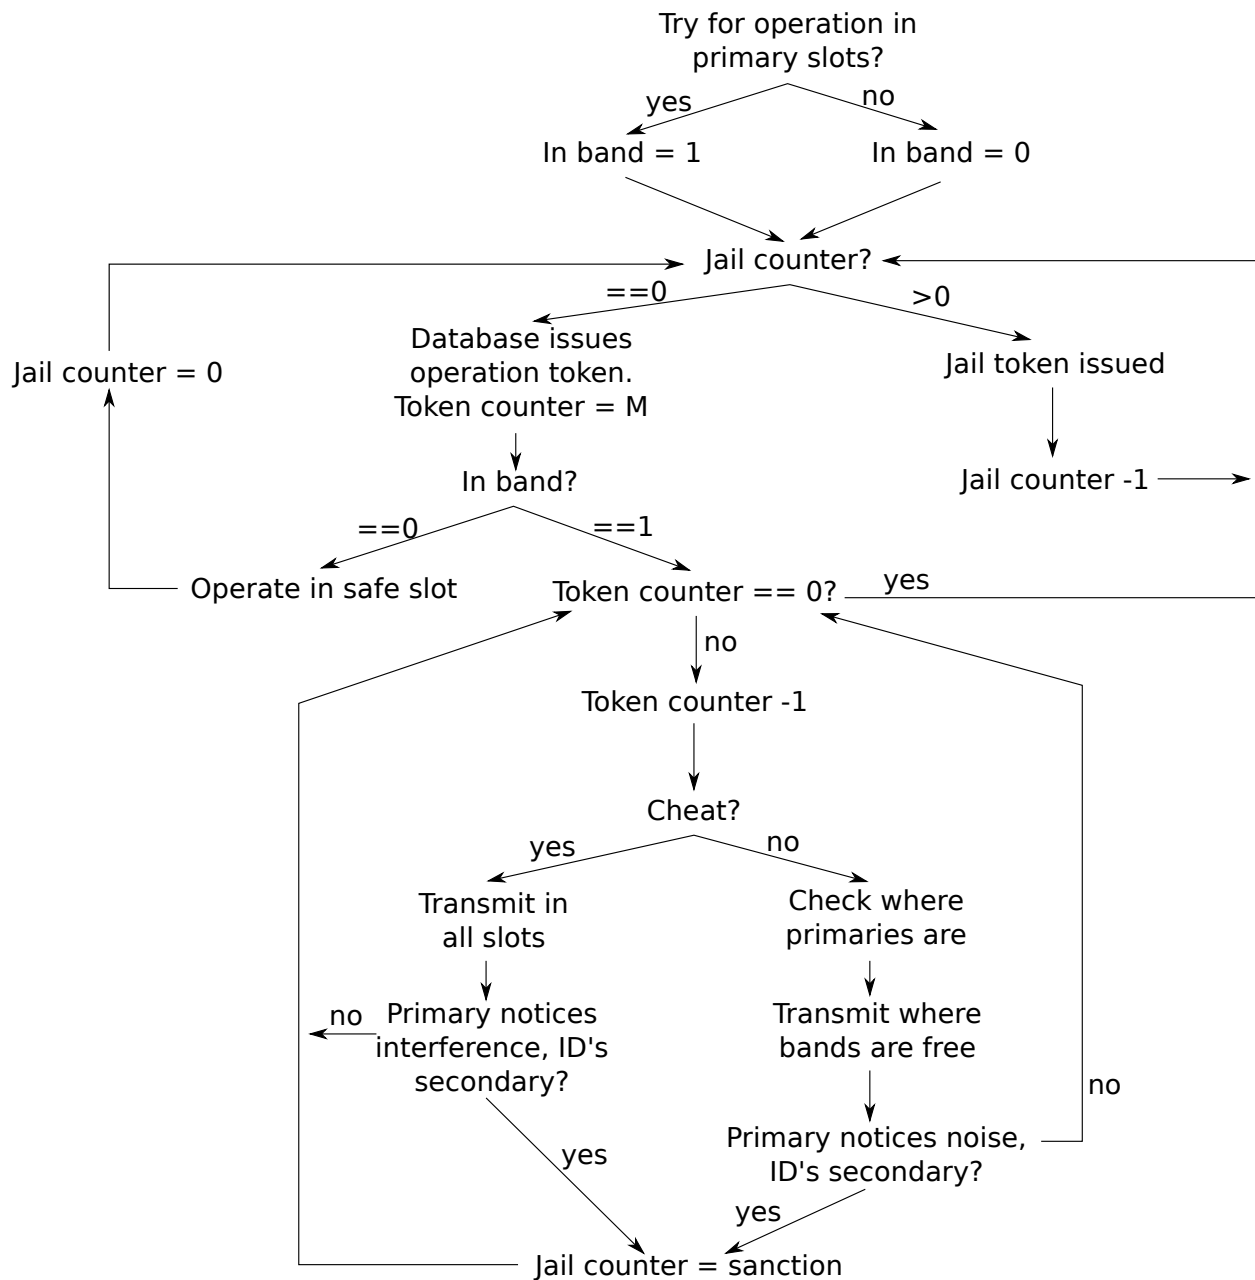


Figure 3.2: Flowchart of events governing the operation of the secondary with a database-token implementation of spectrum jails. The secondary decides whether it will operate in the primary bands; “In band” is the state variable indicating this choice. The database issues operation tokens when the secondary has no jail sentences to serve and jail tokens during jail sentences. The secondary can only transmit when it has an active operation token, which expires after M time steps. Sanctions are issued to secondaries caught causing interference during their operation periods.




Figure 3.3: A jail system run through a database with operation and jail tokens can be modeled as the multiband version of the Markov Chain model, as illustrated here.

The one safe time/frequency slot (which functions as the home band) will still have size β because it is shared between only secondaries and therefore may have different transmission characteristics. All secondaries may use this slot, and so it may be congested. The rules may also allow for higher power in the safe slot. Alternatively, β could account for the existence of another unlicensed band that operates as a second home band.

This model assumes that sensing can only tell the primary's current state, so it needs to be performed at every time slot. Because the secondary cannot do anything else when it is sensing, it loses C_D time out of every available frequency band.

With multiple bands and time steps, the probability of going to jail will also scale. An analog of Thm. 1 holds for the multiband model: the secondary will either be honest in all available slots or cheat in all available slots:

Lemma 4. *With N available slots, the secondary will either cheat in all N , or be honest in all N slots.*

Proof. This is a proof by induction. Assume that for $n = N - 2$ of the slots, with some cheating and some honest, the probability of going to jail at any time step is $P_{tj,n}$. Then, look at the last two bands. In one of the bands, the secondary is honest. In the other, assume that it is in the secondary's best interest to cheat. This means that the cost is lower when cheating than when honest in the last band.

When cheating in the last band, the probability of going to jail is:

$$P_{toJail,1,cheat} = P_{tx}P_{catch} + (1 - P_{tx}P_{catch})(1 - (1 - P_{tj,n})(1 - P_{tx}P_{wrong})) \quad (3.1)$$

Notice that the second term, which is coming from all of the other $N - 1$ bands, does not depend on whether the primary in the last band is active. If the secondary is honest in the last band, the

probability of going to jail is

$$P_{toJail,1,honest} = P_{tx}P_{wrong} + (1 - P_{tx}P_{wrong})(1 - (1 - P_{tj,n})(1 - P_{tx}P_{wrong})) \quad (3.2)$$

If it is in the best interest of the secondary to cheat, this means that (with the cost function modified to include the new probability of going to jail and number of bands):

$$C_{Sec}(P_{cheat} = 1) < C_{Sec}(P_{cheat} = 0) \quad (3.3)$$

$$\frac{1}{\pi_{NJ,cheat}(1 + \beta + n + (1 - P_{tx})(1 - C_{DS}))} < \frac{1}{\pi_{NJ,honest}(\beta + n + 2(1 - P_{tx})(1 - C_{DS}))} \quad (3.4)$$

$$\frac{1 - P_{pen} + P_{toJail,1,cheat}}{(1 - P_{pen})(1 + \beta + n + (1 - P_{tx})(1 - C_{DS}))} < \frac{1 - P_{pen} + P_{toJail,1,honest}}{(1 - P_{pen})(\beta + n + 2(1 - P_{tx})(1 - C_{DS}))} \quad (3.5)$$

$$(1 - P_{pen})(1 - (1 - P_{tx})(1 - C_{DS})) > (P_{toJail,1,cheat} - P_{toJail,1,honest})(\beta + n \quad (3.6)$$

$$+ (1 - P_{tx})(1 - C_{DS})) - P_{toJail,1,honest} \quad (3.7)$$

$$+ P_{toJail,1,cheat}(1 - P_{tx})(1 - C_{DS}) \quad (3.8)$$

This is assumed to be true.

Now, consider the second to last band. In this band, the secondary is choosing whether to cheat or be honest, assuming that in the last band, the secondary is cheating. For this band, the probability of going to jail, while cheating or honest, respectively, is:

$$P_{toJail,2,cheat} = P_{tx}P_{catch} + (1 - P_{tx}P_{catch})(1 - (1 - P_{tj,n})(1 - P_{tx}P_{catch})) \quad (3.9)$$

$$P_{toJail,2,honest} = P_{tx}P_{wrong} + (1 - P_{tx}P_{wrong})(1 - (1 - P_{tj,n})(1 - P_{tx}P_{catch})) \quad (3.10)$$

If it is in the best interest of the secondary to cheat in this band too, the following will be true:

$$C_{Sec}(P_{cheat} = 1) < C_{Sec}(P_{cheat} = 0) \quad (3.11)$$

$$\frac{1}{\pi_{NJ,cheat}(2 + \beta + n)} < \frac{1}{\pi_{NJ,honest}(1 + \beta + n + (1 - P_{tx})(1 - C_{DS}))} \quad (3.12)$$

$$\frac{1 - P_{pen} + P_{toJail,2,cheat}}{(1 - P_{pen})(2 + \beta + n)} < \frac{1 - P_{pen} + P_{toJail,2,honest}}{(1 - P_{pen})(1 + \beta + n + (1 - P_{tx})(1 - C_{DS}))} \quad (3.13)$$

$$(1 - P_{pen})(1 - (1 - P_{tx})(1 - C_{DS})) > (P_{toJail,2,cheat} - P_{toJail,2,honest})(\beta + n + 1) \quad (3.14)$$

$$+ P_{toJail,2,cheat}(1 - P_{tx})(1 - C_{DS}) - P_{toJail,2,honest} \quad (3.15)$$

Notice that the left hand side is the same as what was assumed to be true above. So, it is sufficient to show that the right hand side of Eq. (3.8) is bigger than Eq. (3.15). To do that, just show that

the following is true:

$$0 < (P_{toJail,1,cheat} - P_{toJail,1,honest})(\beta + n + (1 - P_{tx})(1 - C_{DS})) \quad (3.16)$$

$$+ P_{toJail,1,cheat}(1 - P_{tx})(1 - C_{DS}) - P_{toJail,1,honest} \quad (3.17)$$

$$- (P_{toJail,2,cheat} - P_{toJail,2,honest})(\beta + n + 1) \quad (3.18)$$

$$- P_{toJail,2,cheat}(1 - P_{tx})(1 - C_{DS}) + P_{toJail,2,honest} \quad (3.19)$$

$$0 < (P_{toJail,1,cheat} - P_{toJail,2,cheat})(\beta + n + (1 - P_{tx})(1 - C_{DS})) \quad (3.20)$$

$$+ (P_{toJail,2,honest} - P_{toJail,1,honest})(\beta + n + 1) \quad (3.21)$$

$$- (P_{toJail,2,cheat} - P_{toJail,2,honest}) \quad (3.22)$$

$$+ (P_{toJail,1,cheat} - P_{toJail,1,honest})(1 - P_{tx})(1 - C_{DS}) \quad (3.23)$$

$$0 < (*) (\beta + n + (1 - P_{tx})(1 - C_{DS})) + (**) (\beta + n + 1) - (***) + (****) (1 - tx)(1 - C_{DS}) \quad (3.24)$$

where the substitutions in the last line are made to split the problem up for simpler calculation. Plugging in the terms from Eqs. (3.1), (3.2), (3.9) gives:

$$(*) = -P_{tx}(1 - P_{tx}P_{catch})(1 - P_{tj,n})(P_{catch} - P_{wrong}) \quad (3.25)$$

$$(**) = P_{tx}(1 - P_{tx}P_{wrong})(1 - P_{tj,n})(P_{catch} - P_{wrong}) \quad (3.26)$$

$$(***) = P_{tx}(1 - P_{tx}P_{catch})(1 - P_{tj,n})(P_{catch} - P_{wrong}) \quad (3.27)$$

$$(****) = P_{tx}(1 - P_{tx}P_{wrong})(1 - P_{tj,n})(P_{catch} - P_{wrong}) \quad (3.28)$$

Continuing the above, with these substitutions gives:

$$0 < (*) (\beta + n + (1 - P_{tx})(1 - C_{DS})) + (**) (\beta + n + 1) - (***) + (****) (1 - P_{tx})(1 - C_{DS}) \quad (3.29)$$

$$0 < P_{tx}(P_{catch} - P_{wrong})(1 - P_{tj,n})(-(1 - P_{tx}P_{catch})(\beta + n + (1 - P_{tx})(1 - C_{DS}))) \quad (3.30)$$

$$+ (1 - P_{tx}P_{wrong})(\beta + n + 1) - (1 - P_{tx}P_{catch}) \quad (3.31)$$

$$+ (1 - P_{tx}P_{wrong})(1 - P_{tx})(1 - C_{DS}) \quad (3.32)$$

$$0 < P_{tx}^2(P_{catch} - P_{wrong})^2(1 - P_{tj,n})(\beta + n + 1 + (1 - P_{tx})(1 - C_{DS})) \quad (3.33)$$

which is always true, as long as $P_{catch} > P_{wrong}$. \square

With this lemma, the probability of going to jail scales depending on whether the secondary is cheating:

$$P_{tojail,cheat} = 1 - (1 - P_{tx}P_{catch})^{NM-1} \quad (3.34)$$

$$P_{tojail,honest} = 1 - (1 - P_{tx}P_{wrong})^{NM-1} \quad (3.35)$$

This is the real difference between this multiband model and the basic model of the last chapter. Each primary, each slot, and the probability of going to jail in each slot is independent of the others. The total probability of going to jail changes accordingly.

Then, using the same format of cost function from the last chapter (the numerator is the number of time steps per token; the denominator is the average number of messages sent per token), the cost function for the secondary is

$$C_D = \min_i C_{D,i} \quad (3.36)$$

given the following states:

1. Use just the primary slots:

$$C_{D,1} = \frac{M}{\pi_{NJ}(P_{cheat} + (1 - P_{cheat})(1 - P_{tx})(1 - C_{DS}))(NM - 1)} \quad (3.37)$$

2. Use both safe slot and primary slots:

$$C_{D,2} = \frac{M}{\pi_{NJ}((P_{cheat} + (1 - P_{cheat})(1 - P_{tx})(1 - C_{DS}))(NM - 1) + \beta)} \quad (3.38)$$

3. Use just the safe slot:

$$C_{D,3} = \frac{M}{\beta} \quad (3.39)$$

Theorem 10. *With the multiband model above, the necessary sanction that guarantees that no secondary will cheat when the primary is operational more than P_{txMin} of the time is:*

$$P_{pen} > 1 - \frac{(1 - (1 - P_{txMin}P_{catch})^{NM-1})\beta}{NM - 1} \quad (3.40)$$

$$(3.41)$$

Proof. The sanction is found by making cheating less preferable than operation only in the safe slot for a given P_{tx} :

$$C_{D,3} < C_{D,2}(P_{cheat} = 1) \quad (3.42)$$

$$\frac{M}{\beta} < \frac{(P_{to jail, cheat} + 1 - P_{pen})M}{(1 - P_{pen})(\beta + NM - 1)} \quad (3.43)$$

$$P_{pen} > 1 - \frac{P_{to jail, cheat}\beta}{NM - 1} \quad (3.44)$$

$$P_{pen} > 1 - \frac{(1 - (1 - P_{tx}P_{catch})^{NM-1})\beta}{NM - 1} \quad (3.45)$$

Notice that as P_{tx} decreases, the necessary P_{pen} goes to 1. Therefore, the regulator must again define a minimum protected P_{tx} , P_{txMin} . Plugging in the minimum protected P_{tx} gives the result. \square

Notice that as in the simpler jail model, the sanction depends only on parameters that can be known at certification time. This includes P_{catch} , P_{txMin} , β as before, as well as N , M which are the expansion parameters. The secondary has one home time/frequency slot of size β . N , M are extra bands that the secondary may expand into, and they must be known at certification time.

What does certification time mean in the context of a database? The jail sentence is going to the secondary through the database, so presumably it could be changed to suit a more local context. This allows the possibility of some adaptation, if only to the current size (congestion level) of the home band or safe slot. But this adaptation may be on a slower than runtime time scale. The parameters that make up the sanction are certainly known at the slower time scale. For the rest of this thesis, any parameters that are changed on this “database regulatory” time scale will be referred to as certification time.

To connect this sanction to the last chapter, interpret it in terms of the expansion. $1 - (1 - P_{txMin}P_{catch})^{NM-1}$ is the probability of going to jail with the minimum protected P_{tx} . This value increases with N, M , so it can be bounded as the minimum probability of being caught and sent to jail. Call this $P_{toJail,min}$. Then, the sanction becomes:

$$P_{pen} > 1 - \frac{\beta}{NM - 1} P_{toJail,min} \quad (3.46)$$

where $\beta/(NM - 1)$ is the *expansion factor*. In general, then, the sanction is based on how much the secondary is expanding beyond its home band. The more it expands its band use, the more tempting it is to cheat and therefore the higher the sanction must be.

This form of sanction is in fact exactly the same as in the last chapter: before, there was only one primary band, so the expansion was just β . Then, the probability of going to jail in the last chapter was $P_{toJail,min} = P_{txMin}P_{catch}$. So, the sanction in this chapter is just a generalized version of the sanction from the last chapter.

3.2 Performance with improving technology

How are N, M chosen? M , the number of primary time slots per database token, is necessarily constrained by the capability of the database. But the regulator may choose to operate the database slower than its capability. In any case, M can be known at certification time. N , the number of frequency bands the secondary uses, can be chosen in several ways: the regulator can choose a maximum N_{max} at certification time, certify the secondary can only expand this much, and set the sanction accordingly. Alternatively, the secondary could choose N at certification time and have a personalized jail time that it follows when it receives a jail token. If the regulatory body allows the sanction to change at runtime (there is no technical reason why this is not possible if jails are operated through the databases), the sanction could even be adapted to the amount of expansion the secondary chooses at runtime.

This section explores performance given these different choices. The first part assumes that N, M are fixed at certification time, and explores how the performance improves as technology gets better. The result is that spectrum jails run through databases can recover significantly more spectrum holes than even a database with primary look-ahead could alone.

The second part of this section explores expansion. It shows two effects: given a fixed N_{max} set by the regulators, secondaries will choose to fill all available bands as technology improves. Second, if N can be chosen at runtime, how far will secondaries choose to expand as technology gets better and what is the resulting performance?

3.2.1 Recovering spectrum holes with fixed N, M

This part focuses on performance given a database with a fixed capability of sending tokens every M primary time steps. As an example for comparison, remember the back-of-the-envelope computation that resulted in Fig. 3.1. The database knows ahead of time when the primary will transmit, but can coordinate spectrum access only every M time steps. The primary uses each of its time steps with probability 0.3. Only $(0.7)^M$ spectrum holes are recoverable at any time.

How well can spectrum jails recover spectrum holes? This is shown in Fig. 3.4 for different values of C_{DS}, P_{wrong} , assuming $N = 1$. As these two performance parameters go to zero, spectrum

jails can recover all of the spectrum holes, regardless of M . So, as technology improves, the speed of the database matters less.

Quantitatively, how does performance improve with better technology given a fixed N, M ? First, define the appropriate versions of the performance metrics in Chapter 2.

The primary's interests are represented by the probability of collisions:

$$\text{collision} = \pi_{NJ} P_{\text{cheat}} \mathbb{1}_{\text{in primary band}} \quad (3.47)$$

The secondary's interests are captured by the secondary's overhead relative to a secondary user that will never cheat and is not subject to jail. This user has cost function:

$$C_{D,\text{noJail}} = \min_i C_{D,\text{noJail},i} \quad (3.48)$$

given the following states:

1. Use just the primary slots:

$$C_{D,\text{noJail},1} = \frac{M}{(1 - P_{tx})(1 - C_{DS})(NM - 1)} \quad (3.49)$$

2. Use both safe slot and primary slots:

$$C_{D,\text{noJail},2} = \frac{M}{((1 - P_{tx})(1 - C_{DS})(NM - 1) + \beta)} \quad (3.50)$$

3. Use just the safe slot:

$$C_{D,\text{noJail},3} = \frac{M}{\beta} \quad (3.51)$$

The overhead for the secondary is then defined as

$$O_S = \frac{U_{D,\text{noJail}} - U_D}{U_{D,\text{noJail}}} \quad (3.52)$$

$$= 1 - \frac{C_{D,\text{noJail},1}}{C_{D,1}} \quad (3.53)$$

where $U_D = 1/C_{D,1}$ is the utility of a secondary with the jail system and $U_{D,\text{noJail}} = 1/C_{D,\text{noJail},1}$ is the utility of the secondary without jail. As in Chapter 2, the overhead is measured as lost opportunity in the primary bands, not the home slot.

Lemma 5. *The overhead of a secondary subject to jail who does not cheat relative to a user not subject to jail is*

$$O_S = 1 - \pi_{NJ} \quad (3.54)$$

or the time the secondary spends wrongfully in jail.

Proof. Plug in definitions of $C_{D,1}(P_{\text{cheat}} = 0)$ and $C_{D,\text{noJail},1}$ into the definition of O_S . \square

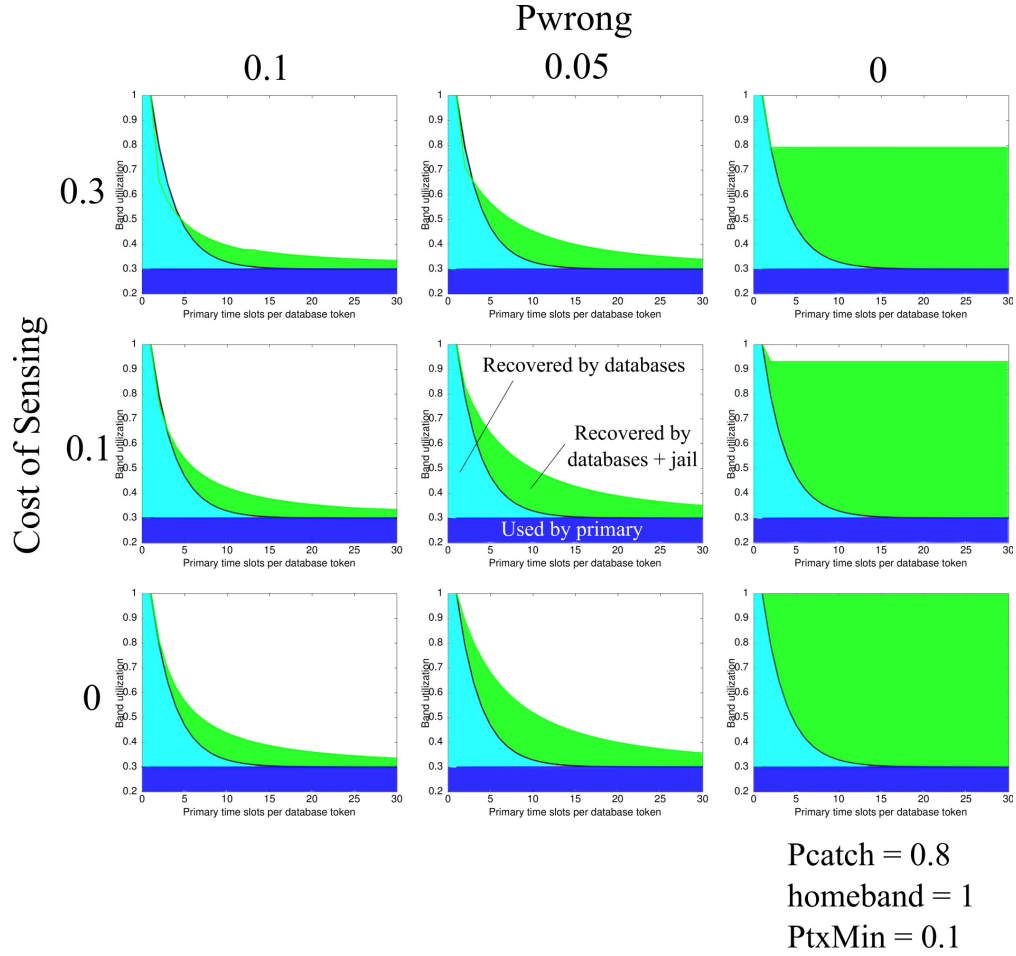


Figure 3.4: Spectrum holes recovered by spectrum jails vs a database alone. If both sensing cost and wrongful convictions rates are high, spectrum jails will recover about the same amount of spectrum holes as a centralized database solution that knows ahead of time when the primary will transmit. However, as technology improves, spectrum jails can recover all spectrum holes, regardless of the speed of the database.

Because the secondary overhead is simply the percentage of time it spends in jail, this overhead can again be captured in an overall metric of spectrum holes due to time in jail:

$$\text{holes}_{\text{jail}} = (1 - \pi_{NJ})(1 - P_{tx})(NM - 1)\mathbb{1}_{\text{in primary band}} \quad (3.55)$$

Finally, spectrum holes due to operation only in the safe slot are captured by:

$$\text{holes}_{\text{homeband}} = (1 - P_{tx})(NM - 1)\mathbb{1}_{\text{in home band}} \quad (3.56)$$

The next theorem shows how these metrics improve as technology improves.

Theorem 11. *Performance improves with technology in the following ways:*

1. $\lim_{P_{\text{wrong}}, C_{DS} \rightarrow 0} \text{collision} = 0 \forall P_{tx}$ where collisions are defined in Eq. (3.47).
2. $\lim_{P_{\text{wrong}}, C_{DS} \rightarrow 0} \text{holes}_{\text{jail}} = 0 \forall P_{tx}$ where the metric is defined in Eq. (3.55)
3. $\lim_{P_{\text{wrong}}, C_{DS} \rightarrow 0} \text{holes}_{\text{homeband}} = 0 \forall P_{tx}$ where the metric is defined in Eq. (3.56)

Proof. This proof treats each metric separately:

1. The collisions metric is defined as $\text{collision} = \pi_{NJ}P_{\text{cheat}}\mathbb{1}_{\text{in primary band}}$, and the sanction is defined to guarantee no cheating for $P_{tx} > P_{tx\text{Min}}$. So it is sufficient to show that the secondary will choose $P_{\text{cheat}} = 1$ for smaller values of P_{tx} as P_{wrong}, C_{DS} improve. The secondary will choose to cheat if the cost is lower than being honest, while operating in both bands:

$$C_{D,2}(P_{\text{cheat}} = 1) < C_{D,2}(P_{\text{cheat}} = 0) \quad (3.57)$$

$$\frac{M(1 - P_{\text{pen}} + P_{\text{toJail,cheat}})}{(1 - P_{\text{pen}})(NM - 1 + \beta)} < \frac{M(1 - P_{\text{pen}} + P_{\text{toJail,honest}})}{(1 - p_{\text{en}})((1 - P_{tx})(1 - C_{DS})(NM - 1) + \beta)} \quad (3.58)$$

$$0 > -C_{DS}(1 - P_{tx})(NM - 1)(1 - P_{\text{pen}} + P_{\text{toJail,cheat}}) \quad (3.59)$$

$$- P_{\text{toJail,honest}}(NM - 1 + \beta) \quad (3.60)$$

$$+ P_{\text{toJail,cheat}}((1 - P_{tx})(NM - 1) + \beta) \quad (3.61)$$

$$- P_{tx}(NM - 1)(1 - P_{\text{pen}}) \quad (3.62)$$

Plugging in $P_{\text{pen}} = 1 - \frac{\beta P_{\text{toJail,cheat}}}{NM - 1}$ gives the condition

$$0 > -C_{DS}(1 - P_{tx})(NM - 1)(1 - P_{\text{pen}} + P_{\text{toJail,cheat}}) - P_{\text{toJail,honest}}(NM - 1 + \beta) \quad (3.63)$$

$$+ P_{\text{toJail,cheat}}(1 - P_{tx})(\beta + NM - 1) \quad (3.64)$$

Given that $P_{\text{toJail,honest}} = 1 - (1 - P_{tx}P_{\text{wrong}})^{NM-1}$, as P_{wrong} and C_{DS} get smaller, the right side negative terms get smaller, so the condition will not be satisfied even with smaller values of P_{tx} . This means that the secondary will choose not to cheat even for smaller values of P_{tx} . As $P_{\text{wrong}}, C_{DS} \rightarrow 0$, the right side becomes positive unless $P_{tx} = 0$, at which point there cannot be collisions. So, $\lim_{P_{\text{wrong}}, C_{DS} \rightarrow 0} \text{collision} = 0 \forall P_{tx}$.

2. The holes due to jail metric is defined as

$\text{holes}_{\text{jail}} = (1 - \pi_{NJ})(1 - P_{tx})(NM - 1)\mathbb{1}_{\text{in primary band}}$, so it is sufficient to show that $\pi_{NJ} \rightarrow 1$ as $P_{\text{wrong}}, C_{DS} \rightarrow 0$ when the secondary is being honest. Plugging in the definition

of π_{NJ} from Eq. (2.2), and substituting the new honest probability of going to jail:

$$\pi_{NJ} = \frac{1 - P_{pen}}{P_{toJail,honest} + 1 - P_{pen}} \quad (3.65)$$

$$= \frac{1 - P_{pen}}{1 - (1 - P_{tx}P_{wrong})^{NM-1} + 1 - P_{pen}} \quad (3.66)$$

Clearly, π_{NJ} gets bigger as P_{wrong} gets smaller, and $\lim_{P_{wrong} \rightarrow 0} \pi_{NJ} = 1$. There is no dependence on C_{DS} .

3. The holes due to retreat to the home band metric is defined as $\text{holes}_{homeband} = (1 - P_{tx})(NM - 1)\mathbb{1}_{\text{in home band}}$. It is sufficient to show that, for any P_{tx} as technology improves, the secondary will eventually not retreat to the home band. The secondary will operate only in the home band if the cost of the home band is less than the cost of honest use in both bands:

$$C_{D,3} < C_{D,2}(P_{cheat} = 0) \quad (3.67)$$

$$\frac{M}{\beta} < \frac{M(1 - P_{pen} + P_{toJail,honest})}{(1 - P_{pen})((1 - P_{tx})(1 - C_{DS})(NM - 1) + \beta)} \quad (3.68)$$

$$0 > -C_{DS}(1 - P_{pen})(1 - P_{tx})(NM - 1) - P_{toJail,honest}\beta + (1 - P_{pen})(1 - P_{tx})(NM - 1) \quad (3.69)$$

As P_{wrong} and C_{DS} get smaller, this condition will be violated only by larger values of P_{tx} . And in the limit as $P_{wrong}, C_{DS} \rightarrow 0$, this condition becomes

$$(1 - P_{pen})(1 - P_{tx})(NM - 1) < 0 \quad (3.70)$$

which is violated only when $P_{tx} = 1$ at which point there are no holes left in the primary band to fill. □

3.2.2 Optimal expansion

Regardless of the sanction chosen at certification time, the secondary chooses its own expansion at run time. This part explores how many bands the secondary will expand into. This is done in two cases: if the regulator chooses N_{max} at certification time to set the sanction, and the secondary chooses $N < N_{max}$ at runtime. The second case assumes that both N and the sanction can be set at runtime and explores how far the secondary will choose to expand.

If N_{max} is chosen at certification time, this sets a fixed sanction P_{pen} for run time. As technology improves, the secondary will expand to fill all the available bands.

Theorem 12. *With a fixed sanction P_{pen} , a fixed M , and a maximum number of bands N_{max} ,*

$$\underset{N \leq N_{max}}{\operatorname{argmin}} \lim_{P_{wrong} \rightarrow 0} C_{D,2}(P_{cheat} = 0) = N_{max} \quad (3.71)$$




Figure 3.5: The necessary sanction goes up as the expansion rises because there is more opportunity and therefore more temptation to cheat as the secondary operates in more primary bands.

Proof.

$$\operatorname{argmin}_{N \leq N_{max}} \lim_{P_{wrong} \rightarrow 0} C_{D,2}(P_{cheat} = 0) = \operatorname{argmin}_{N \leq N_{max}} \lim_{P_{wrong}} \frac{M(1 - P_{pen} + P_{toJail,honest})}{(1 - P_{pen})((1 - P_{tx})(1 - C_{DS})(NM - 1) + \beta)} \quad (3.72)$$

$$= \operatorname{argmin}_{N \leq N_{max}} \frac{M(1 - P_{pen})}{(1 - P_{pen})((1 - P_{tx})(1 - C_{DS})(NM - 1) + \beta)} \quad (3.73)$$

$$= N_{max} \quad (3.74)$$

□

If, however, the sanction can be set based on the N chosen at runtime, the secondary has full control over how many bands it can expand into.

As the secondary expands into more primary bands, the sanction rises. This is shown in Fig. 3.5. At the same time, the probability of getting wrongfully convicted also rises. Intuitively, then, as the expansion rises, so does the secondary overhead and total unrecovered spectrum holes. This is shown in Fig. 3.6 for different P_{wrong} . The metric values are averaged over C_{DS} and the probability of primary transmission P_{tx} , as in the last chapter, Fig. 2.14.

But the secondary makes the expansion decision on the basis of cost, not overhead. Fig. 3.7 gives an example of a particular secondary. M is fixed at 3, and N is varied. The secondary will choose to expand to different numbers of bands depending on P_{wrong} . Notice that if P_{wrong} is high enough, the cost is always 3. This indicates that, regardless of N , this secondary will stay in its home band all the time.

Fig. 3.8 shows the optimal value of N , given $M = 3$ over different values of C_{DS} and P_{tx} . From the regulator's perspective, the sanctions should be designed for the worst case, but performance should be measured for the targeted cases. Sharing is really only desirable for secondaries with reasonable costs of sensing and primaries that have available spectrum. So, this plot only shows a limited section of the total range of C_{DS} and P_{tx} .




Figure 3.6: Overhead metrics as a function of different levels of expansion. As the expansion grows, so do the total spectrum holes because the wrongful conviction rate and the length of sanctions are scaling with the expansion. The collisions are decreasing because the secondary is spending more time in jail.




Figure 3.7: Although the *overhead* rises with expansion, the secondary cost tells a different story. The secondary will expand depending on the wrongful conviction rate, expanding more when the wrongful conviction is lowered.

Notice that if there is plenty of opportunity, the secondary will choose to expand more. This is because when the primary is rarely around, there are few wrongful convictions, and the secondary has a smaller threat of jail.

Averaging over the secondary sensing cost and primary transmission probability from Fig. 3.8 produces summary overhead and expansion statistics. These are shown in Fig. 3.9. Different points along the line have different levels of P_{wrong} , some of which are indicated. The average behavior of the secondary is to expand based on the value of P_{wrong} such that the average percentage of wasted spectrum is around 25%. This seems high, but is a function of allowing the secondary to expand very far — with a wrongful conviction rate of 10^{-4} , the secondary would choose to operate in about 130 primary bands. This suggests that if the regulator wants high utilization, it should set a fixed N_{max} even if the sanction can be changed at runtime.

3.3 Concluding remarks

3.3.1 What we have learned so far

This chapter showed that the database implementation of spectrum jails, advocated in Chapter 1, can be modeled as a multi-band, multi-time version of the basic spectrum jails investigated in Chapter 2. Further, the main ideas from the last chapter hold with this model as well: trust can be guaranteed at certification time, and performance will improve as technology gets better.

The difference in this model is being able to choose the expansion. The sanction is effectively determined by the expansion factor — how much the secondary is expanding into primary bands, relative to the size of its home band. As the sanction grows, for fixed technology, the overhead metrics will get worse as well.

How should the expansion factor be chosen? The capability of the databases themselves




Figure 3.8: This is the optimal expansion for a delay-constrained user over different values of the cost of sensing and the probability of primary transmission. If there is a lot of opportunity (P_{tx} small), the realized number of wrongful convictions will be small, so the desired expansion is very high.




Figure 3.9: Averaging over the range of sensing cost and primary transmission probabilities, this is the average expansion vs average spectrum holes per slot as a function of the P_{wrong} . Notice that the optimal average amount of spectrum holes, from the secondary perspective, is around 25%. The regulator can decrease this by restricting the potential secondary expansion.

determine the length of time a token must last. But the number of primary frequency bands the secondary can access can be chosen more freely. If the regulator chooses the maximum number of bands at certification time, the secondary will expand to fill them all if the wrongful conviction rate is low enough. Likewise, if the secondary is able to choose its own expansion, it will continue to use more bands as the wrongful conviction rate is lowered.

The key observation of this chapter is the comparison between the performance of spectrum jails vs the database alone. If the database is responsible for coordinating all spectrum access, it is limited in its recovery capability by the speed at which it can coordinate actions. If, however, the database is only coordinating spectrum jail tokens, the performance depends on how well secondaries can find spectrum holes and how small the regulator can make the probability of wrongful conviction.

Spectrum jail performance is limited by the best spectrum hole recovery technology, not the speed of the current generation of databases.

3.3.2 Future work – adapting jails to expansion

In Chapter 2, the idea of adaptive jails was discussed in the Future Work. The difficulty encountered there is inherently different than the difficulty in making sanctions adaptive to expansion.

In the last chapter, adaptive jails meant being adaptive to the cost of sensing or primary transmissions. Neither of these can be easily measured or proven at certification time. Adapting at runtime would require either detailed tracking of which sanction applies to which secondary, or it would likely adapt to the worst present primary or secondary characteristic.

Adapting to a different expansion factor is likely an easier problem. While it may be hard to measure C_{DS} or P_{tx} , the secondary will have to choose the number of primary bands it is operating in. The length of the sanction could be programmed into the secondary as a function of the number of bands used. The difficulty in certifying lies not in measuring, but guaranteeing that the secondary cannot fake its band usage. Work should be done to figure out how difficult this is, and whether this kind of adaptation is easily implementable.

Chapter 4

Universal jails

The last two chapters assumed a very simple secondary and harm model: the secondary's goal was to maximize average throughput, and whenever it cheated, it caused the primary to drop a packet. This chapter questions both of these assumptions.

The wireless ecosystem does not include only devices that try to maximize throughput. There are streaming services, file sharing, voice, and even sensor networks that have wildly different service requirements. Will a jail that delays messages deter cheating for all of these disparate services? The first section of this chapter addresses the different kinds of secondaries, and uses energy-sensitive users as a case study for how to extend spectrum jails. The goal will be the same as in the previous chapters: guarantee trust at certification time and then understand how performance improves as technology improves.

The effect of secondary interference was also simplified in the last two chapters so that every instance of interference resulted in a dropped packet. This is not necessarily the case. If the secondary is very far away from the primary, it may cause a small amount of interference that only causes a dropped packet when the fading is particularly bad. The previous chapters also assumed that an honest secondary would cause no interference, but this is not necessarily true either. If the secondary's sensor does not detect the primary, it will cause harm even though the secondary thinks it is being honest. The second part of this chapter explores the implications of variable harm on trust and performance.

Throughout, this chapter will show that only minor modifications are required to extend the “trust first and performance later” property from the last two chapters to the more realistic situations explored here. It will also indicate any new difficult to enforce cases and what new choices need to be made by regulators at certification time.

4.1 The need for different sanctions

In a real wireless ecosystem, there are many different kinds of devices and traffic that have different QOS requirements. For example, there are laptops downloading very large files that care about the overall length of time it takes to download the file. There are also sensor networks, perhaps measuring the stress of joints in bridges [128] or measuring environmental data in inhospitable environments [129]. The nodes of these networks may have very little data to send. If the network is monitoring, say, an active volcano [129], it may not care whether its message is

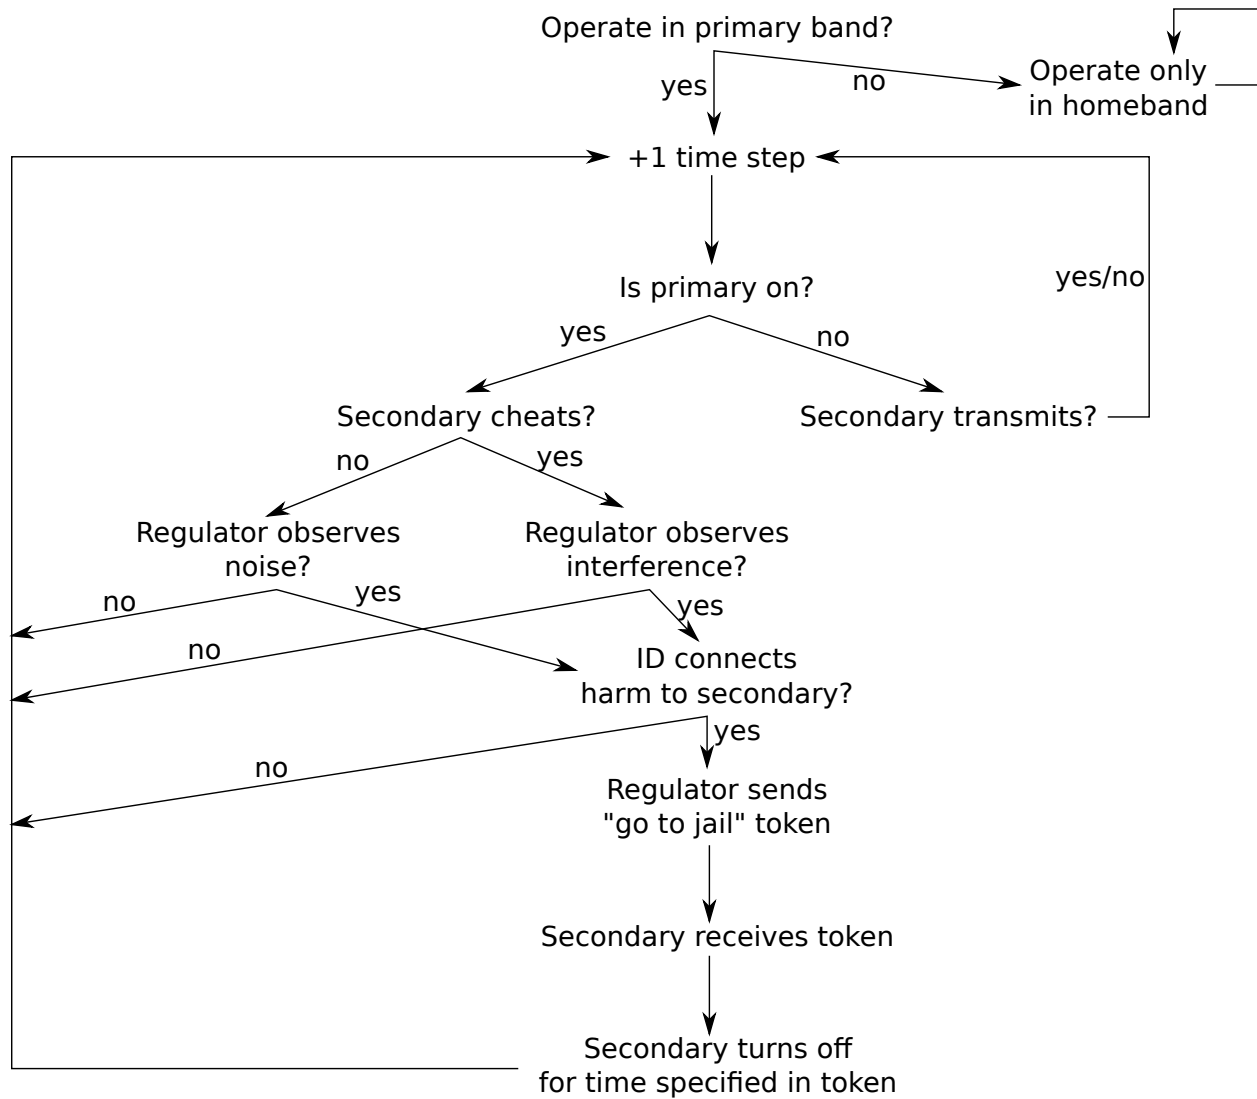


Figure 4.1: This is the flowchart of spectrum jails without the infrastructure of the databases. This is identical to Fig. 2.1 in Chapter 2 in all but the inclusion of the home band — the secondary decides at the beginning whether it will operate only in its home band or in the primary bands as well.




Figure 4.2: The sanction and resulting percentage of time spent in jail for honest users when only delay-constrained secondaries are present.

delayed by a few minutes, or even a day. But, the batteries would be almost impossible to change. If the sensor nodes were smart dust [130], they would have a limit on the speed of recouping energy losses.

What happens if the jail system from Chapter 2 (reviewed in the flowchart in Fig. 4.1) is used to try to deter cheating with a purely energy-sensitive user?¹

At first glance, the jail system should not work at all. No amount of waiting will cause any pain whatsoever to the purely energy-constrained secondary. More precisely, remember the problem of a rarely active primary from Chapter 2 – the energy-constrained user spends energy looking for the primary, but spends no energy sitting in jail. It will never choose to sense when jail cannot harm its operation.

But is jail actually free from an energy perspective? Receiving and decoding the “go-to-jail” command requires some energy, as does waking up at the end of the jail sentence. Assume that this natural energy cost of jail can be measured relative to the energy cost of sensing.

Fig. 4.2 shows the P_{pen} required to keep just the delay user honest by P_{tx} , the amount of time the primary is transmitting. Compare this to Fig. 4.3, which shows the jail sentence needed for different energy costs of jail relative to the energy cost of sensing. If jail is more expensive, the required sanction to keep the energy-constrained user honest is similar to that required by the delay-sensitive user.

Fig. 4.4 shows the same required sanction as a function of the energy cost of jail relative to sensing. Notice that for any P_{tx} , there is a corresponding energy cost of jail and length of sentence that will deter cheating. Naturally, this implies that to deter energy-constrained users from cheating, the regulator must be able to control the energy cost of jail.

There are many ways to do this. The “go-to-jail” command could be designed to be difficult to decode, therefore requiring high processing energy. If the device is a smart phone,

¹For most of this chapter, the model from Chapter 2 will be used because it is simpler. Section 4.2.4 will show how the sanction changes with the database implementation.




Figure 4.3: If energy-constrained users are operational as well, using just a delay-oriented jail to deter cheating requires a potentially much higher sanction depending on the natural energy cost of sitting in jail. This causes much more wrongful jail time for honest users.




Figure 4.4: The sanction required to deter cheating in energy-constrained users as a function of the natural energy cost of jail relative to the energy cost of sensing.

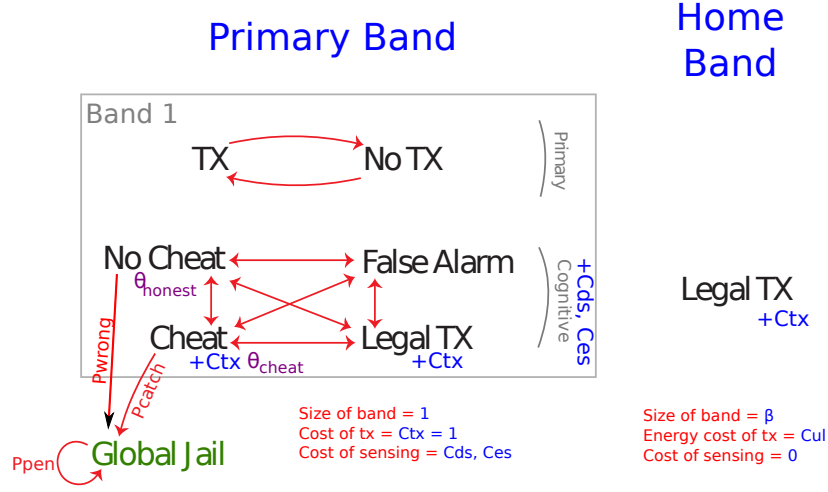


Figure 4.5: Illustration of the Markov model used in this section. The secondary has the option to transmit in a home band along with or instead of the primary band. The transition probabilities and costs for different states are shown.

jail could include a command to turn on the screen and keep it on for the duration. Finally, the device could be required to transmit gibberish on a dedicated gibberish band. Although impractical because of the wasted resources, this last example inspires the terminology for the rest of the chapter – while devices are in jail, they must be forced to “sing” to burn energy, and the energy cost of jail will be called C_{sing} , measured relative to the cost of one unit of transmission.²

Purely delay-constrained, and purely energy-constrained utility functions do not cover all possible service desires for secondary devices. The next parts will cover sanctions for energy-constrained and mixed delay and energy users. Section 4.3 will cover secondary devices with utility functions that are only a small perturbation away from those already covered. Future work, in Section 4.5, will discuss what must be done to extend jails to other kinds of secondary devices.

4.2 Energy and mixed users

4.2.1 The Model

The model for this part is illustrated in Fig. 4.5. The delay-constrained user is the same as in Chapter 2, defined in Eq. (2.19). The energy-constrained user has the same choice to operate in just the primary band, just the home band, or both. The cost function is structured in the same way as the delay user. The denominator is always the average number of messages sent per time step (or database token). The numerator is the average energy spent per time step. This includes the energy sensing cost, and any transmit costs when the secondary is not in jail. It also includes the energy sanction when the secondary is in jail.

²In the future, it may also be possible to use a community service sanction that will require a secondary to sense and report primary activity when it is in jail. For the discussion here, it is assumed that the sanction must be set at certification time and it must work for all realizations of secondary technology. Because it may be difficult to certify the energy cost of community service, it may also be difficult to certify that such a sanction will be effective.

1. Use just the primary band:

$$C_{E,1} = \frac{\pi_{NJ}(P_{cheat} + (1 - P_{cheat})((1 - P_{tx})(1 - C_{DS}) + C_{ES})) + (1 - \pi_{NJ})C_{sing}}{\pi_{NJ}(P_{cheat} + (1 - P_{cheat})(1 - P_{tx})(1 - C_{DS}))} \quad (4.1)$$

2. Use both home band and primary band:

$$C_{E,2} = \frac{\pi_{NJ}(C_{UL}\beta + P_{cheat} + (1 - P_{cheat})((1 - P_{tx})(1 - C_{DS}) + C_{ES})) + (1 - \pi_{NJ})C_{sing}}{\pi_{NJ}(P_{cheat} + (1 - P_{cheat})(1 - P_{tx})(1 - C_{DS}) + \beta)} \quad (4.2)$$

3. Use just the home band:

$$C_{E,3} = \frac{C_{UL}\beta}{\beta} \quad (4.3)$$

Then, the cost for an energy user is:

$$C_E = \min_i C_{E,i} \quad (4.4)$$

Where C_{ES} is the energy cost of sensing, C_{sing} is the energy cost of each unit of jail, and C_{UL} is the energy cost of transmitting in the home band, all measured relative to the energy cost of transmission in the primary band. C_{UL} is allowed to be different than the energy cost of transmitting in the primary band because the home band may be an unlicensed band with congestion. Higher transmit power, better codes, etc, may be required, resulting in a different energy cost of operation. Note that if $C_{UL} < 1$, the energy user would always transmit in the home band. So $C_{UL} > 1$ is the only case that needs to be considered.

Mixed users that care about both energy and delay constraints will have the following cost function:

$$C_{mixed} = \min_i kC_{D,i} + (1 - k)C_{E,i} \quad (4.5)$$

Parameters are listed and defined in Table 4.2.1.

4.2.2 Extending trust to energy and mixed users

This part extends the trust results from Chapter 2 to the energy and mixed user cases. Just as in the last chapter, the secondary will choose to always cheat or never cheat:

Lemma 6. *For the energy user with cost function in Eq. (4.4) and the mixed user with cost function in Eq. (4.5), the optimal $P_{cheat} \in \{0, 1\}$*

Proof. All of the operating modes for both the energy and mixed users have cost functions that satisfy the conditions of Lemma 1. An argument like that in Theorem 3 extends the result to allow the choice between operating modes. \square

Using this lemma, the required sanction to guarantee no secondary will cheat given $P_{tx} > P_{txMin}$ is given in the following theorem.

Theorem 13. *The required sanction to guarantee no secondaries cheat when $P_{tx} > P_{txMin}$ is:*

$$P_{pen} > 1 - P_{txMin}P_{catch}\beta \quad (4.6)$$

$$C_{sing} > \frac{(1 - P_{pen})(C_{UL} - 1)}{P_{txMin}P_{catch}} \quad (4.7)$$

Table 4.1: Parameters definitions for the extended spectrum jails problem. The jail sentence is defined based on those parameters used at certification time, and the performance is determined by those parameters used at runtime.

Parameter	Definition	Timescale when used
P_{tx}	Probability the primary is transmitting	Runtime
P_{pen}	Probability of staying in jail	Certification
C_{sing}	Energy cost of jail per time step	Certification
P_{catch}	Probability cheating secondary is caught	Certification
β	Size of the home band	Certification
C_{UL}	Energy cost home band	Certification
N	Number of primary frequency bands	Certification
M	Primary time steps per database token	Certification
P_{txMin}	Minimum protected probability of primary transmission	Certification
$\theta_{cheatMax}$	Maximum allowed amount of cheating harm	Certification
$\theta_{honestMax}$	Maximum allowed amount of honest harm	Certification
$-\Delta$	Unmodeled secondary benefit for cheating	Certification
θ_{cheat}	Probability of harm caused by cheating secondary	Runtime
θ_{honest}	Probability of harm caused by honest secondary	Runtime
C_{DS}	Delay cost of checking for primary	Runtime
C_{ES}	Energy cost of checking for primary	Runtime
P_{wrong}	Probability honest secondary is sent to jail	Runtime

Proof. First, take the energy-constrained user. Assume there is some set $P_{pen} > 0$. The C_{sing} constraint is found by the same argument as in Thm. 5: first make sure that operating in the home band alone is more attractive than cheating for the energy-only user. Note that the energy-only user will use either the home band or the primary band, but never both. It chooses the one with the better energy use profile.

$$C_{E,1} > C_{E,3} \quad (4.8)$$

$$\Rightarrow C_{sing} > \frac{(1 - P_{pen})(C_{UL} - 1)}{P_{tx}P_{catch}} \quad (4.9)$$

For completeness, if the same calculation is done with $C_{E,2}$ instead of $C_{E,1}$, the same result is achieved. For a fixed P_{pen} , $C_{sing} \rightarrow \infty$ as $P_{tx} \rightarrow 0$, so for any fixed C_{sing} , there exists some P_{tx} at which the sanction will not be sufficient. Therefore, define a P_{txMin} below which protection for the primary will not be guaranteed. This fixes C_{sing} for any defined value of P_{pen} .

To cover the mixed user, P_{pen} must be set to a specific value. The mixed cost function relies on P_{cheat} only through a non-decreasing function of C_D and C_E . So, set P_{pen} to deter pure delay users, and set C_{sing} to deter energy users. Operating in the home band alone is better than cheating for both energy and delay components, and therefore operating in the home band alone is better than cheating for the mixed user. \square

As in Chapter 2, notice the expression for the required sanction. It depends on the characteristics of the home band, β , C_{UL} , the catching ability P_{catch} , and the minimum protected P_{tx} . It does *not* depend on the secondary technology or the operation of the primary. So, again,

trust can be guaranteed for *any* secondary, certifying only its compliance with the jail system and not its technical capability.

4.2.3 Better performance with better technology

As in Chapter 2, the performance of energy and mixed users will improve as technology improves. To investigate this claim, the overhead definitions are generalized to include the energy and mixed users.

First, define an energy-constrained and mixed user that will always act honestly, but that is not subject to wrongful convictions:

$$C_{E,noJail} = \min_i C_{E,noJail,i} \quad (4.10)$$

$$C_{mixed,noJail} = \min_i kC_{D,noJail,i} + (1-k)C_{E,noJail,i} \quad (4.11)$$

where

$$C_{E,noJail,1} = \frac{(1-P_{tx})(1-C_{DS}) + C_{ES}}{(1-P_{tx})(1-C_{DS})} \quad (4.12)$$

$$C_{E,noJail,2} = \frac{(1-P_{tx})(1-C_{DS}) + C_{ES} + C_{UL}\beta}{(1-P_{tx})(1-C_{DS}) + \beta} \quad (4.13)$$

$$C_{E,noJail,3} = \frac{C_{UL}\beta}{\beta} \quad (4.14)$$

for operation in the primary band alone, both bands, and the home band alone, respectively. $C_{D,noJail,i}$ are defined in Eq. (2.27).

The secondary overhead is generalized to reflect these new cost functions. It will still be measured as the extra cost in just the primary band, relative to the honest user with no jail.

Theorem 14. *The generalized overhead for the mixed secondary user is:*

$$O_S = \frac{(1-\pi_{NJ})(k + (1-k)C_{sing})}{(1-\pi_{NJ})(k + (1-k)C_{sing}) + \pi_{NJ}(k + (1-k)(C_{ES} + (1-P_{tx})(1-C_{DS})))} \quad (4.15)$$

Proof. Plug in equations into the definition:

$$O_S = 1 - \frac{C_{mixed,honest}}{C_{mixed}} \quad (4.16)$$

□

This is nicely interpreted as the cost of jail over the total cost when in jail and when not in jail.

The primary collision metric, and the holes due to jail metric will remain the same as in Chapter 2:

$$\text{collision} = \pi_{NJ}P_{cheat}\mathbb{1}_{\text{in primary band}} \quad (4.17)$$

$$\text{holes}_{jail} = (1-\pi_{NJ})(1-P_{tx})\mathbb{1}_{\text{in primary band}} \quad (4.18)$$

The holes due to operation just in the home band must be changed slightly with the new cost functions. The delay-only user was choosing between operating in both bands or just the home band because those were the lowest cost options. The energy user is choosing between operating in either the primary or the home band. The mixed user will use all three options.

To capture the effect of jail, the holes due to retreat to the home band metric will include only those cases when the user with jail is in the home band alone and the without jail user chooses to operate in the primary band (either alone or along with the home band).

$$\text{holes}_{\text{homeband}} = (1 - P_{tx}) \mathbb{1}_{\text{with-jail user in home band}} \mathbb{1}_{\text{without-jail user in primary band}} \quad (4.19)$$

Because the overhead is qualitatively so similar to the last chapter, only the summary plots are included here in Fig. 4.6. These average the metrics over a range of values for P_{tx} , and either C_{DS} or C_{ES} . The secondary overhead into energy and delay cases. Any user with a mixed cost function will have the appropriate mix of these two lines. The other lines average the energy and delay users.

Just like in the last two chapters, once a protection level is chosen, improving technology will improve performance with all metrics. Two theorems will be presented – first performance improvement with just energy-constrained users, then performance improvement for mixed users.

Theorem 15. *For the energy-constrained users, performance improves with technology in the following ways:*

1. $\lim_{P_{\text{wrong}}, C_{ES}, C_{DS} \rightarrow 0} \text{collision} = 0 \forall P_{tx}$ where collisions are defined in Eq. (4.17).
2. $\lim_{P_{\text{wrong}}, C_{ES}, C_{DS} \rightarrow 0} OS = 0 \forall P_{tx}$ where the secondary overhead is defined in Thm. 14 with $k = 0$.
3. $\lim_{P_{\text{wrong}}, C_{ES}, C_{DS} \rightarrow 0} \text{holes}_{\text{jail}} = 0 \forall P_{tx}$ where the metric is defined in Eq. (4.18)
4. $\lim_{P_{\text{wrong}}, C_{ES}, C_{DS} \rightarrow 0} \text{holes}_{\text{homeband}} = 0 \forall P_{tx}$ where the metric is defined in Eq. (4.19)

Proof. This proof treats each metric separately:

1. The collisions metric is defined as $\text{collision} = \pi_{NJ} P_{\text{cheat}} \mathbb{1}_{\text{in primary band}}$, and the sanction is defined to guarantee no cheating for $P_{tx} > P_{tx\text{Min}}$. So it is sufficient to show that the secondary will choose $P_{\text{cheat}} = 1$ for smaller values of P_{tx} as $P_{\text{wrong}}, C_{ES}, C_{DS}$ improve. The energy-constrained secondary's cost is minimized by either operating just in the primary band or just in the home band. Focusing on operation in the primary band alone, the secondary will choose to cheat for any $P_{tx} < P_{tx\text{Min}}$ such that

$$C_{E,1}(P_{\text{cheat}} = 1) < C_{E,1}(P_{\text{cheat}} = 0) \quad (4.20)$$

$$\frac{1 - P_{\text{pen}} + P_{tx} P_{\text{catch}} C_{\text{sing}}}{1 - P_{\text{pen}}} < \frac{(1 - P_{\text{pen}})((1 - P_{tx})(1 - C_{DS}) + C_{ES}) + P_{tx} P_{\text{wrong}} C_{\text{sing}}}{(1 - P_{\text{pen}})(1 - P_{tx})(1 - C_{SD})} \quad (4.21)$$

$$0 > P_{tx} P_{\text{catch}} C_{\text{sing}}(1 - P_{tx}) - C_{DS} P_{tx} P_{\text{catch}} C_{\text{sing}} - C_{ES}(1 - P_{\text{pen}}) \quad (4.22)$$

$$- P_{\text{wrong}} P_{tx} C_{\text{sing}} \quad (4.23)$$

As P_{wrong}, C_{ES} , and C_{DS} get smaller, the right side negative terms get smaller, so the condition will not be satisfied even for smaller values of P_{tx} . This means that the secondary




Figure 4.6: Summary metrics for different levels of wrongful conviction. Energy and delay-constrained secondaries have different lines for their overheads; mixed users will have the appropriate mix of the two overheads. The primary and spectrum hole lines are all averaged between energy and delay-constrained users. All metrics get better as wrongful conviction rates go down.

will choose not to cheat even for smaller values of P_{tx} . As $P_{wrong}, C_{ES}, C_{DS} \rightarrow 0$, the right side becomes positive unless $P_{tx} = 0$, at which point there cannot be collisions. So, $\lim_{P_{wrong}, C_{ES}, C_{DS} \rightarrow 0} \text{collision} = 0 \forall P_{tx}$.

2. The secondary overhead metric with $k = 0$, for energy-constrained users, is:

$$O_S = \frac{(1 - \pi_{NJ})C_{sing}}{(1 - \pi_{NJ})C_{sing} + \pi_{NJ}(C_{ES} + (1 - P_{tx})(1 - C_{DS}))} \quad (4.24)$$

$$= \frac{P_{tx}P_{wrong}C_{sing}}{P_{tx}P_{wrong}C_{sing} + (1 - P_{pen})(C_{ES} + (1 - P_{tx})(1 - C_{DS}))} \quad (4.25)$$

Clearly, this goes to zero as $P_{wrong} \rightarrow 0$.

3. The holes due to jail metric is defined as $\text{holes}_{jail} = (1 - \pi_{NJ})(1 - P_{tx})\mathbb{1}_{\text{in primary band}}$, so it is sufficient to show that $\pi_{NJ} \rightarrow 1$ as $P_{wrong}, C_{DS} \rightarrow 0$ when the secondary is being honest. Plugging in the definition of π_{NJ} from Eq. (2.2), and setting the sanction as $P_{pen} = 1 - P_{tx} \text{Min} P_{catch} \beta$:

$$\pi_{NJ} = \frac{1 - P_{pen}}{P_{toJail} + 1 - P_{pen}} \quad (4.26)$$

$$= \frac{P_{tx} \text{Min} P_{catch} \beta}{P_{tx}P_{wrong} + P_{tx} \text{Min} P_{catch} \beta} \quad (4.27)$$

Clearly, π_{NJ} gets bigger as P_{wrong} gets smaller, and $\lim_{P_{wrong} \rightarrow 0} \pi_{NJ} = 1$. There is no dependence on C_{DS} .

4. The holes due to retreat to the home band metric is defined as

$\text{holes}_{homeband} = (1 - P_{tx})\mathbb{1}_{\text{with-jail user in home band}} \mathbb{1}_{\text{without-jail user in primary band}}$. It is sufficient to show that the secondary will not retreat to the home band for any P_{tx} . The secondary will operate only in the home band if the cost of the home band is less than the cost of honest use the the primary band:

$$C_{E,3} < C_{E,1}(P_{cheat} = 0) \quad (4.28)$$

$$C_{UL} < \frac{\pi_{NJ}((1 - P_{tx})(1 - C_{DS}) + C_{ES}) + \pi_J C_{sing}}{\pi_{NJ}(1 - P_{tx})(1 - C_{DS})} \quad (4.29)$$

$$C_{UL} < \frac{(1 - P_{pen})((1 - P_{tx})(1 - C_{DS}) + C_{ES}) + P_{tx}P_{wrong}C_{sing}}{(1 - P_{pen})(1 - P_{tx})(1 - C_{DS})} \quad (4.30)$$

$$0 > (1 - P_{pen})(1 - P_{tx})(C_{UL} - 1) - C_{DS}(1 - P_{pen})(1 - P_{tx})(C_{UL} - 1) \quad (4.31)$$

$$- C_{ES}(1 - P_{pen}) - P_{wrong}P_{tx}C_{sing} \quad (4.32)$$

As P_{wrong} , C_{ES} , and C_{DS} get smaller, this condition will be violated only by larger values of P_{tx} . In the limit as $P_{wrong}, C_{ES}, C_{DS} \rightarrow 0$, this condition becomes

$$(1 - P_{pen})(1 - P_{tx})(C_{UL} - 1) < 0 \quad (4.33)$$

which is violated only when $P_{tx} = 1$ at which point there are no holes left in the primary band to fill.

□

With mixed users, the same limiting performance can be proved. It is not sufficient to just use the energy and delay user theorems because the energy and delay parts of the mixed cost function may want to operate in different ways: the energy part may want to use just the primary band, while the delay part wants to use both bands. This affects the proof of the primary protection becoming perfect and the secondary choosing not to retreat to the home band.

Theorem 16. *For the mixed users:*

1. $\lim_{P_{wrong}, C_{ES}, C_{DS} \rightarrow 0} \text{collision} = 0 \forall P_{tx}$ where collisions are defined in Eq. (4.17).
2. $\lim_{P_{wrong}, C_{ES}, C_{DS} \rightarrow 0} O_S = 0 \forall P_{tx}$ where the secondary overhead is defined in Thm. 14.
3. $\lim_{P_{wrong}, C_{ES}, C_{DS} \rightarrow 0} \text{holes}_{jail} = 0 \forall P_{tx}$ where the metric is defined in Eq. (4.18)
4. $\lim_{P_{wrong}, C_{ES}, C_{DS} \rightarrow 0} \text{holes}_{homeband} = 0 \forall P_{tx}$ where the metric is defined in Eq. (4.19)

Proof. This proof treats each metric separately:

1. The collisions metric is defined as $\text{collision} = \pi_{NJ} P_{cheat} \mathbb{1}_{\text{in primary band}}$, and the sanction is defined to guarantee no cheating for $P_{tx} > P_{txMin}$. The mixed user case is different from the individual cases because the energy user will only ever use one band, the delay user will always use both bands, but the mixed user may end up in any combination. So, it must stay honest even if the delay component pulls it to use both bands or even if the energy component pulls the user to use only one band. This proof will first show that if the mixed user uses both bands, it will never cheat in the limit of good technology. The second part will address the single band case by showing that if it is worthwhile to use only one band, then it is worthwhile to be honest in that one band.

It is sufficient to consider the limiting case when $P_{wrong} = 0, C_{DS} = 0, C_{ES} = 0$.

When using both bands, the mixed user will not cheat if (with the substitutions $P_{wrong} = 0, C_{ES} = 0, C_{DS} = 0$):

$$C_{mixed,2}(P_{cheat} = 0) < C_{mixed,2}(P_{cheat} = 1) \quad (4.34)$$

$$\frac{k + (1 - k)(C_{UL}\beta + 1 - P_{tx})}{\beta + 1 - P_{tx}} < \frac{k(1 - P_{pen} + P_{tx}P_{catch}) + (1 - k)((C_{UL}\beta + 1)(1 - P_{pen}))}{(1 - P_{pen})(\beta + 1)} \quad (4.35)$$

$$+ \frac{(1 - k)P_{tx}P_{catch}C_{sing}}{(1 - P_{pen})(\beta + 1)} \quad (4.36)$$

$$0 < kP_{tx}(P_{tx}P_{catch}(\beta + 1 - P_{tx}) - P_{tx}(1 - P_{pen})) \quad (4.37)$$

$$+ (1 - k)(P_{tx}(1 - P_{pen})\beta(1 - C_{UL})) \quad (4.38)$$

$$+ (1 - k)P_{tx}P_{catch}C_{sing}(\beta + 1 - P_{tx}) \quad (4.39)$$

Plugging in $P_{pen} = 1 - P_{txMin}P_{catch}\beta$ and $C_{sing} = (C_{UL} - 1)\beta$ gives:

$$0 < kP_{tx}((\beta + 1 - P_{tx})P_{catch} - P_{catch}P_{txMin}\beta) \quad (4.40)$$

$$+ (1 - k)\beta P_{tx}P_{catch}(\beta(C_{UL} - 1)(1 - P_{txMin}) + P_{catch}(C_{UL} - 1)(1 - P_{tx})) \quad (4.41)$$

which is always true for all $P_{tx} > 0$. At this point, there is no primary to interfere with, so the collisions metric does indeed go to zero.

When using just the primary band, the condition for not cheating is:

$$C_{mixed,1}(P_{cheat} = 0) < C_{mixed,1}(P_{cheat} = 1) \quad (4.42)$$

$$\frac{k + (1 - k)(1 - P_{tx})}{(1 - P_{tx})} < \frac{k(1 - P_{pen} + P_{tx}P_{catch}) + (1 - k)((1 - P_{pen}) + P_{tx}P_{catch}C_{sing})}{1 - P_{pen}} \quad (4.43)$$

$$0 < kP_{tx}(P_{catch}(1 - P_{tx}) - (1 - P_{pen})) + (1 - k)P_{tx}(P_{catch}C_{sing}(1 - P_{tx})) \quad (4.44)$$

$$0 < kP_{tx}(P_{catch}(1 - P_{tx}) - P_{txMin}P_{catch}\beta) \quad (4.45)$$

$$+ (1 - k)P_{tx}P_{catch}(C_{UL} - 1)\beta(1 - P_{tx}) \quad (4.46)$$

where the last line comes from plugging in $P_{pen} = 1 - P_{txMin}P_{catch}\beta$ and $C_{sing} = (C_{UL} - 1)\beta$. This is the condition for not cheating. Now, it must be shown to be true when it is in the secondary's best interest to use only one band. The secondary will use just the primary band if the energy cost of the home band is prohibitively high. So, the next part will find the necessary C_{UL} to make using just one band the right choice. This value will then be plugged into the condition here to show the secondary will never cheat.

We already know that the secondary will be honest if using both bands. If the secondary would rather be honest in just the primary band than be honest in both bands, this means:

$$C_{mixed,1}(P_{cheat} = 0) < C_{mixed,2}(P_{cheat} = 0) \quad (4.47)$$

$$\frac{k + (1 - k)(1 - P_{tx})}{1 - P_{tx}} < \frac{k + (1 - k)(C_{UL}\beta + 1 - P_{tx})}{\beta + 1 - P_{tx}} \quad (4.48)$$

$$C_{UL} > 1 + \frac{k}{(1 - P_{tx})(1 - k)} \quad (4.49)$$

Now, plug this into Eq. (4.46):

$$0 < kP_{tx}(P_{catch}(1 - P_{tx}) - P_{txMin}P_{catch}\beta) + (1 - k)P_{tx}P_{catch}(C_{UL} - 1)\beta(1 - P_{tx}) \quad (4.50)$$

$$0 < kP_{catch}(1 - P_{tx} - P_{txMin}\beta) + (1 - k)P_{catch}P_{tx}\beta(1 - P_{tx}) \left(\frac{k}{(1 - P_{tx})(1 - k)} \right) \quad (4.51)$$

$$0 < kP_{catch}(1 - P_{tx} - P_{txMin}\beta - P_{tx}\beta) \quad (4.52)$$

$$0 < kP_{catch}(1 - P_{tx} + \beta(1 - P_{txMin})) \quad (4.53)$$

which is always true.

For completeness, check that even if it is better to cheat in the primary band alone than operate honestly in both bands, the secondary will still not cheat in the single band:

$$C_{mixed,1}(P_{cheat} = 1) < C_{mixed,2}(P_{cheat} = 0) \quad (4.54)$$

$$\frac{k + (1 - k)(\pi_{NJ} + \pi_J C_{sing})}{\pi_{NJ}} < \frac{k + (1 - k)(C_{UL}\beta + 1 - P_{tx})}{\beta + 1 - P_{tx}} \quad (4.55)$$

Plugging in $P_{pen} = P_{txMin}P_{catch}\beta$, $C_{sing} = (C_{UL} - 1)\beta$ and solving for $C_{UL} - 1$ gives:

$$C_{UL} - 1 > \frac{k(P_{txMin}\beta(\beta - P_{tx}) + P_{tx}(\beta + 1 - P_{tx}))}{(1 - k)(P_{txMin} - P_{tx}(\beta + 1 - P_{tx}))\beta} \quad (4.56)$$

Plugging this into Eq. (4.46):

$$0 < kP_{tx}(P_{catch}(1 - P_{tx}) - P_{txMin}P_{catch}\beta) + (1 - k)P_{tx}P_{catch}(C_{UL} - 1)\beta(1 - P_{tx}) \quad (4.57)$$

$$0 < kP_{tx}(P_{catch}(1 - P_{tx}) - P_{txMin}P_{catch}\beta) \quad (4.58)$$

$$+ (1 - k)P_{tx}P_{catch}\beta(1 - P_{tx}) \left(\frac{k(P_{txMin}\beta(\beta - P_{tx}) + P_{tx}(\beta + 1 - P_{tx}))}{(1 - k)(P_{txMin} - P_{tx}(\beta + 1 - P_{tx}))\beta} \right) \quad (4.59)$$

$$0 < kP_{txMin}\beta(\beta + 1/\beta - P_{txMin}) \quad (4.60)$$

This is always true, so the secondary will never cheat in this case.

2. From Thms. 8 and 15, it can be inferred that the secondary overhead goes to zero in the limit of good technology for both energy and delay constrained users. As the mixed user overhead is the appropriate mix of the two, the mixed secondary overhead will also go to zero in the limit of perfect technology.
3. From Thms. 8 and 15, the holes due to jail metric goes to zero in the limit of improved technology. As this metric is no different with the mixed users, the same holds true for this case.
4. The holes due to retreat to the home band metric is defined as $\text{holes}_{homeband} = (1 - P_{tx})\mathbb{1}_{\text{with-jail user in home band}} - \mathbb{1}_{\text{without-jail user in primary band}}$. In the limit of good technology, the with-jail user will never cheat (based on part 1 of this theorem), and so will never be sent to jail. The resulting cost function is the same as the without-jail user. Therefore, they will operate the same way, and the holes due to retreat to the home band metric goes to zero.

□

4.2.4 Database implementation for energy and mixed users

Like in Chapter 3, the database implementation can be captured by a multi-band, multi-time version of the spectrum jail model.

As with the time cost of sensing, the energy cost of sensing will scale with the number of slots used. Sensing must be done at every time step. [131] shows that multiband sensing is more processing-intensive than narrowband sensing, so the energy spent sensing should increase with the number of frequency bands. For simplicity, this model assumes that the energy cost of sensing will scale linearly with the number of bands used. So, if the secondary chooses to sense, it will spend $C_{ES}(NM - 1)$ units of energy for every operation token.

The mixed and energy users have the following cost function:

$$C_{mixed} = kC_D + (1 - k)C_E \quad (4.61)$$

where

$$C_D = \min_i C_{D,i} \quad (4.62)$$

$$C_E = \min_i C_{E,i} \quad (4.63)$$

given the following states:

1. Use just the primary slots:

$$C_{D,1} = \frac{M}{\pi_{NJ}(P_{cheat} + (1 - P_{cheat})(1 - P_{tx})(1 - C_{DS}))(NM - 1)} \quad (4.64)$$

$$C_{E,1} = \frac{\pi_{NJ}((P_{cheat} + (1 - P_{cheat})(1 - P_{tx})(1 - C_{DS}))(NM - 1) + C_{ES}(NM - 1))}{\pi_{NJ}(P_{cheat} + (1 - P_{cheat})(1 - P_{tx})(1 - C_{DS}))(NM - 1)} \quad (4.65)$$

$$+ \frac{(1 - \pi_{NJ})C_{sing}}{\pi_{NJ}(P_{cheat} + (1 - P_{cheat})(1 - P_{tx})(1 - C_{DS}))(NM - 1)} \quad (4.66)$$

2. Use both safe slot and primary slots:

$$C_{D,2} = \frac{M}{\pi_{NJ}((P_{cheat} + (1 - P_{cheat})(1 - P_{tx})(1 - C_{DS}))(NM - 1) + \beta)} \quad (4.67)$$

$$C_{E,2} = \frac{\pi_{NJ}(C_{UL}\beta + (P_{cheat} + (1 - P_{cheat})(1 - P_{tx})(1 - C_{DS}))(NM - 1) + C_{ES}(NM - 1))}{\pi_{NJ}((P_{cheat} + (1 - P_{cheat})(1 - P_{tx})(1 - C_{DS}))(NM - 1) + \beta)} \quad (4.68)$$

$$+ \frac{(1 - \pi_{NJ})C_{sing}}{\pi_{NJ}((P_{cheat} + (1 - P_{cheat})(1 - P_{tx})(1 - C_{DS}))(NM - 1) + \beta)} \quad (4.69)$$

3. Use just the safe slot:

$$C_{D,3} = \frac{M}{\beta} \quad (4.70)$$

$$C_{E,3} = \frac{C_{UL}\beta}{\beta} \quad (4.71)$$

where the delay user cost functions are included for reference. Also for reference, remember that:

$$\pi_{NJ} = \frac{1 - P_{pen}}{1 - P_{pen} + P_{toJail}} \quad (4.72)$$

where

$$P_{toJail} = P_{cheat}P_{toJail,cheat} + (1 - P_{cheat})P_{toJail,honest} \quad (4.73)$$

where

$$P_{toJail,cheat} = 1 - (1 - P_{tx}P_{catch})^{NM-1} \quad (4.74)$$

$$P_{toJail,honest} = 1 - (1 - P_{tx}P_{wrong})^{NM-1} \quad (4.75)$$

It can be inferred from the last section and Chapter 3 that with the appropriate sanction, trust and performance with the database implementation will be like the single band model. So, this section only finds the necessary sanction for trust with energy and mixed users.

Theorem 17. *With the multiband model above, the necessary sanction that guarantees that no secondary will cheat when the primary is operational more than P_{txMin} of the time is:*

$$P_{pen} > 1 - \frac{(1 - (1 - P_{txMin}P_{catch})^{NM-1})\beta}{NM - 1} \quad (4.76)$$

$$C_{sing} > \frac{(1 - P_{pen})(NM - 1)(C_{UL} - 1)}{1 - (1 - P_{txMin}P_{catch})^{NM-1}} = (C_{UL} - 1)\beta \quad (4.77)$$

Proof. The sanctions are found by making cheating less preferable than operation only in the safe slot for a given P_{tx} . The sanction for delay-only users is found in Thm. 10. For energy-only users:

$$C_{E,3} < C_{E,2}(P_{cheat} = 1) \quad (4.78)$$

$$C_{UL} < \frac{(1 - P_{pen})(C_{UL}\beta + (NM - 1)) + P_{to\text{jail},cheat}C_{sing}}{(1 - P_{pen})(\beta + NM - 1)} \quad (4.79)$$

$$C_{sing} > \frac{(1 - P_{pen})(NM - 1)(C_{UL} - 1)}{P_{to\text{jail},cheat}} \quad (4.80)$$

Because mixed users are defined by a linear combination of delay and energy costs, if the sanctions guarantee that the delay and energy users would both rather operate in the home band than cheat, the mixed user will also choose operation in the home band over cheating. So, plugging in for $P_{to\text{jail},cheat}$ and again defining a minimum P_{tx} gives the result. \square

4.3 General delay, energy, and mixed cost functions

The last section showed that the energy-constrained user, with the appropriate sanction, will have the same qualitative properties as the delay-constrained user. The sanction can be set to guarantee trust at certification time, and as technology improves, performance will also improve.

The mixed user considered in the last part was the simplest possible cost function combining the desires of the energy and delay-constrained users. This section considers more general mixed users. It also shows how a margin on the sanctions may be used to guarantee trust even for cost functions that are small perturbations away from the cost functions considered here. Any cost function that includes elements not well covered by delay and energy costs (such as QOS constraints based on packet timing) will be left to future work.

The focus of this part is extending trust, so it will show only the required changes to the sanctions.

4.3.1 Different kinds of mixed users

Mixed users do not necessarily have a cost function limited to a linear combination of the delay and energy costs. If a new mixed cost function is any nondecreasing function of the energy and delay cost functions, then the sanctions already found are sufficient to deter cheating.

Theorem 18. *For cost functions of the form*

$$C = f(C_D, C_E) \quad (4.81)$$

where $f(\cdot)$ is a continuous, non-decreasing function of C_D, C_E , with C_D, C_E as defined in Eqs. (2.19) and (4.4) with no other dependence on P_{cheat} . Then, the sanctions in Thm. 13 are sufficient to guarantee no cheating when $P_{tx} > P_{txMin}$

Proof. If the user with this cost function wants to cheat, based on Thm. 13, both its energy and delay components would be smaller by moving to the home band. Because of the properties of $f(\cdot)$, this user too would have a smaller cost in the home band. \square

4.3.2 Margins for other unmodeled phenomena

The cost function may depend mostly on the energy and delay terms, but may have some light dependence on other things. For example, it may depend slightly on the timing of packets. There may even be some benefit for cheating based on competition between primary and secondary. This part addresses unmodeled dependence on P_{cheat} to see how different a cost function must be from those already considered in order for a new type of sanction to be required.

The dependence on P_{cheat} will have the same form as all of the previous cost functions. The secondary will get a $-\Delta P_{cheat}$ reduction per time step (or database token). This is then normalized by the number of messages sent per time step.

Let the new cost function be

$$C_{unknown} = C_{mixed} - \frac{\Delta P_{cheat}}{\text{messages per step}} \quad (4.82)$$

The necessary sanction to deter cheating can then be found as follows:

Theorem 19. *With the cost function in (4.82), the needed sanction to guarantee no secondaries cheat when $P_{tx} > P_{txMin}$ is:*

$$P_{pen} > 1 - \frac{P_{txMin} P_{catch} \beta (1 - \Delta)}{1 + \beta \Delta} \quad (4.83)$$

$$C_{sing} > \frac{(1 - P_{pen})(C_{UL} - 1 + \Delta)}{P_{txMin} P_{catch}} + \Delta \quad (4.84)$$

Proof. (4.82) satisfies the requirements of Lemma 1, so $P_{cheat} \in \{0, 1\}$. Therefore, the same procedure can be used as in previous theorems. Set the sanction to guarantee that the cost of the home band is less than the cost of cheating for all $P_{tx} > P_{txMin}$. The result follows. \square

Notice that it is not always possible to use P_{pen} and C_{sing} to deter cheating. In fact, it is only possible when $\Delta < 1$. This is exactly the same effect that motivated the singing sanction. The Δ in that case was the energy benefit of not having to sense. Because jail was inexpensive in terms of energy, there was nothing that could temper this extra Δ . An extra sanction was required.

The same thing will happen with other kinds of cost functions that are too different from what the sanction is targeting. Extremely time-sensitive QOS users will need some change to the jail model. Strong competition likely will need to be handled in a different way through extra regulation, etc, instead of trying to fix the problem through the jail system. We believe the QOS dimensions are limited to throughput, energy, and timing, so adding one more dimension to the jail sanction will deter cheating for all kinds of secondaries. This is left to future work.

If Δ is small enough – so the unmodeled phenomena is small or still sensitive to delay or energy-type sanctions – using jail to deter cheating results in the growth in overhead for honest users shown in Fig. 4.7. The overhead gets worse for all users except the primary because higher Δ results in higher sanctions. If the regulator wants to be able to deter cheating but keep the overhead small for honest users, the regulator can raise sanctions and invest in technology to lower P_{wrong} .




Figure 4.7: Overhead metrics as a function of the size of the margin for unmodeled phenomena Δ . The secondary overhead and spectrum hole metrics only rise sharply for high Δ , suggesting some robustness to additional margin. The primary collision metric gets better because higher Δ implies a higher sanction and therefore a larger range of protection P_{tx} .

Theorem 20. *With the cost function in (4.82), to keep the overhead the same as would be possible with C_{mixed} , we need to lower the observed \tilde{P}_{wrong} by the following:*

$$\frac{\tilde{P}_{wrong}}{P_{wrong}} = \left(\frac{1 - \Delta}{1 + \beta\Delta} \right) \left(\frac{k + (1 - k)(\beta(C_{UL} - 1))}{k + (1 - k) \left(\frac{\beta(1 - \Delta)(C_{UL} - 1 + \Delta)}{1 + \beta\Delta} + \Delta \right)} \right) \quad (4.85)$$

Proof. Set the overhead (4.16) equal using the required sanctions for the mixed (4.5) and unknown (4.82) case. Solve for the required P_{wrong} . \square

The required change in P_{wrong} is shown in Fig. 4.8 for energy and delay users.

4.4 A more refined model of harm

Up to this point, the interaction between the secondary and primary was very simple: the primary and secondary are co-located, any secondary transmission will cause all primary packets to drop, the secondary's sensing is perfect, and the two players are perfectly synchronized. Reality is not quite so simple, so this section considers a more refined model for the harm caused by the secondary.

The amount of harm caused by interference changes depending on several different effects. If the primary and secondary are not co-located, the secondary may be far enough away that transmitting at the same time as the primary causes dropped packets only during a particularly bad fade. The secondary may also have variable power, or it may beamform in a changing set of directions. These effects are captured in θ_{cheat} , the probability that a cheating transmission will actually cause harm.




Figure 4.8: The change in P_{wrong} required to keep the secondary overhead equivalent as Δ gets larger. If the regulator wants to increase protection against different kinds of utility functions while maintaining the same performance for the honest secondaries, it can pour money and effort into reducing the wrongful conviction rate.

Even when the secondary is trying to be honest, the primary will experience some degradation in its service. One major source of “honest harm” is when the secondary’s sensing mis-detects the primary. The secondary is doing everything it should, but it still causes harm. There are other sources too. For example, assume time is slotted, but the primary and secondary are not perfectly synchronized. Or, assume that time is not slotted, and the secondary is responsible for detecting when the primary turns on. In either case, there will be a delay between when the primary turns on, and when the secondary turns off. In [90], the authors refer to this effect as a required time margin that will be lost to interference when primary and secondary devices share spectrum. If the margin is too large, the primary will drop packets, even though the secondary is doing everything that the law requires.

To capture this harm the secondary causes when it is trying to be honest, define θ_{honest} as the probability that a primary will drop a packet due to honest operation by the secondary.³

To complete the model of harm to the primary, define θ_N as the probability of a dropped primary packet due to background noise.

This new model of harm changes the conception of trust. This thesis takes a “no harm no foul” stance on punishing secondaries. If cheating causes less harm, the secondary will be punished less often for the crime. This will be captured in an altered effective P_{catch} probability.

The real conceptual change comes in dealing with honest harm – if the primary is being harmed, should it matter whether the secondary thinks it is being honest? High honest harm should be treated no differently than cheating. However, it may be harder to catch. Before this point, P_{wrong} was a product of the identity system. It represented whether noise or other secondary cheating would be linked to an honest secondary. With honest harm, there is another effect: whether honest harm will be linked to the offending secondary. So, split the original P_{wrong} into

³This parameter is meant to capture only direct, physical harm in terms of dropped packets. The secondary may cause other kinds of harm, perhaps because of competition. This is discussed more in the next chapter.

two categories. $P_{wrong,other}$ will be the probability of being caught when the harm is caused by noise or other secondaries. Let $P_{catch,honest}$ be the probability of honest harm resulting in a jail sentence.

$P_{wrong,other}$ is the direct analog of the original P_{wrong} , and one of the goals of the identity system should be to make this as small as possible. $P_{catch,honest}$, on the other hand, should be thought of like P_{catch} . The goal is to have the identity system correctly match *any* harm caused by the secondary to the offending party. $P_{catch,honest}$ and P_{catch} are kept as separate quantities because honest harm may be more or less difficult to catch than cheating harm.

With these new pieces, define $P_{catch,eff}$ as the probability of being caught when cheating, and define $P_{wrong,eff}$ as the probability of being caught when being honest. These have the following equations:

$$P_{catch,eff} = \theta_{cheat}P_{catch} + (1 - \theta_{cheat})\theta_N P_{wrong,other} \quad (4.86)$$

$$P_{wrong,eff} = \theta_{honest}P_{catch,honest} + (1 - \theta_{honest})\theta_N P_{wrong,other} \quad (4.87)$$

Trusting the secondary to not cause too much harm now involves two parts: the secondary must not cause too much cheating harm and must not cause too much honest harm. Effectively, this means that the secondary must act honestly when its honest harm is low enough, and it must operate only in the home band when its honest harm is too high.

To guarantee this behavior, the sanctions must change to the following

Theorem 21. *To guarantee no delay or energy-sensitive secondary will cheat with the extended harm model, the sanctions must be:*

$$P_{pen} > \max \{1 - P_{txMin}\theta_{cheat,max}P_{catch}\beta, 1 - P_{txMin}\theta_{honest,max}P_{catch,honest}\beta\} \quad (4.88)$$

$$C_{sing} > \max \left\{ \frac{(1 - P_{pen})(C_{UL} - 1)}{P_{txMin}\theta_{cheat,max}P_{catch}}, \frac{(1 - P_{pen})(C_{UL} - 1)}{P_{txMin}\theta_{honest,max}P_{catch,honest}} \right\} \quad (4.89)$$

where $\theta_{cheat,max}$ is the maximum allowed cheating harm and $\theta_{honest,max}$ is the maximum allowed honest harm.

Proof. For the delay-sensitive users, first extend the sanction from Thm 5 to include the new catching probability:

$$P_{pen} > 1 - P_{txMin}\beta P_{catch,eff} \quad (4.90)$$

$$= 1 - P_{txMin}\beta(\theta_{cheat}P_{catch} + (1 - \theta_{cheat})\theta_N P_{wrong,other}) \quad (4.91)$$

$$> \max_{\theta_N, P_{wrong,other}} 1 - P_{txMin}\beta(\theta_{cheat}P_{catch} + (1 - \theta_{cheat})\theta_N P_{wrong,other}) \quad (4.92)$$

$$= 1 - P_{txMin}\beta\theta_{cheat}P_{catch} \quad (4.93)$$

The third inequality is because the sanction must be high enough to account for any noise or value of $P_{wrong,other}$ so that it is able to be set at certification time. As with P_{tx} in Chapter 2, if any level of θ_{cheat} must be accounted for, the sanction is an infinite jail sentence. Let $\theta_{cheat,max}$ be the maximum allowable cheating harm. Plugging in gives part of the result. This means that the primary must accept some fixed level of harm caused by the secondary. This can be interpreted, for example, as the allowed interference temperature [22].

The other half of the result comes from guaranteeing that if the secondary causes too much honest harm, it will choose to go to the home band. Using the cost function in Eq. (2.19), the sanction must be:

$$C_{D,3} < C_{D,2}(P_{cheat} = 0) \quad (4.94)$$

$$\frac{1}{\beta} < \frac{1 - P_{pen} + P_{tx}P_{wrong,eff}}{(1 - P_{pen})((1 - P_{tx})(1 - C_{DS}) + \beta)} \quad (4.95)$$

$$P_{pen} > 1 - P_{tx}\beta P_{wrong,eff}(1 - P_{tx})(1 - C_{DS}) \quad (4.96)$$

$$= 1 - \frac{P_{tx}\beta(\theta_{honest}P_{catch,honest} + (1 - \theta_{honest})\theta_N P_{wrong,other})}{(1 - P_{tx})(1 - C_{DS})} \quad (4.97)$$

Again, the sanction should be higher than the maximum of the right hand side:

$$P_{pen} > \max_{\theta_N, P_{wrong,other}, P_{tx} > C_{DS}} 1 - \frac{P_{tx}\beta(\theta_{honest}P_{catch,honest} + (1 - \theta_{honest})\theta_N P_{wrong,other})}{(1 - P_{tx})(1 - C_{DS})} \quad (4.98)$$

$$= 1 - P_{tx} \min_{\beta} \theta_{honest} P_{catch,honest} \quad (4.99)$$

but there must again be a limit to the protected honest harm, called $\theta_{honest,max}$. The second half of the delay sanction becomes:

$$P_{pen} > 1 - P_{tx} \min \theta_{honest,max} P_{catch,honest} \beta \quad (4.100)$$

The full delay sanction is the larger of these two possibilities.

For the singing sanction, first note that the sanction from Thm. 13 must now include $P_{catch,eff}$:

$$C_{sing} > \frac{(1 - P_{pen})(C_{UL} - 1)}{P_{tx} \min P_{catch,eff}} \quad (4.101)$$

$$> \max_{\theta_N, P_{wrong,other}} \frac{(1 - P_{pen})(C_{UL} - 1)}{P_{tx} \min (\theta_{cheat} P_{catch} + (1 - \theta_{cheat})\theta_N P_{wrong,other})} \quad (4.102)$$

$$> \max_{\theta_{cheat} > \theta_{cheat,max}} \frac{(1 - P_{pen})(C_{UL} - 1)}{P_{tx} \min \theta_{cheat} P_{catch}} \quad (4.103)$$

$$= \frac{(1 - P_{pen})(C_{UL} - 1)}{P_{tx} \min \theta_{cheat,max} P_{catch}} \quad (4.104)$$

where again, the maximum allowed θ_{cheat} was needed.

The second half of the energy sanction comes from guaranteeing that operating in the home band alone is preferable to causing too much honest harm to the primary, using the cost function in Eq. (4.4):

$$C_{E,3} < C_{E,1}(P_{cheat} = 0) \quad (4.105)$$

$$C_{UL} < \frac{(1 - P_{pen})((1 - P_{tx})(1 - C_{DS}) + C_{ES}) + P_{tx}P_{wrong,eff}C_{sing}}{(1 - P_{pen})(1 - P_{tx})(1 - C_{DS})} \quad (4.106)$$

$$C_{sing} > \frac{(1 - P_{pen})(1 - P_{tx})(1 - C_{DS})(C_{UL} - 1)}{P_{tx}P_{wrong,eff}} \quad (4.107)$$

$$C_{sing} > \frac{(1 - P_{pen})(1 - P_{tx})(1 - C_{DS})(C_{UL} - 1)}{P_{tx}(\theta_{honest}P_{catch,honest} + (1 - \theta_{honest})\theta_N P_{wrong,other})} \quad (4.108)$$

This sanction should be bigger than the maximum over the different parameters:

$$C_{sing} > \max_{P_{tx}, C_{DS}, \theta_{honest}, \theta_N, P_{wrong, other}} \frac{(1 - P_{pen})(1 - P_{tx})(1 - C_{DS})(C_{UL} - 1)}{P_{tx}(\theta_{honest}P_{catch, honest} + (1 - \theta_{honest})\theta_N P_{wrong, other})} \quad (4.109)$$

$$= \frac{(1 - P_{pen})(C_{UL} - 1)}{P_{tx} \min \theta_{honest, max} P_{catch, honest}} \quad (4.110)$$

The singing sanction must be larger than the max of the two possibilities, giving the result. \square

Notice first that these sanctions still depend only on parameters that can be known at certification time. Also, these sanctions require a further decision by the regulator: how much harm is too much harm? Implicitly, this is the decision the regulators have made in the TV bands by setting the no-talk radius. They have decided that any interference coming from a device a certain number of *km* away is not too much harm. It is also the same spirit of setting a limit on the interference temperature [22].

Finally, this concept is actually one of the first suggestions of how to regulate spectrum. De Vany *et al* [49] propose setting power limits at the edges of licensed regions. The trouble they encountered was how to enforce such a rule – it would need either monitors at the edges of the space or solving the hard problem of certifying such behavior before deployment. Spectrum jails allow regulators to make the decision of how much harm is allowed and then enforce that decision in a light-handed way, without setting explicit no-talk radii, etc.

We will not provide a technology improvement theorem here because in the limit of good technology, the new harm model fits into the theorems already proved (See Thm. 16). Perfect technology means that the honest harm, $\theta_{honest} \rightarrow 0$. Also, $P_{wrong, other} \rightarrow 0$. Because $P_{catch, honest}$ is analogous to P_{catch} for honest harm, it should only get higher with better technology.

Thm. 16 can then be immediately applied because the only other difference is θ_{cheat} , and this only appears as a smaller catching probability, which is already handled in the sanction.

4.5 Concluding remarks

4.5.1 What we have learned so far

This chapter extended the spectrum jails model to include a singing sanction to deter different kinds of secondaries from cheating, and it included a more refined model of harm. The sanction must address the type of QOS constraint the secondary uses, otherwise it will not be effective. But once the sanction is appropriate, the trust and performance results from the last chapter extend naturally.

The refined model of harm gives regulators the opportunity to declare the maximum allowable amount of harm (whether the secondary knows it is causing harm or not), and then enforce that level in a light-handed way. At this point, the prescription for controlling secondary behavior with spectrum jails is now complete.

The regulator chooses the operational parameters including the allowed amount of harm, and the protected level of primary activity. It then figures out the situational parameters – the size and quality of the home band and the allowed expansion by the secondary. Finally, it knows the capability of its identity system, which defines how well honest and cheating harm will be caught and how often the secondary will be sent to jail wrongfully. These parameters together define




Figure 4.9: Spectrum jails guarantee that the primary will be protected for all $P_{tx} > P_{txMin}$ and $\theta > \theta_{max}$ where the θ includes any honest or cheating harm from the secondary.

the required sanction, which can be determined *at certification time* and applied to any secondary *regardless of technical capability*.

The result of this prescription is shown in Fig. 4.9. The primary is protected from any high levels of harm as long as it transmits enough. However, it is expected to accept a certain small amount of harm, and there are no guarantees whatsoever if it is not transmitting enough.

From this guaranteed baseline protection, reality will allow the protection to get better as technology improves. If the secondary can find spectrum holes more effectively, it will have smaller sensing costs and cause a smaller amount of honest harm. If the regulator's identity system becomes better, it will more accurately punish only the secondary that actually caused harm.

In the limit of perfect technology, even using the same sanction set at certification time, there is no regulatory overhead caused by the jail system itself.

4.5.2 Future work – other kinds of secondaries

What other kinds of secondaries can there be? The device's utility is measured in its ability to transmit information with a desired QOS. Because all devices are sending information over a wireless medium, there is a kind of utility bottleneck – there are only so many kinds of QOS. It is therefore reasonable to believe that it is possible to design a spectrum jail that can deter cheating with *any* kind of secondary.

Obviously, energy and delay are important QOS parameters. This chapter gives sanctions to cover these desires. The parameter not considered in this thesis is related to timing. While both a large file and a streaming video desire high throughput, downloading a large file can accept long packet interruptions in the middle. Streaming a video cannot handle very bursty traffic without an unreasonably large buffer.

As another example, consider a transmission that has packets to send infrequently, but when it needs to send messages, it needs to complete the message in a short amount of time. This device will cheat to finish its transmission because the resulting jail sentence will be during the

long interlude when it has nothing to send.

To deter cheating in these kinds of users, spectrum jails may need to start sentences only when the secondary has something to send. Or, the sentences may need to be combined so that the user is sent to jail less often, but for longer amounts of time. Perhaps the jail sentence could even be applied at a random time in the future instead of immediately.

There is lots of work to be done to understand exactly what timing-dependent cost functions would look like, and what kind of sanction would be best to deter cheating by these secondaries. However, there is no reason to believe that the core result of this thesis, trust first and performance later, will break because of this extension.

Chapter 5

The role of the primary

The results of the thesis thus far can be summarized looking at Fig. 5.1. The primary is guaranteed protection from too much harm as long as it transmits often enough. Outside of this protected region, the secondary may cheat at some times or may cause some small amount of harm. As technology improves, the the actual protection region grows to include the entire plot. At the same time, the secondary spends fewer and fewer resources on wrongful conviction.

These results have assumed the the regulator, through the database, will take care of all enforcement activities. It will monitor for interference and refuse operation tokens as appropriate. The primary is treated as a helpless victim that must be protected. This is certainly true for the current primaries – these legacy systems were designed to operate in exclusive bands. New generations of primaries, on the other hand, will be designed to operate within a shared band. How will/should they interact with the enforcement system?

Think about the problem of rarely active primaries. The less often a primary is transmitting, the more tempting it is for a secondary to cheat. Legacy primaries cannot react intelligently to such information. But new generations of primaries may be designed to transmit more often than is actually required in order to protect themselves.

If the primary is already changing its behavior because of the presence of the secondary, can new responsibilities be added to make spectrum jails work better? The primary itself is in the best position to assess whether it is experiencing harmful interference. What if it had the chance to report this interference and be part of its own protection mechanism, instead of having to rely completely on the regulator to detect harm?

If the primary is responsible for reporting interference, the results in the previous chapters can be cast as the result of a game between the primary and the secondary. The secondary has an action space of choosing when to cheat. The primary's action space is limited to reporting interference whenever it drops a packet and not reporting interference otherwise. It has no other choice. This chapter turns this primary into an actual, rational player to question whether the primary has incentive to report interference responsibly.

Intuitively, and unfortunately, the answer is no. The primary will be affected by the presence of the secondary. Be it competition [86], a small rise in interference temperature [22], an overlap in time when the primary turns on before the secondary realizes it is there and turns off [90], or mis-detection in the secondary sensors, the primary will experience some harm. If it costs anything to detect and report interference, the primary is even affected negatively by having to spend resources to report interference. To mitigate these effects, a primary may choose to lie.




Figure 5.1: Spectrum jails guarantee that the primary will be protected for all $P_{tx} > P_{txMin}$ and $\theta > \theta_{max}$ where the θ includes any honest or cheating harm from the secondary.

This chapter formalizes the game referred to above between primary and secondary.¹ The goal is to extend the “trust first, performance as technology improves” theme to this new setting.²

With the game context, “trust” takes on a new meaning: trusting that the primary will be protected, and trusting that the primary will help to protect itself responsibly. The primary will only be protected if it is willing to report interference. The first part of this chapter shows that as long as the cost of reporting is not too high, the primary has incentive to report interference.

Trusting the primary to protect itself responsibly is a more difficult criterion. If the primary is able to report actual interference, it can also report interference when its packet was actually received (a behavior referred to here as “crying wolf”).³ This will raise the wrongful conviction rate, perhaps enough to cause the secondary to retreat to the home band. It can also transmit extra packets to pretend to be more active than it actually is.⁴ If the resulting primary transmission activity is high enough, this behavior also effectively kicks out the secondary.

The primary may have legitimate reasons for wanting to kick out the secondary. The harm

¹Game theory is widely used to understand the equilibrium behavior of cognitive devices in many different situations. For overviews of the general application of game theory to cognitive radio, see [132, 133]. However, it is mostly used as a mechanism to distribute resources. Beyond Etkin *et al* [103], this is the only exploration using game theory to understand trust and enforcement in Cognitive Radio.

²Many of the results in this chapter have been published in [134].

³This is an obviously undesirable behavior. With humans, false accusations are deterred through investigation and evidence, but these processes can put an emotional burden on the victim. So, in sensitive situations, the privacy of the victim becomes of utmost importance [135]. Indeed, [136] argues that abortion is legal partly so that a woman seeking an abortion does not need to face scrutiny of her motives. For this thesis, radios are emotionless entities, so an investigation into an accusation is not emotionally difficult. But it could be very costly to implement. Far better would be a solution that did not require identification of false accusations.

⁴Similar behavior — devices that mis-represent their necessary usage to gain some competitive advantage — has been observed elsewhere in technical literature, where the resulting problem is inefficient use of spectrum. [105] approaches the problem with bargaining, [104] proposes building altruism into its devices, and [137, 138] rely on regulation to enforce etiquettes. For this thesis, we believe extra transmissions may actually be the desired approach to get rid of a bothersome secondary.

caused by the secondary might be small enough to not be deterred by the jail system, but it may be too high for the primary service. This would particularly be true with safety-critical services like public safety networks. What should this primary be allowed to do? Crying wolf is an inherently wasteful behavior, but extra transmission can have a positive interpretation. The primary could sub-lease this extra transmission time to a preferred secondary that the primary trusts to respect its priority. This “band-sitter” can then keep other secondaries honest, or keep them in their home bands if they are harmful.

Trust of good primary behavior, then, can be interpreted as guaranteeing that the primary will choose to sub-lease its spectrum instead of crying wolf. This chapter will show that changing the cost of reporting can achieve this result, for a small enough rate of wrongful conviction.⁵

Finally, this chapter shows how performance will improve with better technology. As technology becomes perfect, the primary and secondary can coexist perfectly, and only the desired equilibria will be obtained.

5.1 The Model

The Qualitative Action Space

The full game is played on the following action space.⁶

Primary Player: The primary is a packet-based transmitter-receiver pair that is able to transmit one packet per unit time. It has something useful to transmit with probability \tilde{P}_{tx} . The packet will be received if it is not stopped by background noise or the action of the secondary.

- If primary has something to transmit, then it will transmit (with a cost of 1 unit of resource).
- If the primary does not have something to transmit, it will choose to stay quiet or to transmit anyway.
- If the primary packet is dropped, the primary chooses to report it as dropped (with a cost of C_{rf} units of resource) or to do nothing.
- If the primary packet is received, the primary chooses to report it as dropped (with cost C_{rf}) or to do nothing.

Secondary Player: The model for the secondary is the same as in Chapter 4, repeated here for ease of reading. The secondary always has something to transmit, and will receive any packet it transmits. It has a home band of size β data units per time, each of which costs C_{UL} units of energy to use. In its home band, the secondary does not have to respect primary priority.

- The secondary chooses to operate in the cognitive band, the home band, or both.
- If in the cognitive band, the secondary chooses to sense or not to sense.
- If it does not sense, the secondary chooses to transmit or not to transmit.
- If the sensing reports a primary packet, the secondary chooses to transmit or not to transmit.
- If the sensing reports a clear band, the secondary chooses to transmit or not to transmit.⁷

The actions are chosen given that the primary packet can be dropped for the following reasons:

⁵Can the cost of reporting actually be controlled? The primary is a radio with QOS constraints, just like the secondary. Reporting could include a mini-jail sentence. The details of this are left for future work.

⁶In future sections, the primary’s action space will be restricted in different ways to better understand the equilibria.

⁷Note that sensing is not always accurate – this effect is captured in the honest harm term introduced in the last chapter, θ_{honest} .

- Background noise causes a dropped packet with probability θ_N .
- A cheating secondary causes a dropped packet with probability θ_{cheat} .
- A secondary causes a dropped packet when honest⁸ with probability θ_{honest} .

The actions are also chosen given the presence of a jail system with the following properties:

- If the secondary is cheating and transmitting regardless of primary action, the secondary will be sent to jail with probability P_{catch} , which relies on the primary choosing to report the interference.
- If the secondary is acting honestly, it will still be sent to jail with probability P_{wrong} , which depends on the level of honest harm, and the primary choosing to report dropped packets.
- Once in jail, the secondary cannot transmit in the primary band or its home band. It stays in jail for an average $1/(1 - P_{pen})$ time steps, and must burn energy at a rate of C_{sing} per unit time while in jail.

Parameterizing the Action Space

The interaction is modeled as a Stackelberg game [139]. The primary chooses its strategy first assuming that the secondary will choose its strategy second to best respond to the primary. After both players have chosen their strategies, nature decides whether the primary has a packet to send, whether the action of the secondary will cause the primary packet to drop, and whether this results in the secondary being sent to jail. The interaction then continues with the same parameter choices for a long time. Both players are trying to minimize their expected cost in the game.

The goal is to find a Nash equilibrium, in which neither player has incentive to deviate from its strategy. This chapter will find a mixed strategy solution in which the primary chooses the marginals

- P_{rf} : the probability of reporting a true dropped packet.
- P_{cw} : the probability of crying wolf (reporting a correctly received packet as dropped).
- P_{tx} : the actual probability of transmitting given that the primary only needs to transmit with probability \tilde{P}_{tx} .

The secondary chooses the marginals

- P_{sense} : the probability that the secondary will sense.
- P_{cheat} : the probability that the secondary will transmit given it does not sense or senses a present primary and chooses to transmit anyway.

Cost Functions

For the examples in this chapter, the secondary will be a mixed user, defined by the cost function in Eq. (4.5). The qualitative results, and the intuition gained from them, will extend to all types of mixed users. However, the difficulty of analytically exploring this game requires that most of the results will be for an example only.

Chapter 4 already gives the best response of the secondary to the primary parameters. The secondary will play a unique pure strategy best response: it will either always sense or never sense. If it never senses, it will always cheat. If it always senses, it will never cheat.

⁸This parameter is meant to capture harm caused to the primary when the secondary is acting legally. The prototypical example for this game is when the secondary mis-detects the presence of the primary, and so transmits when it should stay silent. The secondary thinks it is acting honestly, but in fact causes harm. This effect causes a primary packet to drop with probability θ_{honest} .

The primary is trying to legitimately use the band for constructive purposes,⁹ but it will play the game to allow it to receive its messages most reliably. Its cost function is its total cost (in units of general transmission resources, where one unit of transmission costs 1 unit of resource) per unit of successfully received message:

$$C_P = \frac{P_{tx}(1 + C_{rf}(P_{msg}P_{cw} + (1 - P_{msg})P_{rf}))}{P_{msg}\tilde{P}_{tx}} \quad (5.1)$$

It has higher cost for higher P_{tx} , and spends a cost C_{rf} (cost of red flag) any time it makes an accusation, truthfully or not. In this equation, P_{msg} is the probability that the primary message was received. It is affected by the state of the secondary, and the background noise

$$P_{msg} = (1 - \theta_N)\mathbb{1}_{\text{secondary not in band}} \quad (5.2)$$

$$+ (1 - \theta_N)P_{cheat}(1 - \theta_{cheat}\pi_{NJ})\mathbb{1}_{\text{secondary in band}} \quad (5.3)$$

$$+ (1 - \theta_N)(1 - P_{cheat})(1 - \theta_{honest}\pi_{NJ})\mathbb{1}_{\text{secondary in band}} \quad (5.4)$$

where π_{NJ} is the stationary probability that the secondary is not in jail, given by:

$$\pi_{NJ} = \frac{1 - P_{pen}}{1 - P_{pen} + P_{toJail}} \quad (5.5)$$

where P_{toJail} is the probability of being sent to jail, which depends on how the secondary is caught.

The secondary is not sent to jail without an accusation, so the primary's parameters P_{rf} and P_{cw} determine the probability a secondary will be sent to jail. \tilde{P}_{catch} is the base probability of being caught after an accusation when a dropped packet is caused by cheating, determined by the identity system employed.

Just as in Chapter 4, there are two reasons a secondary may be sent to jail if it is not cheating. It may be caught causing honest harm, which will happen if the secondary does not accurately sense the primary. Let the base probability of being caught because of honest harm be $\tilde{P}_{catch,honest}$. This value should be thought of as comparable to \tilde{P}_{catch} . The secondary may also be sent to jail because of background noise or a false accusation implicate the secondary. Call the base probability of this $\tilde{P}_{wrong,other}$. Both of these probabilities depend on the identity system.

The observed probabilities of being sent to jail given a primary packet was sent are then:

$$P_{catch} = \theta_{cheat}P_{rf}\tilde{P}_{catch} + (1 - \theta_{cheat})\theta_N P_{rf}\tilde{P}_{wrong,other} \quad (5.6)$$

$$+ (1 - \theta_{cheat})(1 - \theta_N)P_{cw}\tilde{P}_{wrong,other} \quad (5.7)$$

if the secondary is cheating, and

$$P_{wrong} = \theta_{honest}P_{rf}\tilde{P}_{catch,honest} + (1 - \theta_{honest})\theta_N P_{rf}\tilde{P}_{wrong,other} \quad (5.8)$$

$$+ (1 - \theta_{honest})(1 - \theta_N)P_{cw}\tilde{P}_{wrong,other} \quad (5.9)$$

if the secondary is being honest.

The probability of being sent to jail is then

$$P_{toJail} = P_{tx}(P_{cheat}P_{catch} + (1 - P_{cheat})P_{wrong}). \quad (5.10)$$

The parameters of the game, when they are set, and which player knows them are shown in Table 5.1.

⁹This primary is not trying to behave like a spectrum troll [10], deriving benefit only from the secondary's misfortune.

Table 5.1: Parameters for the game between primary and secondary: cert. is set at certification time and needed for trust, run. is set at runtime and needed for performance, reg. is known to the regulator, pri. is known to primary, sec. is known to the secondary.

Parameter	Meaning	Time Set	Known to
C_{UL}	Energy cost of home band	cert.	everyone
P_{txMin}	Min protected P_{tx}	cert.	everyone
β	size of home band	cert.	everyone
\tilde{P}_{catch}	base prob. of catching	cert.	reg., pri.
P_{catch}	obs. prob. of catching when cheating	run.	pri., sec.
$\tilde{P}_{catch,honest}$	base prob. of conviction for honest harm	cert.	reg., pri.
$\tilde{P}_{wrong,other}$	base prob. of conviction when secondary does not cause harm	cert.	reg., pri.
P_{wrong}	obs. prob. of catching when honest	run.	pri., sec.
P_{pen}	prob. of staying in jail	cert.	everyone
C_{sing}	energy cost of jail	cert.	everyone
C_{rf}	cost of red flag	cert.	reg., pri.
C_{ES}	energy cost of sensing	run.	pri., sec.
C_{DS}	delay cost of sensing	run.	pri., sec.
k	sec. mix cost parameter	run.	pri., sec.
P_{sense}	prob. sec. senses	run.	pri., sec.
P_{cheat}	prob. sec. cheats	run.	pri., sec.
θ_N	prob of harm from noise	run.	pri.
θ_{cheat}	prob of harm from cheating	run.	pri.
$\theta_{cheat,max}$	max allowed cheating harm	cert.	reg., pri.
θ_{honest}	prob of harm from sensing	run.	pri.
$\theta_{honest,max}$	max allowed honest harm	cert.	reg., pri.
\tilde{P}_{tx}	required pri. tx prob.	run.	reg., pri.
P_{tx}	actual pri. tx prob.	run.	pri., sec.
P_{txMin}	minimum protected P_{tx}	cert.	reg.
P_{rf}	prob. of red flag	run.	pri.
P_{cw}	prob. of crying wolf	run.	pri.

Table 5.2: Values of the parameters used for different examples of the game between primary and secondary.

Parameter	Set (A)	Set (B)
\tilde{P}_{catch}	.6	.6
$\tilde{P}_{catch,honest}$.6	.6
$\tilde{P}_{wrong,other}$.1	.01
C_{UL}	5	5
P_{txMin}	.2	.2
$\theta_{cheat,max}$.2	.2
$\theta_{honest,max}$.2	.2
β	1	1
C_{ES}	.5	.2
C_{DS}	.1	.01
k	.8	.8
θ_N	.01	.005
θ_S	.5	.8
θ_{SI}	.1	.005
$P_{pen} = 1 - \beta P_{txMin} \tilde{P}_{catch} \theta_{cheat,max}$		
$C_{sing} = ((1 - P_{pen})(C_{UL} - 1)) / (P_{txMin} \tilde{P}_{catch} \theta_{cheat,max})$		

5.2 Guaranteeing trust

This section explores the equilibrium behavior of the primary. The goal is twofold: first, if the primary is required to report interference, can the regulator still trust that this primary will be protected? Second, can the primary be trusted to use its reporting power responsibly?

To this end, the section is split into two parts. The first explores the equilibrium behavior if the action space of the primary is restricted only to choosing P_{rf} , its probability of raising a red flag and reporting interference. If the primary can be trusted to report interference, the jail sentence is enough to guarantee the secondary will not cause too much harm.

The second part of this section shows how to use the cost of reporting to guarantee the primary will act responsibly. This is done by restricting the action space again, first to just choosing P_{rf} and P_{cw} and then allowing the full action space. This is done to understand the role of the cost of reporting and when the primary will choose to protect itself using false accusations versus extra transmission. The result is a procedure to make sure the cost of reporting is high enough that the primary will choose to use extra transmissions instead of crying wolf.

The examples in this section will use the parameters given in Table 5.2.

5.2.1 Trusted protection: primary can only raise red flags

When the primary can only raise a red flag, it is operating with the cost function

$$C_P = \frac{1 + C_{rf}(1 - P_{msg})P_{rf}}{P_{msg}} \quad (5.11)$$

The jail sanction is set so that the primary will be protected from too much secondary harm, as long as the primary is willing to report interference. So, the purpose of this section is to understand




Figure 5.2: Cost functions for a primary that can raise red flags, but cannot raise their P_{tx} or cry wolf. The bottom plot is cost by P_{rf} ; the top plot is cost by \tilde{P}_{tx} optimized over P_{rf} . The red flag is correctly used to stop the secondary from cheating and to kick the secondary out of the band when \tilde{P}_{tx} is high enough.




Figure 5.3: This is a guide for reading Figs. 5.4 and 5.9. Above both dashed lines in the pair, the secondary is transmitting only in its homeband. Between the two lines, the secondary is honestly using the cognitive band. Below the two lines, the secondary is cheating in the cognitive band.




Figure 5.4: Comparing the regions from Fig. 5.3 for primaries with different action spaces, using Parameter Set (A) from Table 5.2. The honest secondary (black dashed line) always raises a red flag, and is honest with the other parameters. The primary that can choose its probability of raising a red flag will follow the honest primary as long as the cost of red flag is low enough. The primary that can choose both P_{rf} and P_{cw} will use these to make the secondary not cheat and leave the band as soon as possible. Likewise, a primary that can choose P_{tx} as well will use this tool to kick out the secondary early as long as the cost of red flag is low. Notice that if the primary can choose only P_{rf}, P_{tx} , it coincides with the full game lines for high C_{rf} . This indicates that it is possible to choose C_{rf} such that the primary will prefer using P_{tx} over P_{cw} to control the secondary.

when the primary has incentive to report.

Fig. 5.2 shows the costs to determine the equilibrium P_{rf} . The bottom plot shows that different values of P_{rf} will cause the secondary to be honest and then leave the band. The associated cost drops sharply when the secondary changes state; otherwise, extra reporting costs more. For the bottom plot, the best choice of P_{rf} is the smallest value that will kick the secondary out of the band.

The upper plot of Fig. 5.2 shows the cost optimized over P_{rf} for different values of \tilde{P}_{tx} . The optimal use of reporting depends on the amount of transmission required: if the primary wants to transmit only rarely, keeping the secondary honest is the only choice. If the primary is transmitting more often, it is possible and desirable to kick the secondary out of the band.

The next two Figs. will serve as a comparison point for how the primary's different action spaces affect the secondary behavior. Fig. 5.3 gives the guide: each case will have two lines that show the secondary behavior in response to the primary's choices. Above both lines, the secondary will operate only in the home band. Between the lines, the secondary will honestly operate in the primary's band. Below the lines, the secondary will cheat in the primary's band. The example in Fig. 5.3 shows what the secondary will do in response to a perfectly honest primary that will always report interference and never cry wolf or use extra transmissions.

Fig. 5.4, then, compares the response of a secondary to a rational primary and the response of the secondary to a completely honest primary. The rational primary will use P_{rf} only as much as it needs to. In this example, the result is that the secondary behaves exactly as it would against a completely honest primary. So, this example rational primary can be trusted to protect itself.

With general parameters, to what extent can the primary be trusted to defend itself with reported interference? The primary will only pay as much as it needs to in order to effect the desired secondary behavior. But, if the cost to report is too high, the primary may choose to accept more interference instead of paying to reduce it. To understand the tradeoff better, consider the primary alone, without the complication of a full secondary player. Assume that an action on the primary's part will be able to change its probability of getting a message through from P_{msg} to P'_{msg} where

$$1 - P'_{msg} = (1 - P_{msg})(1 - \gamma). \quad (5.12)$$

In other words, it can reduce its probability of missing a message by a factor of $1 - \gamma$. This parameter could, for example, take into account the difference in packet-drop rate between the secondary cheating and being honest. The difference could be very large if the secondary causes lots of harm when cheating and almost none when honest. The primary is intuitively more likely to actually report interference in this case.

Lemma 7. *With the primary cost function in Eq. (5.11), and with γ defined in Eq. (5.12), the primary is willing to report up to*

$$P_{rf} < \frac{\gamma}{C_{rf}(1 - \gamma)P_{msg}} \quad (5.13)$$

Proof. Compare the primary cost function with P_{rf} and P'_{msg} to the cost function with P_{msg} and $P_{rf} = 0$:

$$\frac{1 + C_{rf}(1 - P_{msg})(1 - \gamma)P_{rf}}{1 - (1 - P_{msg})(1 - \gamma)} < \frac{1}{P_{msg}} \quad (5.14)$$

$$P_{rf} < \frac{\gamma}{C_{rf}(1 - \gamma)P_{msg}} \quad (5.15)$$




Figure 5.5: Maximum acceptable P_{rf} the primary is willing to use to get a γ decrease in dropped packets. If the cost of reporting is very low, the primary is willing to report more often to get the secondary to stop cheating or to exit the band.

□

This is plotted in Fig. 5.5 for different values of C_{rf} . Obviously the primary is willing to report more often if it gets a higher payoff (higher γ), but the payoff is tempered by the cost of reporting.

Therefore, in order to trust that the primary will be protected, the cost of reporting must be kept low enough. To know just how low C_{rf} must be, the primary must be willing to use $P_{rf} = 1$ to either make the secondary honest or kick it out of the band in any protected situation. This includes $\theta_{cheat} > \theta_{cheat,max}$, $\theta_{honest} > \theta_{honest,max}$, and $\tilde{P}_{tx} > P_{txMin}$.

Theorem 22. *To guarantee the primary will report interference to protect itself in any protected situation:*

$$C_{rf} < \min_{\theta_N, \theta_{honest} > \theta_{honest,max}, \theta_{cheat} > \theta_{cheat,max}, \tilde{P}_{tx} > P_{txMin}} \frac{\theta_{cheat} - \theta_{honest}}{(1 - \theta_{cheat})(1 - (1 - \theta_N)(\theta_{cheat} - \theta_{honest}))} \quad (5.16)$$

and

$$C_{rf} < \min_{\theta_N, \theta_{honest} > \theta_{honest,max}, \tilde{P}_{tx} > P_{txMin}} \frac{\theta_{honest}}{(1 - \theta_{honest})(1 - \theta_{honest}(1 - \theta_N))} \quad (5.17)$$

Proof. The two conditions on C_{rf} come from guaranteeing the primary will report enough to make the secondary switch from cheating to being honest, and to switch from being honest to operation in the home band if the honest harm is too high. These cases will be considered separately.

- Reporting enough for the secondary to be honest: to guarantee the primary will report in all cases of interest, assume that $P_{rf} = 1$ is required for the secondary to be honest, compared to $P_{rf} = 0$ when the secondary is cheating. Using Lemma 7, first find the γ with the change in

secondary state. Using Eq. (5.2), $P_{msg} = (1 - \theta_N)(1 - \theta_{cheat})$ when the secondary is cheating because with $P_{rf} = 0$, the secondary is never sent to jail. $P'_{msg} = (1 - \theta_N)(1 - \theta_{honest}\pi_{NJ,honest})$ when the secondary is honest. To find γ :

$$1 - P'_{msg} = (1 - P_{msg})(1 - \gamma) \quad (5.18)$$

$$1 - (1 - \theta_N)(1 - \theta_{honest}\pi_{NJ,honest}) = (1 - (1 - \theta_N)(1 - \theta_{cheat}))(1 - \gamma) \quad (5.19)$$

$$\gamma = \frac{(1 - \theta_N)(\theta_{cheat} - \theta_{honest}\pi_{NJ,honest})}{1 - (1 - \theta_N)(1 - \theta_{cheat})} \quad (5.20)$$

Then, Lemma 7 is used to find the condition on C_{rf} :

$$P_{rf} < \frac{\gamma}{C_{rf}(1 - \gamma)P_{msg}} \quad (5.21)$$

$$1 < \frac{\gamma}{C_{rf}(1 - \gamma)P_{msg}} \quad (5.22)$$

$$C_{rf} < \frac{\theta_{cheat} - \theta_{honest}\pi_{NJ,honest}}{(1 - \theta_{cheat})(1 - (1 - \theta_N)(\theta_{cheat} - \theta_{honest}\pi_{NJ,honest}))} \quad (5.23)$$

$$C_{rf} < \frac{\theta_{cheat} - \theta_{honest}}{(1 - \theta_{cheat})(1 - (1 - \theta_N)(\theta_{cheat} - \theta_{honest}))} \quad (5.24)$$

The last line comes from minimizing right hand side over $\pi_{NJ,honest}$ to remove the dependence on $\tilde{P}_{wrong,other}$. Noting that this condition must work for all protected cases gives the result.

- To guarantee that the primary will report to kick the honestly harmful secondary out of the band, assume that $P_{rf} = 1$ is required to kick out the secondary. Also assume that the secondary will be honest for $P_{rf} = 0$, and compare the two cases. As in the last part, Lemma 7 will be used. $P_{msg} = (1 - \theta_N)(1 - \theta_{honest})$ when the secondary is honest because it is never spending time in jail if $P_{rf} = 0$. $P'_{msg} = (1 - \theta_{noise})$ when the secondary is in the home band. To find γ :

$$1 - P'_{msg} = (1 - P_{msg})(1 - \gamma) \quad (5.25)$$

$$1 - (1 - \theta_N) = (1 - (1 - \theta_N)(1 - \theta_{honest}))(1 - \gamma) \quad (5.26)$$

$$\gamma = \frac{(1 - \theta_N)\theta_{honest}}{1 - (1 - \theta_N)(1 - \theta_{honest})} \quad (5.27)$$

Then, using Lemma 7:

$$P_{rf} < \frac{\gamma}{C_{rf}(1 - \gamma)P_{msg}} \quad (5.28)$$

$$1 < \frac{\gamma}{C_{rf}(1 - \gamma)P_{msg}} \quad (5.29)$$

$$C_{rf} < \frac{\theta_{honest}}{(1 - \theta_{honest})(1 - \theta_{honest}(1 - \theta_N))} \quad (5.30)$$

Again, C_{rf} must work for all values of these parameters of interest, which gives the result. \square

Notice that because $P_{txMin}, \theta_{honest,max}, \theta_{cheat,max}$ are already being chosen at certification time, as per Chapter 4, the cost of reporting can be determined at certification time to meet these requirements as well.

5.2.2 Trusting the primary to act responsibly

The primary has two ways in which it can act irresponsibly: it can cry wolf to report correctly received packets as dropped. It can also transmit extra packets to pretend to be productively using the band more than it actually needs to. This section is split into three parts. To understand when the primary has incentive to cry wolf, the first part restricts the primary's action space to choosing just P_{rf}, P_{cw} . The second part explores the full game to understand when the primary will choose to transmit extra packets. The final part shows explores choosing C_{rf} to encourage the primary to act responsibly.

Primary can raise red flags and cry wolf

In this part, the primary is restricted to choosing only P_{rf}, P_{cw} . It can only transmit when it has a message to send. So, the primary is using the cost function:

$$C_P = \frac{1 + C_{rf}(P_{msg}P_{cw} + (1 - P_{msg})P_{rf})}{P_{msg}} \quad (5.31)$$

Any regulator would want $P_{cw} = 0$. Crying wolf is completely wasted effort, producing completely wasted opportunity while the secondary sits in jail from false convictions. Unfortunately, the primary has incentive to cry wolf – in Chapters 2 to 4, higher wrongful conviction rates make the secondary more likely to operate just in the home band. If the primary can cry wolf and kick out the secondary with little extra cost, it will do so.

Fig. 5.6 shows the action space of the primary broken down. The bottom plot shows the choice of P_{cw} for a given \tilde{P}_{tx} and P_{rf} . Higher probability of crying wolf translates into a higher cost in general, but when the secondary transitions out of the band, the cost falls. The primary no longer has dropped packets because of honest secondary interference.

The middle plot shows the choice of P_{rf} , for a given \tilde{P}_{tx} and optimized over P_{cw} . At some value of P_{rf} , the secondary will stop cheating, and at a higher value it is possible to kick out the secondary by using P_{cw} . Notice that crying wolf will never be used to stop the secondary from cheating – crying wolf can only raise the rate of wrongful conviction. Chapter 2 showed that a higher wrongful conviction rate will actually make the secondary *more* likely to cheat up until the point where the secondary wants to leave the band.

Finally, the top plot shows the cost against \tilde{P}_{tx} , optimized over P_{rf}, P_{cw} . The two parameters are used together to first get the secondary to stop cheating and then to kick it out when possible.

Fig. 5.4 shows the reaction of the secondary to the optimized actions of the primary for the current case. The profile of when the secondary stops cheating does not change because crying wolf cannot change this line. However, the secondary is kicked out of the band for much smaller \tilde{P}_{tx} when the primary can cry wolf. Notice that for higher cost of reporting, the primary will kick out the secondary less quickly because it is getting too expensive to do so.

Intuitively, the cost of reporting can therefore be used to curb false accusations. To get a better handle on how this parameter can be used, consider a toy model for the primary that is similar to the toy used in the last part. Assume that if the primary cries wolf at a particular level, it is able to get a reduction in packet-drop rate of γ :

$$1 - P'_{msg} = (1 - P_{msg})(1 - \gamma) \quad (5.32)$$




Figure 5.6: Illustration of the choices the primary makes when it can report both dropped and received packets as interference. In the bottom plot, the primary can use P_{cw} to kick out the secondary, but it cannot use this tool to make the secondary stop cheating. The middle plot shows the choice over P_{rf} with the optimal P_{cw} . With higher P_{rf} , the secondary stops cheating and then in conjunction with higher P_{cw} , it will leave the band. The top plot shows the overall cost, and the optimal decision given an initial required value \tilde{P}_{tx} .

Then, a similar prescription is followed to understand how much the primary is willing to cry wolf in order to achieve the γ reduction in packet drops.




Figure 5.7: The maximum P_{cw} a primary will use to kick the secondary out of the band given a cost of reporting, and a fixed probability of reporting interference. As the cost of reporting falls, the acceptable P_{cw} rises. However, even if the cost of reporting is high, a big enough change in dropped packets (γ) will still provoke the primary to cry wolf.

Lemma 8. *Given a fixed P_{rf} , the primary is willing to cry wolf up to*

$$P_{cw} < \frac{\gamma(1 - P_{msg})(1 + C_{rf}P_{rf})}{C_{rf}P_{msg}(1 - (1 - P_{msg})(1 - \gamma))} \quad (5.33)$$

where γ controls the reduction in packet drop probability as in Eq. (5.32).

Proof. Assume there is a fixed value of P_{rf} that the primary has committed to in order to control cheating. Then, compare the cost function with $P_{cw} = 0, P_{msg}$ and the cost function with $P_{cw} \neq 0, msg'$. Here, γ can only represent the reduction in packet-drop rate if the secondary chooses to leave the primary band for the home band.

$$\frac{1 + C_{rf}(1 - P_{msg})P_{rf}}{P_{msg}} > \frac{1 + C_{rf}((1 - (1 - P_{msg})(1 - \gamma))P_{cw} + (1 - P_{msg})(1 - \gamma)P_{rf})}{1 - (1 - P_{msg})(1 - \gamma)} \quad (5.34)$$

$$P_{cw} < \frac{\gamma(1 - P_{msg})(1 + C_{rf}P_{rf})}{C_{rf}P_{msg}(1 - (1 - P_{msg})(1 - \gamma))} \quad (5.35)$$

□

The acceptable P_{cw} is shown in Fig. 5.7. Notice that as γ goes up, P_{cw} rises – this is expected because as the benefit to kicking out the secondary goes up, so does the primary's willingness to do anything to make it happen. Notice also that as the cost of reporting goes up, the willingness to cry wolf goes down. However, crying wolf can only be controlled so much by the cost of reporting. For very high C_{rf} , P_{cw} is determined only by γ , P_{msg} , and P_{rf} .

The optimal P_{rf} when crying wolf is simple to calculate. Regardless of the other parameters, the next lemma shows that if the primary is crying wolf in equilibrium, it will be reporting actual interference whenever it happens.

Lemma 9. *In equilibrium, if $P_{cw} > 0$ then $P_{rf} = 1$.*

Proof. Assume that the secondary is being honest, and the primary is trying to kick the secondary out of the band. The probability of going to jail is:

$$P_{toJail} = P_{tx}P_{wrong} \quad (5.36)$$

$$= P_{tx}(\theta_{honest}P_{rf}\tilde{P}_{catch,honest} + (1 - \theta_{honest})\theta_N P_{rf}\tilde{P}_{wrong,other} \quad (5.37)$$

$$+ (1 - \theta_{honest})(1 - \theta_N)P_{cw}\tilde{P}_{wrong,other}) \quad (5.38)$$

The secondary will leave if the probability of going to jail rises by an amount Δ :

$$P_{toJail} = P_{tx}(\theta_{honest}P_{rf}\tilde{P}_{catch,honest} + (1 - \theta_{honest})\theta_N P_{rf}\tilde{P}_{wrong,other} \quad (5.39)$$

$$+ (1 - \theta_{honest})(1 - \theta_N)P_{cw}\tilde{P}_{wrong,other}) + \Delta \quad (5.40)$$

The primary will choose to get this increase by either using P_{rf} or P_{cw} . The proof will show that it is always preferable to achieve the Δ rise using P_{rf} .

Using the equation above, if the primary uses P_{rf} to achieve the Δ rise, it must increase P_{rf} by an amount δ_{rf} :

$$\delta_{rf} = \frac{\Delta}{\theta_{honest}\tilde{P}_{catch,honest} + (1 - \theta_{honest})\theta_N\tilde{P}_{wrong,other}} \quad (5.41)$$

If the primary uses P_{cw} , it must increase it by an amount δ_{cw} :

$$\delta_{cw} = \frac{\Delta}{(1 - \theta_{honest})(1 - \theta_N)\tilde{P}_{wrong,other}} \quad (5.42)$$

Now, assume that the primary has a combination P_{rf}, P_{cw} that will successfully kick the secondary out of the band. Note that this means that the probability of the primary successfully receiving a message, $P_{msg} = 1 - \theta_N$, depends only on the noise. Then, keeping the same probability of going to jail, let the new $\tilde{P}_{cw} = P_{cw} - \delta_{cw}$ and the new $\tilde{P}_{rf} = P_{rf} - \delta_{rf}$. To check that this makes

the primary cost go down:

$$C_P(\tilde{P}_{rf}, \tilde{P}_{cw}) < C_P(P_{rf}, P_{cw}) \quad (5.43)$$

$$\frac{P_{tx}(1 + C_{rf}(P_{msg}\tilde{P}_{cw} + (1 - P_{msg})\tilde{P}_{rf}))}{\tilde{P}_{tx}P_{msg}} < \frac{P_{tx}(1 + C_{rf}(P_{msg}\tilde{P}_{cw} + (1 - P_{msg})\tilde{P}_{rf}))}{\tilde{P}_{tx}P_{msg}} \quad (5.44)$$

$$(1 - \theta_N)\tilde{P}_{cw} + \theta_N\tilde{P}_{rf} < (1 - \theta_N)P_{cw} + \theta_N P_{rf} \quad (5.45)$$

$$\theta_N\delta_{rf} - (1 - \theta_N)\delta_{cw} < 0 \quad (5.46)$$

$$\frac{\delta_{rf}}{\delta_{cw}} < \frac{1 - \theta_N}{\theta_N} \quad (5.47)$$

$$\frac{1 - \theta_N}{\theta_N} > \frac{(1 - \theta_{honest})(1 - \theta_N)\tilde{P}_{wrong,other}}{\theta_{honest}\tilde{P}_{catch,honest} + (1 - \theta_{honest})\theta_N\tilde{P}_{wrong,other}} \quad (5.48)$$

$$\theta_N(1 - \theta_{honest})\tilde{P}_{wrong,other} < \theta_{honest}\tilde{P}_{catch,honest} + \theta_N(1 - \theta_{honest})\tilde{P}_{wrong,other} \quad (5.49)$$

$$0 < \theta_{honest}\tilde{P}_{catch,honest} \quad (5.50)$$

Where the last line is always true. Because any movement to make P_{rf} bigger and P_{cw} smaller will result in a smaller cost, in equilibrium, if $P_{cw} > 0$, then $P_{rf} = 1$. \square

Although the sanction and honest reporting is enough to kick out secondaries with sufficiently high honest harm, this result can be interpreted in a different way. The primary may see the secondary as a threat because of competition, and it may cry wolf in order to hurt its competition. Maybe the temptation for the primary to hurt its competition is very strong. Unfortunately, there is nothing that can be done at this point. In order to make sure the primary will be protected in cases of physical harm, C_{rf} is limited in Thm. 22. It may not be possible to set a high enough C_{rf} to control crying wolf.

In either case, crying wolf is a response by the primary to a secondary that it cannot coexist with easily.

Crying wolf solves any coexistence difficulties by creating spectrum holes. Extra transmission, on the other hand, could solve coexistence difficulties in a way that does not sacrifice usable spectrum. The primary could sub-lease the extra transmission time to a secondary it can more easily coexist with. To understand when it will choose to transmit instead of crying wolf, the next part considers the full game between primary and secondary.

Full game

In this part, the primary can use its full range of actions: it can report dropped packets, honestly or not, and it can transmit as much as it likes above its required amount of transmission. So, it is operating on the full cost function

$$C_P = \frac{P_{tx}(1 + C_{rf}(P_{msg}P_{cw} + (1 - P_{msg})P_{rf}))}{P_{msg}\tilde{P}_{tx}} \quad (5.51)$$

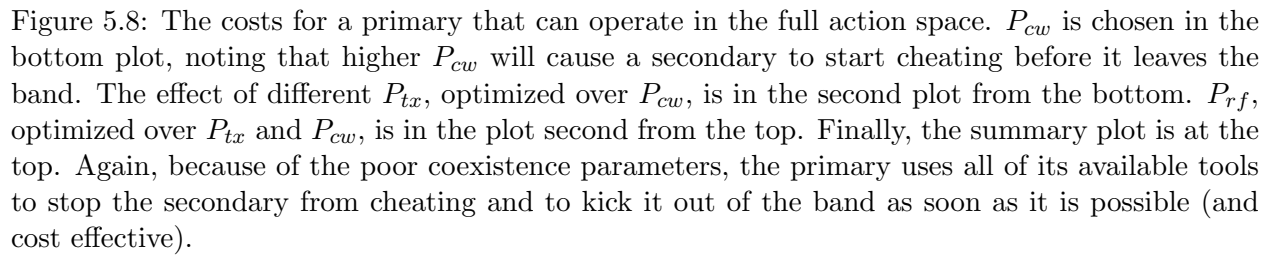


Figure 5.8: The costs for a primary that can operate in the full action space. P_{cw} is chosen in the bottom plot, noting that higher P_{cw} will cause a secondary to start cheating before it leaves the band. The effect of different P_{tx} , optimized over P_{cw} , is in the second plot from the bottom. P_{rf} , optimized over P_{tx} and P_{cw} , is in the plot second from the top. Finally, the summary plot is at the top. Again, because of the poor coexistence parameters, the primary uses all of its available tools to stop the secondary from cheating and to kick it out of the band as soon as it is possible (and cost effective).




Figure 5.9: When better coexistence parameters are used (Parameter Set (B) from Table 5.2), the secondary is easily deterred from cheating and the primary does not try as hard to kick the secondary out of the band.

The role of different primary actions is broken down in Fig. 5.8. At the bottom, P_{cw} is the same as before: it can be used to kick the secondary out of the band, but not to stop cheating. P_{tx} , on the other hand in the second-from-the-bottom plot, can be used in conjunction with P_{cw} to both stop cheating and kick the secondary out of the band. P_{rf} , in the next-up-plot, will be chosen to leverage the other two knobs most efficiently in achieving the desired action from the secondary.

Fig. 5.4 shows the game regions with the full game. Another option is also added: if the primary can change P_{rf}, P_{tx} but must have $P_{cw} = 0$. Notice that when C_{rf} is larger, the full game coincides with the $P_{cw} = 0$ lines. If C_{rf} is large enough, the primary will use P_{tx} alone to control the secondary's actions.

In the previous chapters, time slots when the primary is transmitting and time slots when interference is being policed have been treated as one in the same. This does not need to be true. In this chapter, the two are distinct – \tilde{P}_{tx} represents the time slots that the primary is actually transmitting. P_{tx} represents the time slots that are being protected from secondary interference. During these times, *the primary does not have to be the one that is being protected*. If sensing is implemented through a database, these slots could be left empty or given to a different primary. Alternatively, the primary could be selling these time slots to a chosen alternative secondary that coexists well with the primary.

So, for example, if the primary is not around very often, as with public safety users, the primary could hire a “band-sitter” so that other secondaries will be looking for the public safety primary when it does need the band. If the public safety user is around often enough, but it needs better protection from harm than the sanction alone can provide, the public safety user can also get rid of honestly harmful secondaries by protecting the transmissions of an approved band-sitter.

In any case, this problem is very similar to that which required the home band solution in the last three chapters: if it is too difficult for the primary and secondary to coexist, the enforcement system should not be used as a rationing mechanism. There must be an option for the players to

resolve the situation legally and peacefully. Allowing the primary to sell protected transmission slots to a chosen secondary fills that need here.

Then, the question becomes what level of C_{rf} is required to make the primary want to use P_{tx} instead of P_{cw} to control the actions of the secondary? The next part will address this question by exploring the needed C_{rf} .

To complete this part, assume that C_{rf} has been chosen to make $P_{cw} = 0$ the best option. Then, how much extra transmission is the primary willing to invest in a better equilibrium in the game?




Figure 5.10: The maximum factor increase in P_{tx} the primary is willing to use to produce a γ drop in packet drop rate. This assumes C_{rf} is high enough that it is better to use P_{tx} to control the secondary, rather than P_{cw} .

The improvement for the primary is modeled as a γ reduction in the rate of dropped packets, as in Eq. (5.32). Then

Lemma 10. Define A as a factor increase in transmission ($P_{tx} = A\tilde{P}_{tx}$) that the primary is willing to pay in order to get a γ reduction in packet drop rate. Then

$$A < \frac{(1 + C_{rf}(1 - P_{msg})P_{rf})(1 - (1 - P_{msg})(1 - \gamma))}{(1 + C_{rf}(1 - P_{msg})(1 - \gamma)P_{rf})P_{msg}} \quad (5.52)$$

Proof. Define $P_{tx} = A\tilde{P}_{tx}$. Then, assume a fixed P_{rf} . Compare the cost function with the reduction and higher P_{tx} to the cost function with original \tilde{P}_{tx} and no packet-drop reduction:¹⁰

$$\frac{\tilde{P}_{tx}(1 + C_{rf}(1 - P_{msg})P_{rf})}{P_{msg}\tilde{P}_{tx}} > \frac{A\tilde{P}_{tx}(1 + C_{rf}(1 - P_{msg})(1 - \gamma)P_{rf})}{(1 - (1 - P_{msg})(1 - \gamma))\tilde{P}_{tx}} \quad (5.53)$$

$$A < \frac{(1 + C_{rf}(1 - P_{msg})P_{rf})(1 - (1 - P_{msg})(1 - \gamma))}{(1 + C_{rf}(1 - P_{msg})(1 - \gamma)P_{rf})P_{msg}} \quad (5.54)$$

□

¹⁰Remember that the cost of any transmission is 1 unit of resource.

This gives a region of parameters P_{rf}, P_{tx} that the primary is willing to use to keep the secondary honest given a particular cost of reporting. This space is shown in Fig. 5.10. The primary will optimize the parameters within this space.

As a final note, the example parameters being used are particularly bad in order to emphasize the effects. Fig. 5.9 shows the game regions when the coexistence parameters are good (set (B) in Table 5.2). If there are good coexistence parameters, unless it is essentially free to kick out the secondary, the primary will happily use only those parameters required to keep the secondary from cheating and allow it to remain in the band. This effect will be discussed more fully while analyzing performance in Section 5.5.

5.2.3 Setting C_{rf} to deter crying wolf

It is difficult to get an analytical solution to what C_{rf} is required to make P_{tx} a better option for the primary than P_{cw} . So, this section walks through an example case with plots to explore the problem. Analytically finding the necessary C_{rf} is left to future work.

Because the primary is trying to keep its cost low, it will choose the smallest cost combination of P_{cw}, P_{tx}, P_{rf} to kick out the secondary. Lemma 9 showed that if $P_{cw} > 0$ then $P_{rf} = 1$, so it is sufficient to look at combinations of P_{cw}, P_{tx} with $P_{rf} = 1$ that successfully make the secondary move to the home band. These are shown in Fig. 5.11 for different levels of $\tilde{P}_{wrong,other}$. Any parameter pair above the line will make the secondary transmit only in the home band. Any parameter pair below the line will have the secondary in the primary band. The examples all use parameter set (A) from Table 5.2 with $\theta_{honest} = 0.01$ to make the secondary more difficult to kick out.




Figure 5.11: For an example secondary, and different values of $\tilde{P}_{wrong,other}$, the solid lines show the tradeoff between P_{tx} and P_{cw} above which the secondary will move to operating only in its home band. The dashed line shows the minimum P_{tx} required achieve a switch. Note that the original \tilde{P}_{tx} does not affect the line – to kick out the secondary, the primary must raise its transmission probability to this level, regardless of what activity level is actually desired.

The primary is trying to minimize its cost, so it will choose the lowest cost (P_{cw}, P_{tx})

pair along the line. Fig. 5.12 takes the (P_{cw}, P_{tx}) pairs that can kick out the secondary when $\tilde{P}_{wrong,other} = 0.01$, and shows the cost against the level of P_{cw} for different values of C_{rf} . Notice that as C_{rf} gets higher, the lowest cost choice changes from $P_{cw} = 1$ to $P_{cw} = 0$. The necessary C_{rf} , then, is the lowest value for which $P_{cw} = 0$ is the best choice. In this example, $C_{rf} > 0.359$.




Figure 5.12: These are example primary costs for different levels of C_{rf} . The plot is done against P_{cw} , but P_{tx} is also changing to trace out the line in Fig 5.11 that will actually make the secondary move to the home band. For this example, the $\tilde{P}_{wrong,other} = 0.01$ line is used. The cost of reporting must be high enough that the lowest cost is at $P_{cw} = 0$. In this example, that level is $C_{rf} > 0.359$.

For any value of $\tilde{P}_{wrong,other}$, the minimum necessary C_{rf} can be found such that the primary will prefer using P_{tx} rather than P_{cw} to force the secondary to retreat to the home band. This is shown in Fig. 5.13. Notice that as the $\tilde{P}_{wrong,other}$ rises, so does the needed cost of reporting. This makes intuitive sense: if $\tilde{P}_{wrong,other}$ is high, then crying wolf even a little bit is very effective at raising the effective P_{wrong} . A larger C_{rf} is therefore needed to make the cost of crying wolf prohibitive.

This result leads to two natural conclusions. First, the regulator should choose C_{rf} as high as is allowed to encourage the primary to use sub-leasing instead of crying wolf for a larger range of $\tilde{P}_{wrong,other}$. In other words, the cost of reporting must be the highest allowed in Thm. 22.

Second, this result leads naturally to an interpretation for $\tilde{P}_{wrong,other}$. These wrongful convictions are coming from wrongful accusations. If the secondary is able to build a strong alibi for why it was actually not causing interference, the rate of wrongful accusations turning into jail sentences will go down. Perhaps, then, the identity system should focus on catching correctly, but building in a method to allow alibis.




Figure 5.13: The C_{rf} needed to make the primary want to sub-lease its band instead of crying wolf is dependent on the value of $\tilde{P}_{wrong, others}$. This means that it is not possible to define a C_{rf} that will deter crying wolf for all parameter combinations. The best the regulator can do is set C_{rf} at the highest value that does not violate the trust requirement in Thm. 22. Then, when the identity system is good enough, the correct primary behavior will happen.

5.3 Improving Performance

5.4 Choosing to split the band

Extra transmissions were desirable because they allow the primary to more easily defend its band from secondary harm. But, if the primary is incapable of selling those extra transmissions, perhaps P_{tx} can be raised in a different way – the primary could declare that it will never use a part of its band (here modeled as a percentage of time). For example, it could declare that it will use only even time steps, and leave the odd time steps always open for secondary use. If the database or other entity can propagate this information, the secondary has look-ahead information on the primary, and the secondary can use the open slots as part of its home band.¹¹

For illustration here, assume that the band can only be split in half. So, the primary first chooses whether it will use all or only half of its band. Then the game is played as before.

The result of this new game is shown in Fig. 5.14. The original game regions, and the regions for a completely honest primary are shown for comparison. When the band is split, it is far easier to defend because the primary is active more often and the secondary has more to lose when it goes to jail. So the secondary will stop cheating at a much lower \tilde{P}_{tx} . Likewise, it is not important to kick the secondary out at the odd times, so that half of the band is still used until $\tilde{P}_{tx} \geq 0.5$, at which time splitting the band is no longer an option. Fig. 5.15 shows the cost to the primary in the original game and after allowing splitting for different values of \tilde{P}_{tx} . Being able to split the band is obviously beneficial to both primary and secondary.

¹¹In Chapter 3, the home band was modeled as a safe slot within the primary bands. One way to implement the safe slot is by having the primary declare this slot always empty.




Figure 5.14: If the primary is allowed to split its band and give up half the time to unlicensed use, it is left with a piece of spectrum that is much more easily defensible than its original allocation. So, it more easily stops the secondary from cheating. It also is able to share its band more easily with the secondary, so the primary does not try to kick out the secondary.




Figure 5.15: When the primary is allowed to split its band, it does not have to work as hard to protect itself. So, it has lower cost.

5.5 Improvement with technology

As in the last chapters, the goal here is to understand how the performance metrics change as technology changes. Within the game setting, the technology parameters for performance include those from the last chapters, $\tilde{P}_{wrong,other}, C_{DS}, C_{ES}$. They also include the coexistence parameter θ_{honest} , because if technology allows the secondary to more effectively find spectrum holes, this same technology should cause less impact to the primary player.

Some of the metrics from the last chapters are also important here:

- Secondary time in jail, $\pi_J = \frac{P_{toJail}}{1 - P_{pen} + P_{toJail}}$ where P_{toJail} is defined in Eq. (5.10)
- Total spectrum holes, $holes_{total} = holes_{jail} + holes_{homeband}$ where the two components are defined in Eqs. (4.18) and 4.19.

To determine performance for the primary player, the collisions metric from the last chapters is replaced by the primary's packet drop probability, $1 - P_{msg}$. Two more metrics are added:

- Extra transmissions, $P_{tx} - \tilde{P}_{tx}$
- Primary cost spent on reporting, $C_{rf}(P_{msg}P_{cw} + (1 - P_{msg})P_{rf})$

to complete the picture of whether the primary is accepting the current secondary and how hard it is working to keep that secondary in line.

Fig. 5.16 shows these metrics against P_{pen} for different values of $\theta_{cheat}, \theta_{honest}$. For simplicity, the secondary is assumed to be purely delay-constrained ($k = 1$). The plots are generated by doing a Monte Carlo simulation to average the metrics over the types of secondaries, C_{DS} , and the desired transmit rate of the primary, \tilde{P}_{tx} .

The best case is when $\theta_{cheat} = 1, \theta_{honest} = 0$, as it should be: the secondary is easy to catch and the primary has no incentive to kick it out of the band. On the other extreme, $\theta_{cheat} = 0.3, \theta_{honest} = 0.5$, the secondary is so hard to catch and causes so much honest harm, that either the secondary is cheating, or it is in the home band. There is little opportunity for honest coexistence.

In general, given a level of cheating harm, all of the metrics seem to get better as θ_{honest} goes down. The next two theorems show that this indeed the case. The first theorem gives the equilibrium behavior of the primary in the limit of good technology and shows that the primary will be honest. The second theorem shows that as technology becomes perfect, the performance becomes perfect as well.

Theorem 23. *If $\theta_{honest}, \tilde{P}_{wrong,other}, C_{DS}, C_{ES} \rightarrow 0$, the primary's equilibrium action choices are:*

1. $P_{cw} = 0$
2. $P_{tx} = \tilde{P}_{tx}$

Proof. The two parts of the theorem are taken separately. Thm. 16 gives that in the desired limit, the secondary will not cheat for any $P_{tx} > 0$ and will go to the home band only if $P_{tx} = 1$, as long as there is some chance of going to jail when cheating. With this action space for the secondary, the effect of the primary's choices are limited, which leads to the result:

1. P_{cw} affects the probability of going to jail through P_{catch} (Eq. (5.6)), and P_{wrong} (Eq. (5.8)). In both cases, P_{cw} only changes the probabilities if $\tilde{P}_{wrong,other} > 0$. So, as $\tilde{P}_{wrong,other}, P_{cw}$ has no effect on the probability of being sent to jail. Using Lemma 8, this means that $\gamma = 0$, and the primary is not willing to choose anything other than $P_{cw} = 0$.

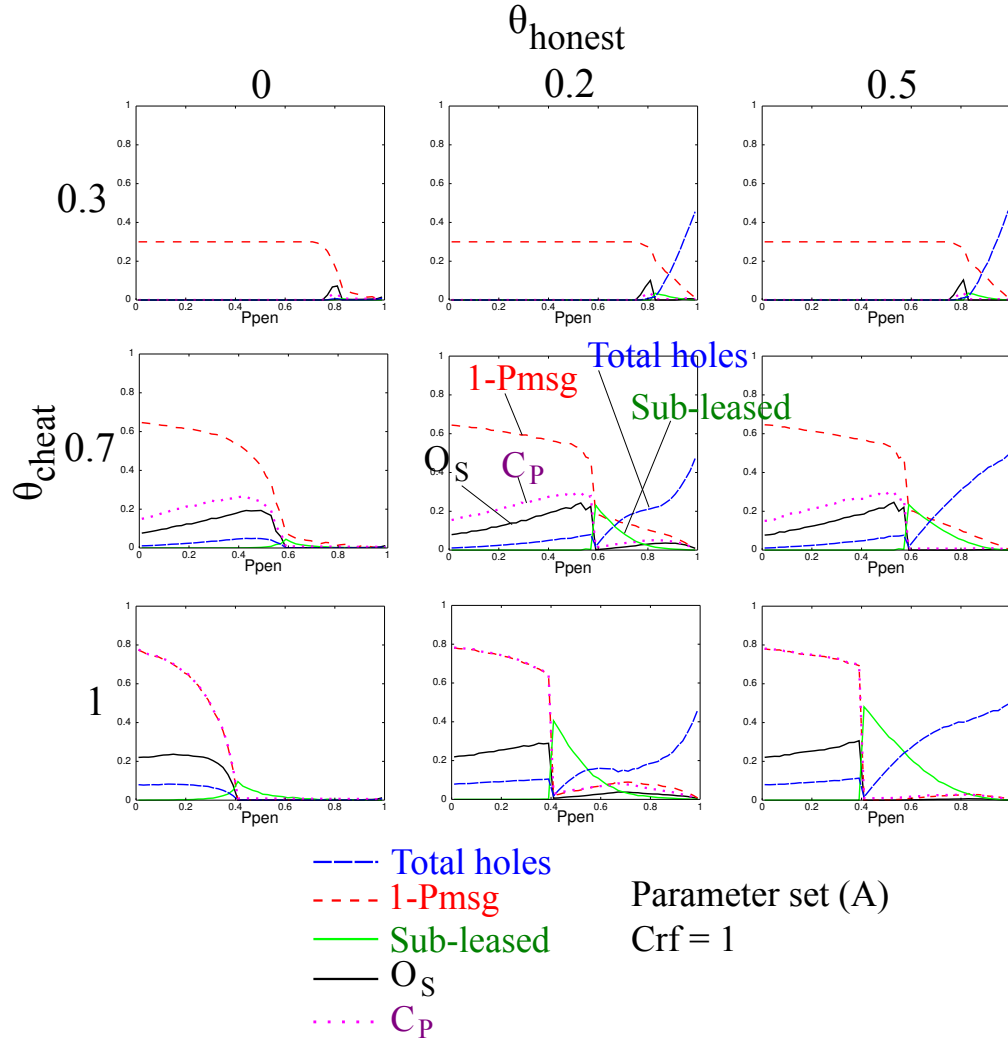


Figure 5.16: Performance metrics against P_{pen} assuming the secondary is purely delay-constrained ($k = 1$). The plots use the Monte Carlo method to average the metrics over the types of secondaries (parameterized by C_{DS}) and the types of primaries (parameterized by P_{tx}). In general, the most important factor for improving performance is to reduce the level of honest harm, θ_{honest} .

2. In the desired limit, the secondary will never cheat and will only go to the home band if $P_{tx} = 1$. Because $\theta_{honest} \rightarrow 0$, the packet drop probability does not change regardless of what the primary chooses as P_{tx} . Using Lemma 10, with $\gamma = 0$, the primary will choose $A = 1 \Rightarrow P_{tx} = \tilde{P}_{tx}$.

□

Theorem 24. *If $\theta_{honest}, \tilde{P}_{wrong,other}, C_{DS}, C_{ES} \rightarrow 0$, the metrics change as follows:¹²*

1. $\pi_J \rightarrow 0$
2. $P_{msg} \rightarrow (1 - \theta_N)$
3. $holes_{total} \rightarrow 0$
4. $(P_{tx} - \tilde{P}_{tx}) \rightarrow 0$

Proof. Using the equilibrium primary and secondary actions given in Thms. 23 and 16:

1. As in Thm. 8, it is sufficient that $P_{wrong} \rightarrow 0$:

$$\lim_{\tilde{P}_{wrong,other}=0, \theta_{honest}=0} P_{wrong} = \lim_{\tilde{P}_{wrong,other}=0, \theta_{honest}=0} \theta_{honest} P_{rf} \tilde{P}_{catch,honest} \quad (5.55)$$

$$+ (1 - \theta_{honest}) \theta_N P_{rf} \tilde{P}_{wrong,other} \quad (5.56)$$

$$+ (1 - \theta_{honest}) (1 - \theta_N) P_{cw} \tilde{P}_{wrong,other} \quad (5.57)$$

$$= 0 \quad (5.58)$$

2. If the secondary is always in the primary band and honest, and $\theta_{honest} \rightarrow 0$, and using the last part:

$$\lim_{\theta_{honest}, P_{cheat}, \pi_J \rightarrow 0} P_{msg} = \lim_{\theta_{honest}, P_{cheat}, \pi_J \rightarrow 0} (1 - \theta_N) (1 - P_{cheat}) (1 - \theta_{honest} \pi_{NJ}) \quad (5.59)$$

$$= (1 - \theta_N) \quad (5.60)$$

3. For $holes_{total}$, note that the secondary will never leave for the home band. As the first part of this Thm. shows, it will also spend no time in jail. So, the total holes must go to zero.
4. Thm. 23 shows that in the desired limit, $P_{tx} = \tilde{P}_{tx}$.

□

5.6 Concluding remarks

5.6.1 What we have learned so far

The primary is in the best position to determine whether harmful interference has occurred, so it is natural to wonder whether the primary should participate in its own protection. This chapter has explored whether the primary should be responsible for reporting interference.

Even though the primary only controls accusations, and does not control exactly which secondaries end up in jail, the primary has several choices at its disposal. It can report interference; it can report interference even when its packet was successfully received; it can even pretend that

¹²Letting the $\tilde{P}_{wrong,other} \rightarrow 0$ can be thought of as the secondary having the ability to create a perfect alibi. Even if the primary accuses the secondary of causing harm, the secondary will not suffer jail if it did not actually cause harm.

it needs to transmit more than is actually required. These tools can be used to keep the secondary honest, and kick out any secondaries that are too bothersome.

This chapter has shown that this primary can be trusted to correctly defend itself if the regulator chooses the cost of reporting correctly. It must be low enough that the primary wants to report actual interference. However, the cost of reporting should be as high as allowed so that the primary has more incentive to use a band-sitter, instead of crying wolf, to kick the secondary out of the band.

Once the cost of reporting has been set, the equilibrium achieved in the game, and therefore the performance of the players, is determined by the current state of technology. Even with the same sanction and cost of reporting, as technology becomes perfect, all spectrum holes can be filled and the primary's only source of dropped packets is noise.

5.6.2 Future work

Understanding the primary's role in enabling spectrum sharing is currently in its infancy. Future work in this direction includes a more accurate model of the primary player, better limits on the needed cost of reporting, understanding direct competition, and exploring other ways to improve this game.

Fleshing out the primary cost function

Right now, the primary cost is defined relative to its personal cost of transmitting a message. This is done in this thesis to help understand how the enforcement game works without the complication of a full characterization of the primary. But primaries are radios, just like the secondary devices, so they too have real cost functions based on their QOS. The results here need to be extended to understand the behavior of different kinds of primaries.

The cost of reporting

Related to the different kinds of primaries is the characterization of the cost of reporting. This chapter has shown that the regulator needs to have control over the primary's cost of reporting. One way to do this would be to implement a primary mini-jail that requires some fixed degradation in QOS for every interference report the primary sends. This implementation needs to be explored in light of a more realistic QOS-based model for the primary.

This chapter has also left incomplete the determination of the required cost of reporting. We have provided a prescription for determining the smallest cost that guides a primary to sell time slots instead of crying wolf. This needs to be made analytical to ensure that the correct cost of reporting can in fact be determined at certification time.

Direct competition

This chapter has ignored any effects of direct competition between primary and secondary. If dropped packets are not the only harm a secondary can cause, the primary will have more incentive to kick out the secondary. On one hand, the results here still hold: the primary will still prefer to use extra transmission instead of crying wolf. However, direct competition may still present a problem.

If the primary feels it is in direct competition with all secondaries, it may choose to transmit gibberish instead of selling the extra transmission slots. This is extremely wasteful as the spectrum is not being used for productive purposes and transmission resources are being wasted. What should the regulator's stance be? Is the primary entitled to have a band to itself if it is willing to constantly transmit to protect itself? Is it even possible to stop such behavior?

This thesis also argues that it is good for the primary to be able to choose the secondaries it is willing to coexist with. From a technical standpoint, this is certainly true: primaries should not have to coexist with harmful secondaries if there are other secondaries that will not interfere as much.

From a competition viewpoint, this stance is more questionable. Should the primary be allowed to coexist only with its favored business partners? Light handed regulations that provide low barriers to entry may not actually give enough spectrum to new services if primaries are capable of denying access to their bands.

This is going to be a problem even if the primary is not involved in enforcement. The current approach to whitespaces makes no provision against primaries transmitting gibberish to keep out the secondaries. PCAST [79], in suggesting receiver standards, has started the conversation about limiting the primary's ability to hoard its band by dishonest means. Involving the primary in enforcement will take away the technical incentive to hoard the band. But dealing with the effects of competition is an important open problem.

Other ways to improve the game

This chapter has shown that if the primary is capable of declaring that it will only use half of its original band, everything works better: the resulting smaller band is easier to defend, and the vacated band is easy for secondary devices to use. Are there other strategies that can ease coexistence without requiring difficult improvements to existing technology?

Chapter 6

Sharing risk amongst multiple secondaries

With a single secondary and single primary, enforcement through spectrum jails is about choosing a maximum level of harm and designing a sanction so that it is in the best interest of the secondary to respect that choice. With multiple secondaries, it is clear that the same notion of harm should apply: the primary should be protected from too many dropped packets. But what does this mean for the secondaries?

Dropped packets occur roughly when the interference temperature [22] is too high, which means the aggregate interference is too much. The acceptable interference temperature can therefore be thought of as a resource that must be split amongst the transmitting secondaries. What is the desired split? Should all secondaries beyond a no-talk radius be able to transmit? Do the transmitting secondaries need to be thinned so that the density is not too high? How should this thinning occur? Without even a concept of the desired equilibrium, more basic questions need to be answered before spectrum jails can be extended to multiple secondaries.¹

This chapter begins the discussion of spectrum jails for multiple secondaries by exploring three pieces of the puzzle: first, this chapter explores the distribution of the aggregate interference that spectrum jails are supposed to control. Second, it investigates different choices in splitting the interference resource amongst secondaries based on the aggregate interference. Last, it shows a way forward in adapting spectrum jails to deal with multiple secondaries.

Modeling aggregate interference to explore placement risk

The FCC's current TV whitespace rules ignore aggregate interference completely [6], leading to [31] claiming that the current rules will not actually protect TV receivers sufficiently. Other options for rules, including no-talk radii and power control have been explored [31, 140, 141]. But the current literature ignores a potentially significant effect: what happens if the secondaries are all clustered at the edge of the no-talk radius? This thesis refers to the risk of high interference due to clusters of secondaries as *placement risk*.

The problem with the current literature is that the models are mostly deterministic – the secondaries are all placed in grids or seas of interferers. So, fading and shadowing can be accounted

¹These same questions will need to be answered for *any* enforcement system. The enforcement system cannot enforce undefined rules.

for [140], but placement cannot be properly addressed. Recent work in stochastic geometry has provided the tools to analyze aggregate interference coming from randomly placed interferers [142, 143]. There has been some work applying these ideas to cognitive radio [144], but no work addressing how these insights can be used to help design good sharing rules.

Section 6.2 explores the distribution of aggregate interference, focusing on how it changes as the no-talk radius moves. There is a qualitative change that occurs, depending on the no-talk radius and the density of secondaries. The results show that rural and urban environments require different levels of conservativeness in designing a no-talk-radius.²

Distributing placement risk

Section 6.3 explores who should bear the risk of poor placement of secondary . Three possibilities are analyzed, with the choice hinging on what decisions the FCC will allow the database to make at runtime.

If the database cannot make decisions at runtime, then a fixed no-talk radius must include a margin so that the probability of too much interference from poor placement is kept low. The risk is then shared between the primaries when poor placements do occur, and the close secondaries who cannot transmit because of a very conservative rule.

If the no-talk radius can be changed, then it can be moved based on the placement information held by the database. In this case, the primary can be protected from poor placement, and the secondaries close to the primary are affected only when necessary.

Finally, if the power can be changed based on the database's placement information, again the primary can be protected. The secondaries can share the risk of poor placement amongst *all* secondaries, instead of just those close to the primary.³

Adapting spectrum jails

Once the desired equilibrium has been chosen, can spectrum jails be used to achieve it? A very simple toy model is shown in the future work, Section 6.4.3, to illustrate how the theory of spectrum jails can be extended in the future using a queuing model.

6.1 The models

The generic model is consistent with TV whitespace ideas, and is shown in Fig. 6.1. There is a TV at the origin, TV receivers within a protected radius r_p of the origin, and secondary devices (interferers) outside of the no-talk radius r_n . The goal is to understand the total interference experienced at the TV receivers.

For simplicity, assume that $r_p = 0$, so TV transmitter and receivers are co-located at the origin. This is the worst case from an interference perspective: if $r_p = 0$, then the interferers are equally far away in all directions. If $r_p \neq 0$, then consider a receiver at the edge of the protected radius (this is the receiver most vulnerable to interference). In one direction, the interferers are $r_n - r_p$ away, but in the other direction, the interferers are $r_n + r_p$ away. The $r_p \neq 0$ configuration will produce less total interference.

²Most of the results from this section have been published in [145].

³Work with a similar goal, which does not use stochastic geometry tools is here [146].




Figure 6.1: Model setup: there is a TV transmitter at the origin, receivers within the protected radius r_p , no transmitters between r_p and r_n , and whitespace devices outside r_n placed according to a specified placement model.

Table 6.1: Parameter definitions for aggregate interference from multiple secondaries.

Parameter	Definition
α	Path loss exponent
λ	Poisson density parameter
I_T	Total interference seen at the origin
P_j	Transmit power of interferer j
r_j	Distance of interferer j from the origin
r_n	No-talk radius
P	Power density
μ	Exponential fading parameter

A standard path-loss propagation model is used to find the interference from three secondary placement models: a deterministic model, a dithered placement model, and a 2D Poisson model. For reference, the parameters are all listed in Table 6.1.

6.1.1 Per-node interference power model

The path-loss model assumes the power observed at the origin from interferer j at a distance r_j is given by:

$$I_j(r_j) = X P_j r_j^{-\alpha} \quad (6.1)$$

where P_j is the power of the interferer, α is the path-loss exponent, and X is the fading parameter. For all calculations with fading, $X \sim \text{Exp}(\mu)$ (Rayleigh fading). For numerical calculations, $P_j = 1$, $\alpha = 2.5$, and $\mu = 1$. The total amount of interference seen at the origin from all the interferers as

$$I_T = \sum_j I_j(r_j). \quad (6.2)$$

where r_j is determined by the placement model.

6.1.2 Node placement models

Deterministic models

This thesis uses two deterministic models: a sea and a modified grid.

In the sea model, there exists a secondary power density P per unit area over the entire space outside of the no-talk radius r_n .

The modified grid model is shown in Fig. 6.2. The area outside the no-talk radius, r_n , is divided into rings of width d . Within each ring, interferers are placed around the inside edge of the ring such that each interferer is the only node within an area d^2 .

To investigate the effect of placing nodes at different places within the ring, define $m = r_n/d$. Assume that the first ring is placed at a distance $(m + \delta)d$ from the origin. Each interferer has power Pd^2 .




Figure 6.2: 2D deterministic placement model: the first ring of interferers is located $(m + \delta)d$ away from the center, and each subsequent ring is d away from the previous ring. The number of nodes in each ring is chosen so that there is one node per d^2 area.

Dithered model

The dithered model takes the modified grid and randomly dithers δ to allow each interferer to have a limited random placement. Each interferer will have an *iid* $\delta_{j,k}$ for interferer j within ring k .

Let δ be distributed such that the interferer is placed uniformly in an area d^2 that is a slice of annulus surrounding it, as illustrated in Fig. 6.2. The distribution of δ is given by

Lemma 11.

$$f_\delta(x) = \frac{2(m + k + x)}{2(m + k) + 1}, \text{ for } 0 \leq x \leq 1. \quad (6.3)$$

Proof. The area is a wedge of the annulus with radius $(m+k)d$ to $(m+k+1)d$ and angle θ . The area can be made to be d^2 by appropriate choice of θ . This does not affect the distribution of δ , so in general, the area is:

$$\frac{\theta}{2}(((m+k+1)d)^2 - ((m+k)d)^2) \quad (6.4)$$

$$= \frac{\theta}{2}d^2(2(m+k)+1) \quad (6.5)$$

The uniform density over this area is

$$\frac{2}{\theta d^2(2(m+k)+1)} \quad (6.6)$$

To get the distribution of δ , for $0 \leq \delta \leq 1$:

$$P(\delta \leq x) = \int_0^\theta \int_{(m+k)d}^{(m+k+x)d} \frac{2}{\theta d^2(2(m+k)+1)} r \, dr \, d\phi \quad (6.7)$$

$$= \frac{d^2}{d^2(2(m+k)+1)} ((m+k+x)^2 - (m+k)^2) \quad (6.8)$$

$$= \frac{x^2 + 2(m+k)y}{2(m+k)+1} \quad (6.9)$$

The PDF is obtained by differentiating. □

All other parameters are the same as in the deterministic modified grid model.

2D Poisson Model

Secondary devices are placed outside the no-talk radius r_n according to a 2D Poisson process [142] with rate λ , which can be thought of as a density parameter. For comparison with the dithered model, set $\lambda = 1/d^2$. Each interferer has a per-node power $P_j = P/\lambda$, $\forall j$ so that the average power density remains P .

6.2 Near vs far field: understanding the distribution of aggregate interference

From a regulator's perspective, being able to accurately model the aggregate interference is crucial to being able to select a no-talk radius that adequately protects the primary receivers. So, the goal of this section is to understand the distribution of aggregate interference and how it relates to the no-talk radius.

Regardless of the placement model, the interference distribution will converge to a Gaussian distribution as the node density rises. Near and far field refer to the closeness of the actual distribution to Gaussian, with far field being very close to Gaussian. This Gaussian insight very useful for regulators. Calculating the actual distribution is resource intensive; the results in this section give guidance as to when a simpler approximation is sufficient.

6.2.1 Calculating the interference power

This part calculates and plots the expressions for the total interference power for each of the models. Expected value and variance are also given to provide a first idea how the distribution changes with λ and r_n .

Deterministic models

The total interference with deterministic placement models is

Theorem 25. *The interference experienced from a sea of interferers is:*

$$I_{sea} = \frac{2\pi P}{d^{\alpha-2}} \frac{1}{m^{\alpha-2}(\alpha-2)} \quad (6.10)$$

For the modified grid model, the total interference I_{grid} is given by:

$$I_{grid} = \frac{\pi P}{d^{\alpha-2}} \sum_{k=0}^{\infty} \frac{2(m+k)+1}{(m+k+\delta)^\alpha} \quad (6.11)$$

Proof. P is taken as a power density, and the area of the annulus from radius $(m+k)d$ to radius $(m+k+1)d$ is $\pi d^2(2(m+k)+1)$. Therefore, the interference can be calculated as:

$$I_{grid} = \sum_{k=0}^{\infty} \frac{P\pi d^2(2(m+k)+1)}{((m+k+\delta)d)^\alpha} \quad (6.12)$$

$$= \frac{\pi P}{d^{\alpha-2}} \sum_{k=0}^{\infty} \frac{2(m+k)+1}{(m+k+\delta)^\alpha} \quad (6.13)$$

$$I_{sea} = \int_0^{2\pi} \int_{md}^{\infty} \frac{P}{x^\alpha} x \, dx \, d\theta \quad (6.14)$$

$$= 2\pi P \int_{md}^{\infty} x^{1-\alpha} \, dx \quad (6.15)$$

$$= \frac{2\pi P}{\alpha-2} \frac{1}{(md)^{\alpha-2}} \quad (6.16)$$

□

Fig. 6.3 shows the aggregate interference from the best and worst case grids ($\delta = 1$ and 0 , respectively) compared to the sea model. This plot keeps the power density $P = 1$ and the no-talk radius $r_n = 1$ constant.

Notice that the best and worst case grids converge to the sea as d decreases. The closer together the nodes, the more sea-like they appear. This is the first indication of near and far field — in near field, the nodes are spaced far enough apart that their individual placement matters. In far field, they are so close together that changes in δ do not affect the total interference much.




Figure 6.3: Keeping the power density and the no-talk radius constant, this is the interference power for the deterministic placement models as m grows (nodes get closer together). The closest grid has $\delta = 0$ and the farthest has $\delta = 1$. As the nodes are packed more closely, the interference from all deterministic models converges to the sea model.

Dithered model

The total interference is the convolution of the interference distribution coming from each of the interferers. Their distributions can be found as the following.

Theorem 26. *For the 2D dithered grid, each interferer j within ring k has a distribution of interference:*

$$f_{I_{j,k}}(y) = \frac{2P^{2/\alpha}}{\alpha d^2(2(m+k)+1)} \frac{1}{y^{2/\alpha+1}} \quad (6.17)$$

For all y such that:

$$\frac{P}{(d(m+k+1))^\alpha} \leq y \leq \frac{P}{(d(m+k))^\alpha} \quad (6.18)$$

Proof. For one secondary in the 2D dithered grid, the interference $I_{k,j}$ is:

$$P(I_{k,j} \leq y) = P\left(\frac{P}{((m+k+\delta)d)^\alpha} \leq y\right) \quad (6.19)$$

$$= P\left(\delta \geq \left(\frac{P}{d^\alpha y}\right)^{1/\alpha} - (m+k)\right) \quad (6.20)$$

$$= \int_{\frac{P}{d^\alpha y}^{1/\alpha} - (m+k)}^1 \left(\frac{2(m+k)}{2(m+k)+1} + \frac{2x}{2(m+k)+1}\right) dx \quad (6.21)$$

$$= 1 + \frac{(m+k)^2}{2(m+k)+1} - \frac{1}{2(m+k)+1} \left(\frac{P}{d^\alpha y}\right)^{2/\alpha} \quad (6.22)$$

using Lemma 11. Therefore, by differentiating,

$$f_{I_{j,k}}(y) = \frac{d}{dy} P(I_{j,k} \leq y) \quad (6.23)$$

$$= \frac{2P^{2/\alpha}}{\alpha d^2(2(m+k)+1)} \frac{1}{y^{2/\alpha+1}} \quad (6.24)$$

This is true for values of y between:

$$y \leq \frac{P}{(d(m+k))^\alpha} \quad (6.25)$$

$$y \geq \frac{P}{(d(m+k+1))^\alpha} \quad (6.26)$$

$$(6.27)$$

by the fact that $\delta \in [0, 1]$. \square

The distribution can be numerically approximated by convolving the distribution from the first sufficiently many interferers. Let k is the index of the ring (with $k = 0$ being the first ring). Clearly, as k grows, the difference in interference power between $\delta = 0$ and $\delta = 1$ shrinks. At the same time, the interference power is getting smaller. Therefore, the interference power from ring k converges to a delta function at 0 as $k \rightarrow \infty$. As long as a sufficient number of interferers are included in the convolution, the rest will have a negligible effect on the shape or statistics of the distribution.




Figure 6.4: Example aggregate interference distributions for the 2D dithered grid model. If d is small (nodes are packed tightly), the distribution resembles a Gaussian. If d is large (nodes are spread out), then the distribution is heavy-tailed.

This numerical approximation is shown in Fig. 6.4. Notice that for large values of m , the distribution looks qualitatively Gaussian. We refer to this distribution as being in far field.

When m is small, the distribution is heavy-tailed. We refer to this distribution as being in near field. Is there a corresponding approximation to the near-field distribution? When m is small, d is very large compared to r_n , so the interference is dominated by the closest interferer. The right approximation is, therefore, the distribution of the interference coming only from the closest secondary. Let I_{NFD} be the interference power of this closest secondary in the dithered model.

Theorem 27. *The distribution of the interference from the single closest secondary given the dithered model, I_{NFD} is:*

$$f_{I_{NFD}}(x) = \left(\left(\frac{\lambda}{P} \right)^{2/\alpha} (2/\alpha) I^{2/\alpha-1} + 2(m - r_n) \left(\frac{\lambda}{P} \right)^{1/\alpha} (1/\alpha) I^{1/\alpha-1} \right) (2m + 1)^{-1}, \quad (6.28)$$

$$\text{for } \frac{P}{\lambda(r_n + d)^\alpha} \leq x \leq \frac{P}{\lambda(r_n)^\alpha} \quad (6.29)$$

Proof. Let r be the radial distance of the closest secondary, and N be the number of secondaries in the annulus closest to the primary. So, $N = \lceil \pi(1 + 2r_n/d) \rceil$. Using the distribution of δ from Lemma 11, with $k = 0$:

$$P(I_{NFD} < I) = P\left(r > R \mid \frac{P}{\lambda R^\alpha} = I\right) \quad (6.30)$$

$$= P\left(\text{no secondaries closer than } R \mid \frac{P}{\lambda R^\alpha} = I\right) \quad (6.31)$$

$$= P(\delta > R - r_n)^N \quad (6.32)$$

$$= \left(1 - \frac{(R - r_n)^2 + 2m(R - r_n)}{2m + 1}\right)^N \quad (6.33)$$

$$= \left(1 - \frac{((\frac{I\lambda}{P})^{1/\alpha} - r_n)^2 + 2m((\frac{I\lambda}{P})^{1/\alpha} - r_n)}{2m + 1}\right)^N \quad (6.34)$$

Because the radial width of the closest ring is d , this is only valid for $0 \leq R \leq d$, or $\frac{P}{\lambda(r_n + d)^\alpha} \leq I \leq \frac{P}{\lambda(r_n)^\alpha}$. The density is obtained by differentiating. \square

The approximations are shown in Fig. 6.5 for very large and very small d respectively, keeping the power density constant.

As a final note, the expected value of interference from the dithered grid is the same as that coming from the sea model. Define $I_{d.grid}$ as the total interference coming from the dithered model.

Theorem 28. *For the dithered model*

$$E[I_{d.grid}] = I_{sea} \quad (6.35)$$

Proof. Split the sea into wedges of area d^2 such that these wedges coincide with the wedges over which each node is dithered in the dithered grid model. Then it is sufficient to show that the interference coming from a subset of wedges has the same expected value for sea and dithered grid.

Because the dithered grid has the same distribution for all secondaries in a ring, the whole ring can be grouped together and thought of as an interference power coming from a ring $(m + k + \delta)d$




Figure 6.5: Approximate distributions for the dithered grid placement model. If the nodes are placed close together, the right approximation is a Gaussian distribution. If the nodes are spread out, the right approximation is the interference caused by a single randomly-placed interferer.

away from the center. The distribution of the dither in the radial direction was given in Lemma 11. Therefore, the expected value of the interference from one interferer in the 2D dithered grid is

$$E[I_{d,grid,ring\ k}] = E\left[\frac{P\pi d^2(2(m+k)+1)}{(m+k+\delta)^\alpha d^\alpha}\right] \quad (6.36)$$

$$= \frac{P\pi(2(m+k)+1)}{d^{\alpha-2}} E[(m+k+\delta)^{-\alpha}] \quad (6.37)$$

$$= \frac{P\pi(2(m+k)+1)}{d^{\alpha-2}} \int_0^1 \frac{2(m+k+x)}{(m+k+x)^\alpha(2(m+k)+1)} dx \quad (6.38)$$

$$= \frac{2\pi P}{d^{\alpha-2}} \int_0^1 (m+k+x)^{1-\alpha} dx \quad (6.39)$$

$$= \frac{2\pi P}{(\alpha-2)d^{\alpha-2}} \left(\frac{1}{(m+k)^{\alpha-2}} - \frac{1}{(m+k+1)^{\alpha-2}} \right) \quad (6.40)$$

For the sea:

$$I_{sea,ring\ k} = \int_0^{2\pi} \int_{(m+k)d}^{(m+k+1)d} \frac{P}{r^\alpha} r dr d\theta \quad (6.41)$$

$$= 2\pi P \int_{(m+k)d}^{(m+k+1)d} r^{1-\alpha} dr \quad (6.42)$$

$$= \frac{2\pi P}{(\alpha-2)d^{\alpha-2}} \left(\frac{1}{(m+k)^{\alpha-2}} - \frac{1}{(m+k+1)^{\alpha-2}} \right) \quad (6.43)$$

□

This theorem, along with the law of large numbers, implies that as $d \rightarrow 0$, the interference

coming from the dithered grid will converge to its mean, and will therefore be the same as the interference produced by a sea.

Poisson model

Given a per-node power $P_j = P/\lambda$ (which has no fading), the distribution of the total interference with the the Poisson model, I_{Pois} is given through its Laplace transform

Theorem 29. *The Laplace Transform of the interference seen at a point from a set of interferers distributed as a 2D Poisson process with rate λ outside of a ring r_n from the point of observation is*

$$L_{I_{Pois}}(t) = \exp \left(-\lambda 2\pi \int_{r_n}^{\infty} z \left(1 - e^{-P_j t / z^\alpha} \right) dz \right). \quad (6.44)$$

Proof. See Corollary 2.3.2 in [142] where the Laplace Transform of the power from each interferer as

$$L_P(tP_j/z^\alpha) = e^{-tP_j/z^\alpha} \quad (6.45)$$

because there is no fading, so the CDF is a step function at the power of the interference with path loss. \square




Figure 6.6: Example aggregate interference distributions for the Poisson placement model with no fading. As λ gets larger (nodes are more dense), this distribution also closely resembles a Gaussian.

Fig. 6.6 uses inversion equation (4.5) of [147] to numerically invert this Laplace transform for a range of λ , keeping the power density fixed. Notice that as in the dithered grid case, small λ distributions have heavy tails, while higher λ produce distributions that look qualitatively Gaussian.⁴

Like the dithered model, the heavy tails in near field are caused by the distribution being dominated by the interference from the closest secondary. To make a precise approximation to the true distribution, the distribution of the interference coming from the closest secondary, I_{NF} :

⁴The heavy tail in near field was noticed also in [144].

Theorem 30. *With the Poisson placement model, the distribution of the power of the interference coming from the closest interferer, I_{NF} is:*

$$f_{I_{NF}}(x) = \frac{2\pi P^{2/\alpha} \lambda^{1-2/\alpha}}{\alpha x^{2/\alpha+1}} \exp(-\lambda \pi ((P/(\lambda x))^{2/\alpha} - r_n^2)) \quad (6.46)$$

$$\text{for } x \leq \frac{P}{\lambda r_n^\alpha} \quad (6.47)$$

Proof. Let r be the radial distance of the closest secondary. Then

$$P(I_{NF} < I) = P(r > R | I = P/(\lambda R^\alpha)) \quad (6.48)$$

$$= P(\text{no interferers within radius } R | I = P/(\lambda R^\alpha)) \quad (6.49)$$

$$= \exp(-\lambda \pi (R^2 - r_n^2)) \quad (6.50)$$

$$= \exp(-\lambda \pi ((P/(\lambda I))^{2/\alpha} - r_n^2)) \quad (6.51)$$

Because the closest interferer can be no closer than r_n , the interference I is bounded: $0 \leq I \leq P/(\lambda r_n^\alpha)$.

The distribution is obtained by differentiating. \square




Figure 6.7: Approximate distributions for the interference caused with the Poisson placement model. For high λ (nodes are dense), the approximation is a Gaussian. If λ is small, the right approximation is the distribution of the interference coming from a single interferer.

Fig. 6.7 shows the near and far field approximations for the Poisson model, given very small and very large λ , keeping the power density constant.

The following theorem finds the mean and the variance of the interference distribution.

Theorem 31. *For a set of 2-D $\text{Poisson}(\lambda)$ distributed interferers with no-talk radius r_n , the mean*

and variance of the total interference I_{Pois} is given by

$$E[I_{Pois}] = \frac{2\pi\lambda P_j}{(\alpha - 2)r_n^{\alpha-2}} \quad (6.52)$$

$$var[I_{Pois}] = \frac{2\pi\lambda P_j^2}{(2\alpha - 2)r_n^{2\alpha-2}} \quad (6.53)$$

Proof. Divide the space into concentric annuli of width dr , starting from r_n . Let r_j be the distance from the center to annulus j and $N_j \sim Poiss(\lambda 2\pi r_j dr)$ be the number of nodes in annulus j . Then

$$E[I_T] = E \left[\sum_{j=1}^{\infty} \sum_{i=1}^{N_j} \frac{P_j}{r_j^{\alpha}} \right] \quad (6.54)$$

$$= \sum_{j=1}^{\infty} \frac{P_j}{r_j^{\alpha}} E[N_j] \quad (6.55)$$

$$= \sum_{j=1}^{\infty} \frac{P_j}{r_j^{\alpha}} \lambda 2\pi r_j dr \quad (6.56)$$

$$= \int_{r_n}^{\infty} \frac{2\pi\lambda P_j}{r^{\alpha}} dr \quad (6.57)$$

$$= \frac{2\pi\lambda P_j}{(\alpha - 2)r_n^{\alpha-2}} \quad (6.58)$$

Using the same division of space into annuli:

$$var[I_T] = var \left[\sum_{j=1}^{\infty} \sum_{i=1}^{N_j} \frac{P_j}{r_j^{\alpha}} \right] \quad (6.59)$$

$$= var \left[\sum_{j=1}^{\infty} \frac{N_j P_j}{r_j^{\alpha}} \right] \quad (6.60)$$

$$= \sum_{j=1}^{\infty} \frac{P_j^2}{r_j^{2\alpha}} var[N_j] \quad (6.61)$$

$$= \sum_{j=1}^{\infty} \frac{P_j^2}{r_j^{2\alpha}} 2\pi r_j dr \lambda \quad (6.62)$$

$$= 2\pi\lambda P_j^2 \int_{r_n}^{\infty} r^{1-2\alpha} dr \quad (6.63)$$

$$= \frac{2\pi\lambda P_j^2}{(2\alpha - 2)r_n^{2\alpha-2}} \quad (6.64)$$

□

Notice that the mean is the same as the interference from the sea in Thm. 25, given $P = P_j\lambda$. This mean decreases with r_n as $1/r_n^{\alpha-2}$, and the standard deviation (square root of the

variance) is decreasing as $1/r_n^{\alpha-1}$. Because the standard deviation is falling faster than the mean, the distribution *concentrates* about the mean as r_n increases. However, there is no difference in the qualitative behavior when r_n is small versus when it is large.

The effect of fading

What happens if fading is included? To keep the narrative simple, this thesis will only show fading with the Poisson model. The total interference from a Poisson model with exponential fading ($I_{Pois,fade}$) is:

Theorem 32. *The Laplace Transform of the interference seen at a point from a set of interferers distributed as a 2D Poisson process with rate λ outside of a ring r_n from the point of observation with $\text{Exp}(\mu)$ fading and per node power $P_j = P/\lambda$ is*

$$L_{I_{Pois,fade}}(t) = \exp \left(-\lambda \int_0^{2\pi} \int_{r_n}^{\infty} \frac{ztP_j}{P_jt + \mu z^\alpha} dz d\theta \right) \quad (6.65)$$

Proof. See Corollary 2.3.2 in [142] with the Laplace Transform of the power from each interferer as

$$L_P(tP_j/z^\alpha) = \frac{\mu}{P_jt/z^\alpha + \mu} \quad (6.66)$$

because the fading is exponential (Rayleigh). □




Figure 6.8: Distribution of the interference coming from nodes placed according to the Poisson model with fading. Even with fading, the distribution is more Gaussian-like for higher λ , and more heavy-tailed for smaller λ .

Fig. 6.8 shows the inverted transform (the PDF) of the distribution with fading. Again, the distribution seems to converge to a Gaussian.

With fading, the mean and variance can be calculated as follows:

Theorem 33. For a set of 2-D Poisson-distributed interferers of density λ with “no-talk” radius r_n , each with independent exponential fading $X_i \sim \text{Exp}(\mu)$, the mean and variance of the total interference I_T is given by:

$$E[I_T] = \frac{2\pi\lambda P_j}{\mu(\alpha - 2)r_n^{\alpha-2}} \quad (6.67)$$

$$\text{var}[I_T] = \frac{4\pi\lambda P_j^2}{\mu^2(2\alpha - 2)r_n^{2\alpha-2}} \quad (6.68)$$

where $P_j = P/\lambda$ is the power of each interferer, and α is the path-loss exponent.

Proof. Divide the space into concentric annuli of width dr , starting from r_n . Let r_j be the distance from the center to annulus j and $N_j \sim \text{Poiss}(\lambda 2\pi r_j dr)$ be the number of nodes in annulus j . Then

$$E[I_T] = E \left[\sum_{j=1}^{\infty} \sum_{i=1}^{N_j} \frac{X_i P_j}{r_j^{\alpha}} \right] \quad (6.69)$$

$$= \sum_{j=1}^{\infty} E[N_j] E[X_i] \frac{P_j}{r_j^{\alpha}} \quad (6.70)$$

$$= \frac{2\pi\lambda P_j}{\mu(\alpha - 2)r_n^{\alpha-2}} \quad (6.71)$$

Using the same division into annuli:

$$\text{var}[I_T] = \text{var} \left[\sum_{j=1}^{\infty} \sum_{i=1}^{N_j} \frac{X_i P_j}{r_j^{\alpha}} \right] \quad (6.72)$$

$$= \sum_{j=1}^{\infty} \frac{P_j^2}{r_j^{2\alpha}} \text{var} \left[\sum_{i=1}^{N_j} X_i \right] \quad (6.73)$$

$$= \sum_{j=1}^{\infty} \frac{P_j^2}{r_j^{2\alpha}} \left[\lambda A \frac{1}{\mu^2} + \lambda A \frac{1}{\mu^2} \right] \quad (6.74)$$

$$= \frac{4\pi\lambda P_j^2}{\mu^2(2\alpha - 2)r_n^{2\alpha-2}} \quad (6.75)$$

□

If $\mu = 1$, the mean of the distribution does not change. However, the variance doubles compared to the expression in Thm. 31. This is an effect of choosing the Exponential distribution for the fading, and may not extend to other models.

6.2.2 Near vs far field

All of the random placement models have shown a convergence to Gaussian for small enough d (or equivalently, large enough λ). This part examines that convergence more closely.

Testing that the PDFs are indeed converging to a Gaussian requires a metric of measuring distance. The metric should remain the same even when the distributions are scaled by the same constant. At the same time, convergence in this metric should imply convergence in distribution. The metric of variational distance (L_1 norm of the difference between the PDFs) fits the bill⁵.




Figure 6.9: The variational distance between the interference distribution and the appropriate Gaussian without fading, $\alpha = 2.5$. Notice that there are lines of equal variational distance to a Gaussian.

Over the two-dimensional space of (r_n, λ) , which regions correspond to near field, and which correspond to far field? A thought experiment can help: choose an r_n and λ , and fix a realization of the location of Poisson distributed nodes. “Zoom-in” on this realization so that r_n grows and the interferers spread out, effectively decreasing the density of the interferers. While this can scale the interference random variable, intuitively, it should not affect the *shape* of the distribution.

To verify this intuition rigorously, zoom in such that r_n is scaled up by a factor c , so that the new no-talk radius $r'_n = cr_n$. The density of interferers scales *down* by $1/c^2$, *i.e.* the new density $\lambda' = \frac{\lambda}{c^2}$. The total interference, originally $I_T = \sum \frac{P_j}{r_j^\alpha}$ changes to $I'_T = \frac{1}{c^2} \sum \frac{P}{(cr)^\alpha} = \frac{1}{c^{2+\alpha}} I_T$. Because I_T and I'_T are related to each other by a constant, their variational distance from the Gaussians of the same means is equal! Therefore, the sets of points $(r'_n, \lambda') = (cr_n, \frac{\lambda}{c^2})$, for all c

⁵Given X and Y with PDFs $f_X(x)$, $f_Y(y)$, and some positive scaling factor c , the variational distance $d_{var}(\cdot, \cdot)$ between cX and cY for some constant c is given by:

$$\begin{aligned}
 d_{var}(cX, cY) &= \int_{\mathbb{R}} \left| \frac{1}{c} f_X(cx) - \frac{1}{c} f_Y(cx) \right| dx \\
 &\stackrel{(t=cx)}{=} \int_{\mathbb{R}} \left| \frac{1}{c} f_X(t) - \frac{1}{c} f_Y(t) \right| c dt \\
 &= \int_{\mathbb{R}} |f_X(t) - f_Y(t)| dt \\
 &= d_{var}(X, Y).
 \end{aligned}$$

Further, convergence in variational distance is stronger than convergence in distribution [148].

have the same variational distance from their corresponding Gaussian distribution. The locus of these points is the equation $r_n^2 \lambda = b$ for some constant b , and thus these curves can be used to define the contours of equal variational distance from Gaussian (see Fig. 6.9).




Figure 6.10: Variational distance by $r_n^2 \lambda$ with and without fading and with $\alpha = 2.5$. Notice that fading does not change the speed at which the distribution converges to Gaussian. This is because fading does not change *how many* nodes are contributing significantly to the interference.

Fig. 6.10 shows how quickly the distribution converges to Gaussian as $r_n^2 \lambda$ grows. By visual inspection, the variational distance (d_{var}) follows the relationship $(r_n^2 \lambda) d_{var}^{1/slope} = b$, where b is a constant (when plotted on a log-log scale, the sum of these two terms is a constant, so δ is the slope of Fig. 6.10).

Notice the lines for the dithered grid and the Poisson model with fading. A second thought experiment can help to understand the relationship between the different lines, based on the CLT. Given a set of N *iid* random variables, the CLT says that N will determine how closely the sum can be approximated by a Gaussian. For the dithered grid, consider just the first ring of nodes. The number of interferers in that ring is the area divided by d^2 or $N = \pi((r_n + d)^2 - r_n^2)/d^2 = \pi + 2\pi(r_n/d)$. The number of nodes that contribute to the CLT in this ring is determined by $m = r_n/d$. If the first ring is close to Gaussian, all subsequent rings will be closer to Gaussian (they have greater area, and therefore more nodes). Sums of Gaussians are Gaussian, so the total variational distance will depend only on m , and will decrease as m increases.

Applying the same logic to the Poisson case reveals a similar dependence on $r_n \sqrt{\lambda}$, or $r_n^2 \lambda$ as in the plot. When fading is added, *the fading does not affect the number of nodes contributing to the CLT*. Fading only changes the randomness of the individual terms. So, the distribution with fading converges to Gaussian at the same rate as without fading. The fading only changes the scaling factor.

Finally, consider Fig. 6.11, which plots the convergence to Gaussian for different path-loss exponents, α . α determines how similar in power the interference level from different nodes is, so it will affect the scaling. But, it does *not* affect the convergence, because α does not change how many more nodes are included in the CLT as $r_n^2 \lambda$ changes.




Figure 6.11: Variational distance to Gaussian for different values of α . As with fading, the path loss exponent does not change how many more nodes are contributing significantly to the interference as $r_n^2 \lambda$ gets large, so it does not change how quickly the distribution converges to Gaussian.

Conceptually, defining these “equivalence classes” in (r_n, λ) -space could be very useful:

- While designing policies, the policies for heavy-tailed distributions could be very different from light-tailed ones: a slight decrease in transmit power (or increase in distance) for light-tailed distributions can decrease the outage-probability significantly.
- To find the actual distribution of interference, one needs to work with just one value of (r_n, λ) for each fixed $r_n^2 \lambda$ and then scale appropriately. This will reduce the amount of computational overhead.
- In a far-field situation, when the Gaussian approximation is quite accurate, this approximation can be appreciably simpler than the actual distribution. This is because expressions for actual distributions are usually not even available in closed form, but only as Laplace-inverses of expressions that are in turn not expressed in a closed form (see e.g. (6.44) or expressions in [144, 149] etc.). For this reason, it would be nice to obtain rigorous bounds on the rate of this convergence to the Gaussian.

In the next section, this understanding of the distributions is used to explore different ways to share risk of bad placement realizations between primaries and secondaries.

6.3 Dealing with placement risk

This section uses the aggregate interference distributions from the last section to understand what should be done about clusters of secondaries close to the no-talk radius. The risk of these clusters causing too much interference is called “placement risk.”

In spirit, this section will help quantify Section 1.2 of the Introduction Chapter, which argues how performance and trust can improve as regulation decisions are moved closer to runtime. In this Section, as the TV whitespace database is able to use more of the information *it already has* about secondary device locations, it is able to better protect the primary and share the risk of

bad placement amongst secondary devices.

The Section will explore three strategies: setting a fixed no-talk radius at certification time, setting the per-node power at runtime, and setting the no-talk radius at runtime. These leverage database information about secondary locations and secondary node density to various degrees. They also put the burden of poor secondary placement on different primaries and secondaries in different ways. The examples are chosen to give an idea of the possible gains, but there are certainly many more approaches to using database location information. These will be left for future work.

6.3.1 r_n fixed at certification time

As shown in Section 1.2, the current approach to rules in the TV whitespaces do not allow the database to use *any* of the information it can gather from secondary requests. Without any runtime information, the only rule available is to set a fixed no-talk radius and fixed per-node power limit that should work for all realizations of secondary interference. This is obviously not possible, even if the secondaries are placed according to a regular grid, or somehow look like the interference is coming from a sea. Even fading and shadowing will force the fixed no-talk radius to be set according to a probabilistic guarantee,⁶ like:

$$P(I_T > I_{lim} | r_n = R) < \gamma \quad (6.76)$$

$$F_{I_T | r_n=R}(I_{lim}) < 1 - \gamma \quad (6.77)$$

where I_T is the total interference from the secondaries, I_{lim} is the allowed interference level at the primary receiver, r_n is the no-talk radius, F_{I_T} is the CDF of the total interference distribution, and $\gamma < 1$. This same rule will be used in this part to find the no-talk radius needed given random interference due to secondary placement.

This kind of rule distributes risk of something going wrong between both primaries and secondaries. The primaries will suffer whenever the fading or placement produces very high interference. But the secondaries closer to the primaries will suffer *all the time*. They are being denied transmission closer to the primary because of the fear of something going wrong.

Fig. 6.12 shows the effect of random placement on the needed no-talk radius to maintain this probabilistic guarantee on interference. To focus on placement, all of the examples in this Section will assume that there is no fading or shadowing.⁷ So, the distributions used for F_{I_T} in this plot are the Poisson distribution in Thm. 29, and a Gaussian with mean and variance given in Thm. 31. The problem is obvious: as the density of secondary devices increases, so does the needed no-talk radius.⁸

If this is the only option, the only way forward is to guess the largest expected secondary density, perhaps accounting for their self-interference, and setting the no-talk radius accordingly.

If, however, the database can use some of the information it can glean from the secondary transmission requests, it can provide better performance. The database certainly knows how many requests it is serving, and can figure out the secondary node density. Then, the power limit can

⁶The needed *margin*, or extension of the no-talk radius to account for fading and shadowing was explored in [140]. Another approach to finding the margin, which assumed nothing was known about the distribution of interference beyond the mean, was explored in [141].

⁷The results in [140] can be used to add the appropriate extra margin for these effects.

⁸The problem of high secondary node density is not unique to random placement. Using a deterministic model, this same problem was shown in [31].




Figure 6.12: The required r_n to ensure that the probability of total interference being too high is less than γ for different γ , and both the Poisson distribution and Gaussian approximation. This is for fixed per node power, and no-talk radius fixed at certification time. As nodes get more dense, the required no-talk radius must get larger. Therefore, with a fixed per-node power limit, and no control on the node density, there is no way to guarantee protection for primary users.

be a *power density limit* instead of a per-node limit.⁹ This requires the database to be able to set the secondary node power. This capability is already included in the OFCOM standards for TV whitespaces [151].

With a power-density limit, the needed no-talk radius is shown vs node density in Fig. 6.13. The effect of different γ 's is shown in Fig. 6.14. The no-talk radius with this rule seems like it must be set now for very low node densities. But this is slightly misleading – as the nodes become more sparse, a power-density rule will lead to extremely large per-node powers. The problem with low node densities is that when a node is next to the no-talk radius, its transmit power is so high that this node alone will overwhelm the primary.

Fig. 6.15 shows a more realistic example. The per-node power cannot reasonably go to infinity as the node density goes to zero. There must be some maximum power. With this maximum power, the no-talk radius can be set according to that required at the density where the maximum per-node power is reached.

With a power density rule and added maximum per-node power, it is possible to set a no-talk radius that will be sufficient to provide a probabilistic guarantee of protection to the primary receivers for *any* density of randomly placed secondary devices.

On the surface, this looks like a reasonable, albeit inefficient solution. The no-talk radius is being conservatively set for the rare events when clusters of secondaries are close to the no-talk radius. But the primary *is* being protected up to γ of the time. But what does this actually mean?

Placement uncertainty is a qualitatively very different than fading and shadowing uncertainty. Presumably, fading and shadowing change quickly as individual secondaries move, so a

⁹Making rules based on power density was suggested to the FCC, but ultimately was not accepted into the current rules [150]. [31, 32] show the gains from a power density rule, especially for rural locations, using a deterministic placement model.




Figure 6.13: If the no-talk radius and the maximum *power density* is fixed at certification time, the database can allow secondaries to have higher per node power at runtime. However, as the node density gets smaller, the per node allowed power goes up, which means the required no-talk radius must be larger for smaller node densities. Again, no guarantee can be offered for primary users.




Figure 6.14: The required no-talk radius from Fig. 6.13 for different levels of protection γ . As γ gets smaller, the margin over the no-talk radius required for the sea gets larger. The difference from the Gaussian distribution also gets larger, for the same node density.




Figure 6.15: The per-node power cannot be increased indefinitely. If there is a fixed maximum per-node power, as well as a maximum power density and no-talk radius fixed at certification time, this is the required no-talk radius by node density. In this case, a protection guarantee can be offered to the primary user: set the no-talk radius at the highest required point, where the maximum per-node power and the per-node power determined by the power density are equivalent.

primary will be able to average its performance over time to account for bad realizations.

On the other hand, how fast will placement change? Secondary towers will never move. Mobile devices will cluster around the towers, so if a tower is close to the no-talk radius, even the movement of the mobile devices will not change the poor placement. It may be a rare event for the tower to be poorly placed, but as soon as it is, the affected primary will always be harmed.

To account for placement risk appropriately, the database needs to use the information it has more effectively. It will certainly be capable of logging the information received from secondary nodes. It knows where these devices are placed. If it is able to change the node power, it can do this based on the actual realization of the placement. This is explored in the next part.

6.3.2 r_n set at certification time; power set at runtime

If the power can be set at runtime based on the actual location of secondary devices, a different rule is possible:

- Set the no-talk radius at certification time. For the example here, the no-talk radius is set to be sufficient for a sea with a desired average power density.
- Set the actual power density of nodes at runtime, depending on the realized placement, so that $I_T = I_{lim}$ when fading is ignored.¹⁰

The risk of poor placement is now spread amongst the participants in a very different way: the primary is *never* affected by poor placement. The secondaries share the burden of poor

¹⁰Fading can be added to this rule – it would turn into a slightly larger no-talk radius, and a probabilistic guarantee that the harm due to fading will not happen too often. Also, the rule here assumes the power will be the same for all nodes. A rule like that in [31] could be applied that changes the power based on radius to maximize efficiency. This is left to future work.

placement among all secondaries at all radial distances from the primary. Poor placement will result in a lower transmit rate for all secondary devices, instead of forcing outage at small radial distances.




Figure 6.16: If the per-node power can be set at runtime based on the realized node placement (information that the database has from secondary access requests), the CDF of the achieved power density is shown here. The sea would achieve a power density of 5. Sparse secondary nodes (low λ) would see the most gain from favorable placements allowing much higher power densities.

Fig. 6.16 shows the CDF of the power density that results from such a rule for different node densities. The plot is calculated using:

$$P(P_j < p | r_n = R) = P(I_T > I_{lim} | r_n = R, P_j = p) \quad (6.78)$$

$$= 1 - F_{I_T | r_n = R, P_j = p}(I_{lim}) \quad (6.79)$$

where the Poisson version of $F_{I_T | r_n = R, P_j = p}$ is given by Thm. 29, and the Gaussian has mean and variance given in Thm. 31. P_j is the per-node power. To plot, the power density $P = P_j \lambda$.

When the node density is very low, the power density has a much higher variance, because when a secondary node is located close to the fixed no-talk radius, the power must be much lower than when that first secondary is randomly placed far from the fixed no-talk radius. But this means that the secondaries can take advantage of much higher powers if they are all placed farther from the primary. They only have to lower their powers when their placement is worse.

The database is the gateway between the secondary devices and the rules that govern their behavior. If the database is trusted to make decisions, it can use this same location information to change the no-talk radius instead of (or along with) the power. This is done in the next part.

6.3.3 r_n set at runtime; power density set at certification time

If the database is allowed to set the no-talk radius at runtime, it can do so based on the actual realization of secondary placement. To isolate the effect of changing the no-talk radius, assume the rule is now the following:

- Power density is set at certification time
- No-talk radius is set at runtime according to the realization of secondary placement, such that $I_T = I_{lim}$ when fading is ignored.¹¹

Again, the risk of poor placement is divided among the secondaries – the primary will never experience too much interference because of secondary placement. However, the only secondaries that will be hurt by poor placement are the ones closest to the primary. So, instead of smooth degradation for all primaries as in the last part, there will be a probability of outage for each secondary.




Figure 6.17: If the power density is set at certification time, but the no-talk radius can be set at runtime based on the realization of the node placement, this is the CDF of the achieved no-talk radius. The sea would have a no-talk radius of 10. Sparse nodes see a higher no-talk radius because the power density rule allows these nodes to transmit at very high powers.

Fig. 6.17 shows the result of such a rule. It is the CDF of the probability that a secondary node at radius r from the primary will be able to transmit. It is found using the following:

$$P(r_n < r | \text{node at } r) = P(I_T < I_{lim} | r_n = r, \text{node at } r) \quad (6.80)$$

$$= P\left(I_{T, \text{no node at } r} + \frac{P_j}{r^\alpha} < I_{lim} | r_n = r\right) \quad (6.81)$$

$$= P\left(I_{T, \text{no node at } r} < I_{lim} - \frac{P_j}{r^\alpha} | r_n = r\right) \quad (6.82)$$

$$(6.83)$$

where P_j , the per-node power is $P_j = P/\lambda$.

As with the variable power density rule above, the more sparse the nodes are, the more they can take advantage of the closest node being far from the primary. But, sparse nodes will also have to give up more area when they happen to be clustered closer to the primary.

¹¹ Again, in this rule, fading can be added, and will show up as an extra margin on the required no-talk radius set at runtime. The power and no-talk radius could also be changed together to optimize secondary performance. This exploration is left to future work.

How do these different rules compare to each other? The next part gives a performance metric for evaluating the goodness of different rules and evaluating the value of runtime decisions.

6.3.4 Performance evaluation

To compare the performance of the different rules on a more quantitative basis, this part proposes and evaluates a performance metric. This concept is to understand the performance of a node at a radial distance r , including both power and transmission probability. To capture the power aspect, and to reduce the dependence on secondary transmission range, the metric will use a modified Gaussian capacity formula that assumes the transmit power is the received power, and the noise is 1. So, the “capacity” will be $\log(1 + P_j)$.

To capture the value of being able to transmit, the metric follows the insights from [90,91]: people tend to cluster around population centers, and the primary is assumed to be a TV tower. So, the metric will assume that transmitting closer to the primary is better. So, if the secondary node at r can transmit, it will get a value evaluation of r^{-w} where w controls how much higher of a value the secondary can get from being closer to the primary.

This will be integrated over the space of possible secondaries. So, the metric is:

$$\text{Performance} = \int_{\theta=0}^{2\pi} \int_{r=0}^{\infty} E_{P_j, P(TX)} [\text{capacity} \cdot \text{value} | \text{node at } r] r dr d\theta \quad (6.84)$$

$$= 2\pi \int_{r=0}^{\infty} E_{P_j, P(TX)} \left[\log(1 + P_j) \cdot r^{-w} \mathbb{1}_{\text{node transmitting} | \text{node at } r} \right] r dr \quad (6.85)$$

$$= 2\pi \int_{r=0}^{\infty} \int_{p=0}^{\infty} \log(1 + p) r^{1-w} P(\text{node transmitting}) f_{P_j}(p) dp dr \quad (6.86)$$

For the different versions of the rule, the metric can be simplified in different ways:

1. r_n and P_j fixed at certification time: to evaluate the metric, notice that any node placed at $r > r_n$ will have $P(\text{node transmitting}) = 1$. Also, $f_{P_j}(p)$ is a delta function at $p = P_j$. So, the metric can be evaluated as:

$$\text{Performance}_1 = 2\pi \int_{r_n}^{\infty} \log(1 + P_j) r^{1-w} dr \quad (6.87)$$

$$= 2\pi \log(1 + P_j) \frac{r_n^{2-w}}{w-2} \quad (6.88)$$

The r_n will be picked to be sufficient in Fig. 6.12 for $\lambda = 1$ and $\gamma = 0.01$.

2. r_n , the power density P , and the maximum node power $P_{j,max}$ fixed at certification time: again, any node at $r > r_n$ will always be able to transmit. The power will be fixed at the minimum of the power density rule or the maximum power rule. So, the evaluation here is similar to the above:

$$\text{Performance}_2 = 2\pi \log(1 + \min(P/\lambda, P_{j,max})) \frac{r_n^{2-w}}{w-2} \quad (6.89)$$

The r_n will be picked to be sufficient in Fig. 6.15 for all λ with $\gamma = 0.01$.

3. r_n set at certification time, power density set at runtime: to evaluate the metric, any node with $r > r_n$ will be able to transmit, but the power will follow the density given in Fig. 6.16:

$$\text{Performance}_3 = 2\pi \int_{r=r_n}^{\infty} \left(\int_{p=0}^{\infty} \log(1+p) f_{P_j}(p) dp \right) r^{1-w} dr \quad (6.90)$$

$$= 2\pi \frac{r_n^{2-w}}{w-2} \int_{p=0}^{\infty} \log(1+p) f_{P_j}(p) dp \quad (6.91)$$

4. r_n set at runtime, power density set at certification time: for this case, $f_{P_j}(p)$ will be a delta function at $p = P_j$, and the density of the no-talk radius will be that given in Fig. 6.17:

$$\text{Performance}_4 = 2\pi \log(1+P_j) \int_{r=0}^{\infty} P(\text{node at } r \text{ will transmit}) r^{1-w} dr \quad (6.92)$$

The performance metric is evaluated in Fig. 6.18 for different node densities. Regardless of node density, the rule with r_n and P_j set at certification time will have the same performance. When the node density gets large ($\lambda > 1$ in this case), the primary is no longer being protected up to $\gamma = 0.01$. This results in better performance than for the adaptive rules because the adaptive rules protect the primary despite the node density being high. The rule with r_n , power density, and maximum per node power set at certification time flattens for very small λ because the maximum per node power limit is reached.

In general, moving decisions closer to runtime does significantly better than forcing those decisions to be made conservatively at certification time. These are just introductory examples. Optimizing these rules to take advantage of the available information is left for future work.

6.4 Concluding remarks

This section will first include the conclusions and future work for the chapter thus far. The last part of this chapter will show an initial model for adapting spectrum jails to handle multiple secondaries.

6.4.1 What we have learned so far

This chapter gives a framework and simple examples for understanding how to move forward with the theory of spectrum jails with aggregate interference.

First, the nature of the interference must be understood. This chapter looked at the distribution from randomly placed nodes to understand what effect changing the no-talk radius will have on the interference. When the interferers are sparse from the perspective of the primary, the interference behaves as though it is coming from a single interferer. When the interferers are dense, the distribution behaves like a Gaussian. This suggests that urban and rural environments may need different no-talk radii, especially if the FCC chooses to allow a power-density rule instead of a per-node power rule.

Next, the desired rule must be chosen. This chapter showed examples of how the TV whitespace database could use the information it gets from the secondary devices to make regulatory decisions at runtime. Whether it is changing the no-talk radius or the power density based on the actual realization of nodes, the primary is better protected and the secondaries are better able to recover spectrum holes if the database can adapt decisions at runtime.




Figure 6.18: The performance as a function of the node density. Different plots have a different weighting, w , for the desire to transmit closer to the primary. The rules that allow modification based on the realization perform significantly better than those that have most parameters fixed at certification time. The rule with per-node power and no-talk radius fixed at certification time does well at high λ only because it is no longer offering sufficient protection to the primary. If it is important to have secondaries close to the primary, it is better to change the no-talk radius at runtime. If distance is less of an issue, it is better to change the power density at runtime. The optimal rule will be some combination of these two.

6.4.2 Future work

Because this chapter is really just a framework for future study into aggregate interference, there is a lot of work to be done in this space. Obviously finding better sharing rules and creating better models need to be done. So, this future work section will address only those bigger issues that have been ignored thus far in the chapter.

Bounds on the distribution for near field

The Gaussian distribution is a very nice approximation in far field, and qualitatively the single interferer model is a good approximation in near field. However, these have not yet been united to provide good and simple bounds for the actual distribution in all forms. Because the actual distribution is computationally difficult to work with, simple bounds will be very helpful for regulators trying to design good but simple rules.

Multiple database providers

When discussing the information the database has on secondary placement, it was implicitly assumed that there was only one database provider. Right now, in the U.S., there are several approved providers for the TV whitespaces. What is the right way to handle these providers?

If the database providers share all their secondary information, the problem reduces to the problem considered here. If however, the providers do not share information, there are several choices.

The interference limit I_{lim} could be split between the different providers, so that each only has to pay attention to their secondaries. Then, if one provider is already saturated with requests, the secondaries could shop around until they found one that gave them better transmission characteristics. There would be overhead on the part of the secondaries, but otherwise the problem would again reduce.

Alternatively, the databases could query each other to estimate the what information the others have. This is a much riskier proposition, as the primary protection would be dependent on the databases correctly estimating the others. But it would be interesting to know if such an approach could work.

Dealing with $r_p \neq 0$

Finally, throughout the chapter, it was assumed that the protected radius, $r_p = 0$. This was done partly for simplicity, but also because we believe this to be the worst case. If $r_p \neq 0$, then it will have full interference from the secondaries on one side of the no-talk radius, but the secondaries on the other side will be much farther away.

But how would one show that the no-talk radii found in this chapter are sufficient? How do the run-time decision gains change if $r_p \neq 0$? As a first note, the no-talk radius must be adapted to be $r'_n = r_n + r_p$, or the needed no-talk radius is actually the one found here plus the protected radius.

The new no-talk radius must be sufficient to protect the entire disk determined by r_p . One way to do this is to assign *sentinels* – points around the circumference of the protected disk that will be used to test total interference. The sentinels should be placed densely enough so that

the interference correlation between different sentinels is small but existent so that they capture the local area, but do not miss large sections of the disk.

Work needs to be done to prove that protecting these sentinels implies protection for the entire disk.

Now, how do the runtime decision results change? If r_p is very large, performance gains from runtime decisions may be lost. If r_p is very large, the sentinels will have essentially independent realizations of the placement of local interferers. If the runtime decision must work for all sentinels simultaneously, there may be enough independent realizations for the runtime decision to be equivalent to the certification time decision. This tradeoff between size of r_p and gains from runtime decisions must be studied.

If, however, the database can make local decisions on a sentinel-by-sentinel basis, the runtime gains are still possible. Work needs to be done to figure out how to effectively make such local decisions.

6.4.3 A first toy model for enforcement

Once the sharing rule for aggregate interference has been set, the next step is figuring out whether spectrum jails can enforce that rule. How should the spectrum jails model from chapters 2 to 5 be adapted to deal with multiple secondaries?

This part will present a very simple model for spectrum jails with multiple secondaries. The intent is to show that there is a way forward through queuing theory and to illustrate some of the problems that must be addressed. However, solving these issues will be left for future work.

The model

Queuing theory is the natural multiple-agent analog to the Markov chains used earlier. However, adapting the problem to use traditional queuing theory tools requires significant simplifying assumptions.




Figure 6.19: Placement model for the queuing example to extend spectrum jails. The primary receiver is in the center, and all secondaries are located a fixed distance away so that they all cause the same amount of interference.

First, assume that the secondaries are placed according to Fig. 6.19, so that all secondaries are the same fixed distance from the primary transmitter. This removes the problem of tracking individuals that cause different levels of interference.

Also assume that all secondaries are part of the same network, so that the decision to cheat will be made collectively instead of individually. This assumption completes the homogenization of the secondaries, so that they do not have different priorities and are not competing with each other. It can also, therefore, be assumed that they can coordinate with each other to choose who is allowed to turn on when.

With all secondaries in the same network, the sanction can also be a group sanction – it does not matter which secondary caused the interference to tip over the allowed value. Sending any of the participating secondaries to jail will hurt the network as a whole.

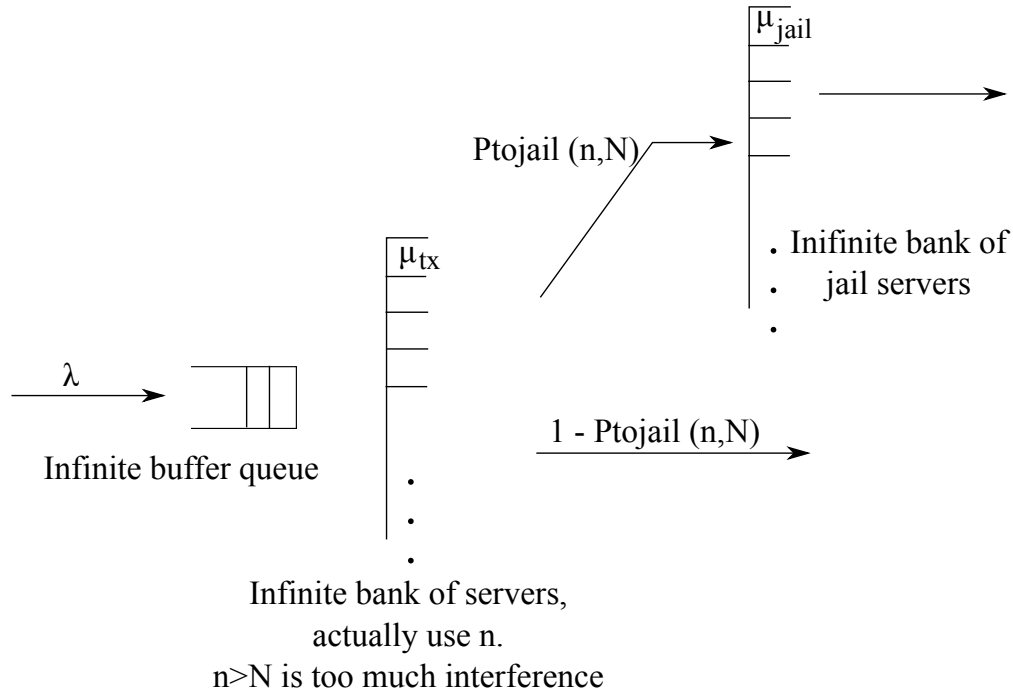


Figure 6.20: The toy model to extend spectrum jails using queues. Secondaries enter the first buffer when they have something to transmit. As a group, they decide how many interference servers they will use, where N is the tolerable amount of interference at the primary. When finished transmitting, a secondary will either leave the system or go to jail depending on how much total interference is being caused. The goal of the regulator is to choose the length of the jail sentence (or the speed of the jail server μ_{jail} to make it in the secondaries' best interest to not cause too much interference.

The model is shown in Fig. 6.20. The aggregate interference is modeled as a bank of servers at the primary receiver. When a secondary is transmitting, it is taking up one server slot, or one chunk of the allowed interference. The total allowed interference is equivalent to N simultaneously transmitting secondaries. Cheating is defined as more than N secondaries transmitting at any given time.

The secondaries as a group will choose how many of them are able to simultaneously

transmit, N_{opt} . If all N_{opt} servers are being used, any secondary that wants to transmit will sit in an infinite buffer queue. The secondaries have a Poisson arrival process governing when they want to transmit, with rate λ . The service time is exponential with rate μ_{tx} , which depends on the length of the secondary transmissions.

Let n be the number of servers that are actually being used at any given time. Depending on N_{opt} , when a secondary is finished transmitting, it may be sent to jail with probability P_{toJail} , defined as

$$P_{toJail} = P_{catch}P(n > N|N_{opt}) + P_{wrong}P(n \leq N|N_{opt}) \quad (6.93)$$

As in previous jail models, a secondary can be sent to jail rightfully if the secondaries as a group are causing too much interference, and wrongfully if they are not. This probability depends on whether the secondaries are willing to cheat ($N_{opt} > N$). The probability of going to jail when willing to cheat is:

$$P_{toJail,cheat} = P_{catch}P(n > N|N_{opt} > N) + P_{wrong}P(n \leq N|N_{opt} > N) \quad (6.94)$$

$$= P_{catch} \left(\sum_{k=N+1}^{N_{opt}} \pi_k \right) + P_{wrong} \left(\sum_{k=0}^N \pi_k \right) \quad (6.95)$$

where π_k is the probability of having k users in the queue or being served at any given time. For this queue, it is given by:

$$\pi_0 = \left(\sum_{k=0}^{N_{opt}-1} \frac{(N_{opt}\rho)^k}{k!} + \frac{(N_{opt}\rho)^{N_{opt}}}{N_{opt}!} \frac{1}{1-\rho} \right)^{-1} \quad (6.96)$$

$$\pi_k = \begin{cases} \pi_0 \frac{(N_{opt}\rho)^k}{k!}, & \text{if } 0 \leq k < N_{opt} \\ \pi_0 \frac{\rho^k N_{opt}^{N_{opt}}}{N_{opt}!}, & \text{if } k \geq N_{opt} \end{cases} \quad (6.97)$$

where

$$\rho = \frac{\lambda}{N_{opt}\mu_{tx}} \quad (6.98)$$

This is a standard result for an $M/M/N_{opt}$ queue, found in [152]. The probability of going to jail when honest is:

$$P_{toJail,honest} = P_{catch}P(n > N|N_{opt} \leq N) + P_{wrong}P(n \leq N|N_{opt} \leq N) \quad (6.99)$$

$$= P_{wrong} \quad (6.100)$$

Jail is modeled as an infinite bank of servers, with exponential service rate μ_{jail} which controls the length of sentence. It is an infinite bank of servers because any number of secondaries may sit in jail at a time, and the sentences should not depend on how many secondaries are in jail.

The secondary network wants to choose N_{opt} to minimize the expected time through the system, which is equivalent to the amount of time required for each transmission. So, the overall cost function is

$$C_S = C_{queue} + \frac{1}{\mu_{tx}} + C_{jail} \quad (6.101)$$

where

$$C_{jail} = E[\text{time in jail}] \quad (6.102)$$

$$\frac{1}{\mu_{jail}} \quad (6.103)$$

is the average time spent in jail, and

$$C_{queue} = E[\text{time in queue}] \quad (6.104)$$

$$= \left(\frac{\left(\frac{(N_{opt}\rho)_{opt}^N}{N_{opt}!} \right) \left(\frac{1}{1-\rho} \right)}{\sum_{k=0}^{N_{opt}-1} \frac{(N_{opt}\rho)^k}{k!} + \left(\frac{(N_{opt}\rho)_{opt}^N}{N_{opt}!} \right) \left(\frac{1}{1-\rho} \right)} \right) \quad (6.105)$$

is the average time spent waiting in the queue to transmit. The expression for the amount of time spent in the queue is a standard result for an $M/M/N_{opt}$ queue, found in [152].

The goal of the secondary is to choose N_{opt} such that:

$$N_{opt} = \underset{N_{opt}}{\operatorname{argmin}} C_S \quad (6.106)$$

6.4.4 Observations and the way forward

With this toy model of spectrum jail, the necessary sanction to guarantee $N_{opt} \leq N$ can easily be found. Define $N_{opt,cheat}$ as the optimal number of servers if cheating ($N_{opt,cheat} > N$), and $N_{opt,honest}$ as the optimal number of servers when honest ($N_{opt,honest} \leq N$). Then:

Lemma 12. *For the secondaries to choose $N_{opt} \leq N$, the sanction must be such that:*

$$C_{jail} > \frac{C_{queue}(N_{opt,honest})}{P_{toJail,cheat} - P_{wrong}} \quad (6.107)$$

Proof. The sanction must be such that the cost when honest (choosing $n \leq N$) is less than the cost when cheating (choosing $n > N$). Given an optimal $n = N_{opt,cheat} > N$ when cheating and an optimal $n = N_{opt,honest} \leq N$ when honest:

$$C_S(N_{opt,honest}) < C_S(N_{opt,cheat}) \quad (6.108)$$

$$C_{queue}(N_{opt,honest}) + \frac{1}{\mu_{tx}} + P_{toJail,honest} C_{jail} \leq C_{queue}(N_{opt,cheat}) + \frac{1}{\mu_{tx}} + P_{toJail,cheat} C_{jail} \quad (6.109)$$

$$C_{jail} > \frac{C_{queue}(N_{opt,honest}) - C_{queue}(N_{opt,cheat})}{P_{toJail,cheat} - P_{wrong}} \quad (6.110)$$

$$C_{jail} > \frac{C_{queue}(N_{opt,honest})}{P_{toJail,cheat} - P_{wrong}} \quad (6.111)$$

$$(6.112)$$

The last line is obviously true, but it also justified: the more secondaries that can turn on at a given time (higher N_{opt}), the smaller the queue wait time will be. So, if it is worthwhile to cheat, it is worthwhile to clear the buffer and have the wait time be 0. \square

With this Lemma, a number of observations can be made about issues that will need to be addressed in future work on spectrum jails for multiple secondaries:

- **Stability** – If N is such that $\lambda/(N\mu_{tx}) > 1$, the queue is unstable, and the expected wait time for secondaries goes to infinity. There is no way that jail can be used to keep secondaries honest in this case. This is potentially a big problem. If secondary networks have high market penetration, and the secondary demand is high, there will be incentive to overwhelm the primaries. It is possible that this could be solved with appropriate homebands that guarantee enough spectrum for all the secondaries that need it. But as in previous chapters, jail is not an appropriate rationing mechanism.
- **Adaptive jails** – The time waiting in the queue is a function of the arrival rate λ , and the service time μ_{tx} . Both of these are parameters of the secondary network. Even assuming that the system was stable, the honest wait times could be extremely long. The only way to cover all stable realizations of the secondary is to have an essentially infinite jail sentence. With multiple secondaries, adaptive jails may be required to provide trust without unacceptable overhead.

Extending this model to something more realistic will require innovative modeling in the following directions:

- **Different (potentially random) placement** – Truly capturing aggregate interference requires capturing different placements of secondaries. Within a spatial-queuing model for the aggregate interference, this means that secondaries will require different numbers of servers depending on their respective distances from the primary. Alternatively, the formulation could include continuous servers that represent a block of server resource. Each secondary would consume as large a portion of the available servers as it needed, with the service time being the same for any block of server resource. The author could not find any examples in the queuing literature for how to do this.
- **Individual decisions** – What happens if the secondaries are not actually part of a single network? If a secondary would be the only one using a server above the limit, would it have a different incentive to cheat than if the interference were already over the limit? How would one track the desires of individuals in different situations?
- **Group vs individual punishment** – The punishment assumed here was really for the group. If the interference is too high, the next secondary to finish transmitting would be sent to jail, regardless of whether it was the one to tip over the limit. Individual punishments would certainly be required if the secondaries are not part of the same network with a common goal.

Chapter 7

Jails vs fines

This thesis has been built on in-kind punishment. It is natural in the spectrum sharing context to think of spectrum jails because they can be implemented through databases using the identity system and kill-switch the TV whitespace databases will already need. It is also conceptually easier to develop sanctions in the same units as the behavior they are trying to prevent. For radios, transmission and jail use the same time and energy resources, so they can be combined to measure the total cost of sending a message.

But in-kind punishments are not the traditional sanction; fines are. This is because of a widely held belief in the law and economics¹ community that in-kind punishments and fines can create the same amount of dis-utility, but fines move resources around while in-kind punishments waste resources [153].² So, on one hand, if the appropriate conversion were made, fines should provide the same basic trust results that jails have been able to provide.

At the same time, the wasted resource argument does not quite apply. For humans, imprisonment wastes resources because the state now has to support the inmates with no productive output. Fines require significantly less infrastructure. For radios, *fines* have a more difficult implementation because they require a billing system, while jails can be implemented in the radio itself. While a radio is in jail, it does burn energy resources, but it also frees spectrum to be used by the primary or by a more honest secondary.

Beyond the question of resources, are fines and jails actually equivalent from a deterrence perspective? The law and economics literature can also be instructive here. Different kinds of people respond to different kinds of sanctions [112], so the type of sanction should be tailored to the individual to have the greatest effect. This is particularly true when fines can be treated by the rich like prices [154]. For example, parents may gladly pay a fine to pick up their kids from daycare late. They are effectively just paying for an extra hour of child care without caring about the label of a “fine” [116]. Other work suggests that social stigma may actually be more effective with companies who can easily pay the fine because it hurts something more important to the company: their reputation [115].

In spectrum regulation, primary users have the priority right to use the band. It is not an option in the current framework to use fines as a way of moving spectrum resources to higher valued uses. So, if fines are used, they must follow the ideas in [155] for applying a sanction when

¹This literature grew from the seminal papers [106–108]. See [109–111] for a general introduction to this area.

²There are many, like [114, 115] that counter this belief because in-kind punishments may have positive effects beyond anything that can be monetarily measured.

the action is definitely undesirable.

Another difference between fines and jails may appear when it is difficult to observe the person's level of wealth [112]. For spectrum regulation, this is a particularly important distinction. Unlike with humans, radios have a quality of service that they must provide. The benefits of cheating can be measured in terms of this quality of service. As argued in the rest of the thesis, it should be easy to bound the in-kind punishment required. Fines, however, would require knowing the monetary value of a transmission. How is this measured? By how much the user is willing to pay for the service? By how the company's market share changes with slight increases in reliability? The uncertainty in estimating the correct fine may be excessive.

So far, these differences are all qualitative. The purpose of this chapter is to show how this comparison can be made quantitative to really understand when in-kind punishments should be preferred over fines.

As with the rest of the thesis, the goal is to choose the type of sanction that most effectively addresses the problem of cheating. As mentioned in the Introductory Chapter, this idea is coming into vogue in the practice of enforcement through compliance theory [99]. There are many options to bring about compliance with rules within a population, and the goal is to put together a systematic way of choosing amongst the options.

With radios, the problem lends itself to systematic study. Because quality of service can be quantitatively measured as easily as monetary values, the two can be more directly compared. Moreover, the process of identifying and catching cheating radios will be a technical process with measurable performance specifications. So, the comparison can be done quantitatively in terms of *overhead*. Overhead is defined here as the cost to an honest user of the sanction system. Therefore, the overhead is directly tied to wrongful conviction.

Within the law and economics literature, wrongful conviction is almost taboo. It is obviously bad and obviously important [156, 157], but it is not be quantified or traded off the same way it can be in an engineered system.

This chapter uses overhead to understand the salient differences between fines and jails, including the effect of uncertainty and when a fine can be treated as a price. Through a general model and two toy examples, this chapter will show that for spectrum, in-kind punishments are the right choice.

7.1 The general problem

Fig. 7.1 gives an abstract idea of what must be considered when choosing how to apply a sanction. There are essentially three levels at which the sanction can be applied. Perfect deterrence refers to the actual valuation of the crime in the mind of the culprit. If it were possible to know all motivations for the crime, applying sanctions at those points of motivation would provide perfect deterrence with no overhead. However, this is not possible, so a sanction can only be applied through some kind of abstraction.

The in-kind abstraction applies a sanction based on the perceived utility of the crime. Spectrum jails to combat interference are a good example – the value of the crime is being able to transmit with a higher quality of service. So, by reducing the quality achieved by cheating through a jail sentence, it is possible to make the crime no longer worthwhile. But this involves estimating the cheating QOS, which has some uncertainty.

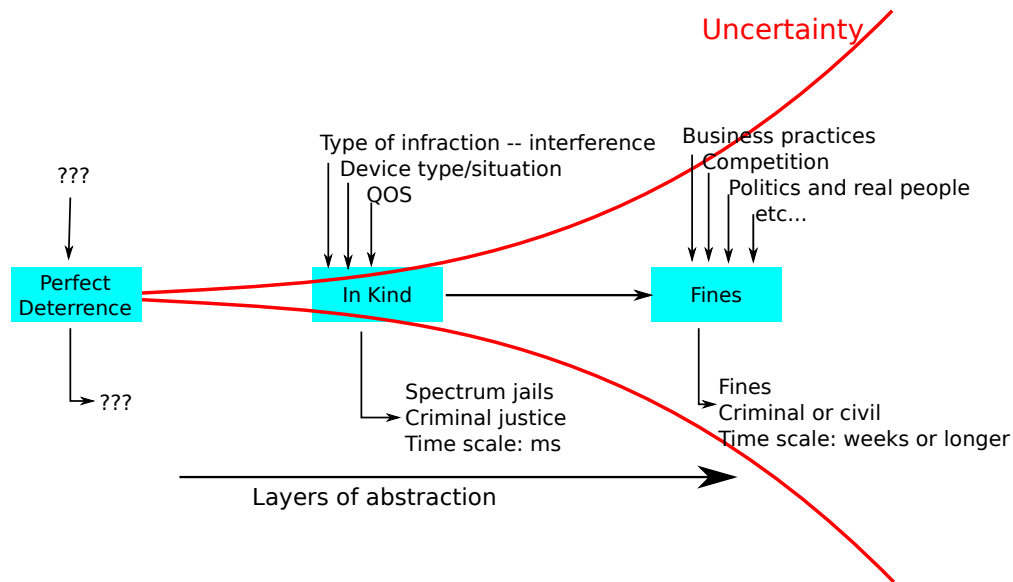


Figure 7.1: An abstract idea of the considerations of choosing in-kind vs monetary sanctions. There exists an actual valuation on the crime, but there is no way to estimate or attack that actual valuation. So, sanctions could be applied at different layers of abstraction. For spectrum sharing, the in-kind abstraction includes all technical inputs that determine the quality of service available, and the sanction looks like spectrum jails. The monetary abstraction layer includes all business and personal inputs that affect the monetary value of cheating. It also includes the available quality of service. With more inputs, the uncertainty increases, so while it may be possible to understand the range of in-kind valuations, it may not be possible to reliably quantify the range of monetary valuations.

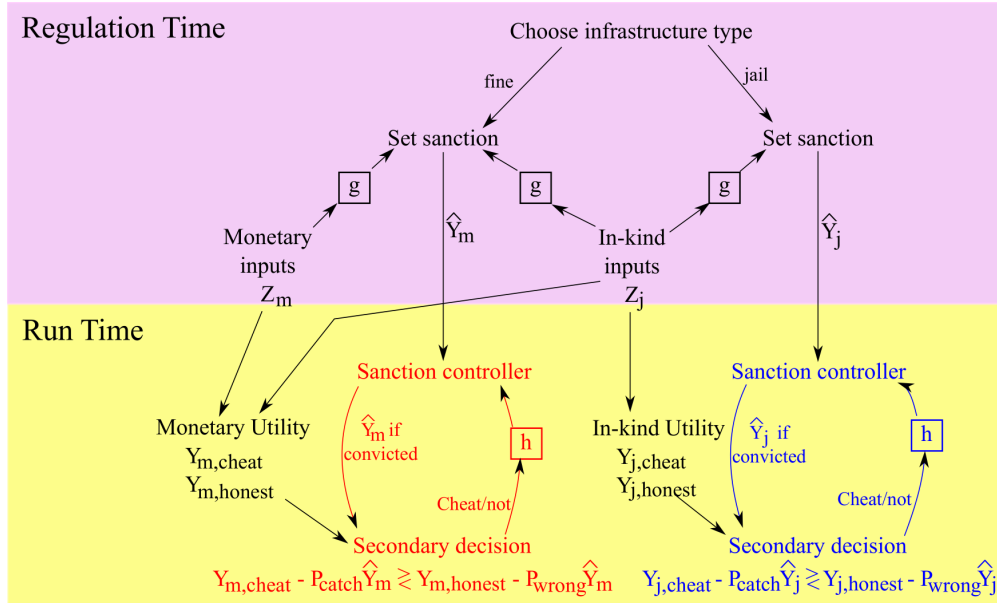


Figure 7.2: A model of how the type of sanction should be chosen. The sanction must be chosen at the regulation time scale with some noisy understanding of the monetary and in-kind inputs that are determining the utility function. These sanctions are then applied at runtime to any rightfully or wrongfully convicted agent. The agent decides whether to cheat based on the applied sanction, and its base utility while cheating or being honest.

The fines abstraction requires taking the value of the crime and abstracting it to monetary units. This includes all of the QOS valuation terms (and corresponding uncertainty) because these are determining the underlying value of the crime. But, converting to monetary units also involves understanding how the company translates the cheating QOS into a price for its customers or an edge over its competition. So, the monetary abstraction includes all of the uncertainty from the QOS *and* all the uncertainty related to business practices.

But this diagram comes from intuition, and does not allow a quantitative understanding of overhead and when fines or jails would be preferable.

Fig. 7.2 casts the problem as something closer to a quantitative model. It also includes elements of a flow-chart to understand when decisions need to be made. The decision of the type of sanction must be made at a regulation time scale, because jails and fines require significantly different infrastructure to implement. The diagram also assumes that the sanction must be set at regulation time, consistent with the rest of the thesis, so that trust can be guaranteed for all realizations of the actor at runtime. Setting the sanction is done with reference to monetary or in-kind inputs that are obscured through some function g . This is done because the regulator does not know the actual realization of the actor, so it cannot observe the actual inputs.

These decisions are made to allow the enforcement control systems to effectively deter cheating at runtime. The control systems for fines and jails are highlighted in Fig. 7.2. The sanction controller will apply the sanction if the actor is convicted, based on its observation of the decision to cheat, which may be noisy. So, conviction can happen wrongfully. The actor decides whether to cheat based on the sanction, and a utility function which takes in the actual monetary

and in-kind inputs.

The goal for the regulator is to choose the correct path (fines or jail) and the correct sanction so that the actor will choose not to cheat and the overhead is low. The overhead will be defined as the cost of wrongful convictions.

This general problem is closer to a quantitative model, but it is still difficult to analyze. So, the rest of this chapter includes two toy simplifications of this general model to show when jails should be preferred over fines.

7.2 A data-processing toy

This very simple first model will assign values to the elements of Fig. 7.2. Assume that the actor is choosing to cheat or do nothing. Also assume that the base valuation of cheating is v to which the monetary and in-kind inputs will be added:

$$Y_{m,honest,orig} = Y_{j,honest,orig} = 0 \quad (7.1)$$

$$Y_{m,cheat,orig} = v + z_m + z_j \quad (7.2)$$

$$Y_{j,cheat,orig} = v + z_j \quad (7.3)$$

where z_m and z_j are the realizations of the inputs for this particular actor. The regulator knows the distribution of the inputs, Z_m , Z_j , but not the realization. The regulator must choose the sanctions, \hat{Y}_m , \hat{Y}_j such that the realized actor has only a γ chance of choosing to cheat. After applying the sanction, the utility functions are:

$$Y_{m,honest} = -P_{wrong}\hat{Y}_m \quad (7.4)$$

$$Y_{m,cheat} = Y_{m,cheat,orig} - P_{catch}\hat{Y}_m \quad (7.5)$$

$$Y_{j,honest} = -P_{wrong}\hat{Y}_j \quad (7.6)$$

$$Y_{j,cheat} = Y_{j,cheat,orig} - P_{catch}\hat{Y}_j \quad (7.7)$$

The actor will try to maximize its utility, so it will choose to cheat if:

$$Y_{i,cheat,orig} > \rho\hat{Y}_i \quad (7.8)$$

where $i \in \{m, j\}$ is determined by the regulator's choice of whether to apply a fine or an in-kind sanction. Also, $\rho = P_{catch} - P_{wrong}$. So, the regulator must choose the sanction such that:

$$P(Y_{i,cheat,orig} > \rho\hat{Y}_i) < \gamma \quad (7.9)$$

It is always possible to find a sanction that will deter cheating $(1 - \gamma)$ of the time, so the choice between fines and jails depends on overhead. The overhead is coming only from wrongful convictions and is tied to the size of the sanction. Choosing the better option therefore requires comparing the two sanctions.

But this cannot be done directly – money and in-kind sanctions are fundamentally different units, even though this model obscures that fact. The fair comparison is in terms of monetary utility. The fine sanction is already in these units. To convert the in-kind sanction to monetary

units requires adding the realization of the monetary input. So, the in-kind sanction in monetary units is:

$$\hat{Y}_{j,converted} = \hat{Y}_j + Z_m \quad (7.10)$$

Both sanctions are guaranteed to work $1 - \gamma$ of the time. But they may not work the *same* $1 - \gamma$ of the time. Realizations with high monetary inputs and low in-kind inputs may cheat with a monetary sanction but not with an in-kind sanction. The regulator must decide which kind of failure is worse or whether it matters. Both kinds of failures will be considered equivalent here. So, the sanctions will be compared just based on average cost, $E[\hat{Y}_i]$ for $i \in \{m, j\}$.

For a concrete example, let $V = 0$, $Z_m \sim N(0, \sigma_m^2)$, $Z_j \sim N(0, \sigma_j^2)$.

Theorem 34. *Given the all-Gaussian example, the monetary abstraction always has a worse expected value of the sanction.*

Proof. The valuations have the following distributions:

$$Y_j \sim N(0, \sigma_j^2) \quad (7.11)$$

$$Y_m \sim N(0, \sigma_j^2 + \sigma_m^2) \quad (7.12)$$

And to satisfy the sanction requirement:

$$\hat{Y}_j = \sigma_j Q^{-1}(\gamma) / \rho \quad (7.13)$$

$$\hat{Y}_m = \sqrt{\sigma_j^2 + \sigma_m^2} Q^{-1}(\gamma) / \rho \quad (7.14)$$

$$E[\hat{Y}_j + Z_m] = \sigma_j Q^{-1}(\gamma) / \rho < E[\hat{Y}_2]. \quad \square$$

For the all-Gaussian example, it is preferable choose an in-kind sanction.³

7.3 QOS vs prices toy

While the first toy does capture the dependence between in-kind and monetary utility functions, it does not respect that the two can be fundamentally different. This second toy model captures some of those differences. It also more closely represents an abstract model of the spectrum sharing context.

Assume that there exists some service that is being provided on an ongoing basis. There is a base level of service available at a fixed price (this is the home band in previous chapters), but there exists an opportunity for better service if the service provider and user are willing to pay for it. The service provider can invest in a better service, either honestly or dishonestly. For example, this could be honestly or dishonestly investing in cognitive radio capability by choosing whether or not to actually look for a primary. The user will then have to pay a higher price for the better service.

There is some relationship between the quality of service (QOS) provided and the cost either to provide it (for the service provider) or the cost a user is willing to pay. For example, the relationship could be like that in Fig. 7.3. Assume that the provider will sell the service at the

³This result may change if the inputs are dependent, particularly if they are inversely correlated. This is left to future work.




Figure 7.3: A general model to understand operational differences between jails and fines. A provider can offer a service of a particular quality for a particular price based on whether or not it cheats. The user will buy the service as long as it is below the user's QOS-price tradoff. The sanction will either reduce the QOS available or raise the price at which that service is offered in order to make cheating not worthwhile.




Figure 7.4: A simple version of the general model in Fig. 7.3 in which all relationships are assumed to be linear.

user's price.⁴ So, it has a profit margin of the difference between what the user is willing to pay, and the cost to provide the service. The provider will provide a QOS at the lowest possible cost to maximize its profit margin.

When the sanctions are applied, they only operate in one direction. Any monetary sanction will raise the price of the service provider's lines for a given QOS. So, it operates in the vertical direction. If an in-kind sanction is applied, it does not affect the price, but it does lower the QOS available at that price. So, the in-kind sanction operates in the horizontal direction.

This general model is complicated, so the simplified linear model in Fig. 7.4 will be used instead.

7.3.1 The model

There exists a base level of quality providable at a fixed price. The user is always willing to pay this price. For simplicity, this point is normalized to $(0,0)$.

⁴This assumption is only true in a monopoly situation, but it captures the desired effect: the provider will pay the lowest cost it can to provide the service.

The price the user is willing to pay is defined by a function of the QOS, q :

$$P_{user}(q) = p_{user}q \quad (7.15)$$

The cost to provide a QOS while honest or cheating have the same form:

$$P_{honest}(q) = p_{honest}q \quad (7.16)$$

$$P_{cheat}(q) = p_{cheat}q \quad (7.17)$$

Assume that $p_{user} \geq p_{honest} \geq p_{cheat}$. In other words, in the absence of a sanction, it is cheaper to cheat than operate honestly, and the user is willing to pay either of these costs. The provider will provide every possible QOS at the lowest possible price, assuming the user is still willing purchase it.

The possible QOS is not unbounded. Let q_h and q_c be the maximum possible QOS providable when honest and cheating respectively. So, the honest price is defined on $q \in [0, q_h]$ and the cheating price is defined on $q \in [0, q_c]$. The user's price is defined on the entire range: $q \in [0, \infty)$.

What do the cheating and honest lines really mean? Think of the example of spectrum sharing. The honest line can be interpreted as a range of possible technologies. For example, better sensing technology will be better able to recover spectrum holes, so it will produce a higher QOS. But, it will also be more expensive. Think of the cheating line as using more band-expansion to provide higher QOS. The technology is more expensive to accommodate a larger bandwidth. But, the cheating also causes harm to more primaries. Therefore, when a sanction is applied, higher QOS implies that the secondary will get caught more often. For the same jail sentence or fine, higher QOS will have a higher effective sanction. A linear relationship will be assumed.

Define the fine sanction as p_{fine} and the in-kind sanction as q_{jail} . The cheating price reflects the full force of the sanction, and the honest price reflects a fraction δ of the sanction to model wrongful conviction. The user's price does not change, because it is a limit on what the user will pay, not a price at which the QOS is able to be provided. Because the in-kind sanction alone affects the QOS, the in-kind sanction will affect the maximum possible QOS, while the monetary sanction will not.

With the monetary sanction, the price functions become:

$$P_{user,fine}(q) = p_{user}q, q \in [0, \infty) \quad (7.18)$$

$$P_{honest,fine}(q) = p_{honest}q + \delta p_{fine}q, q \in [0, q_h] \quad (7.19)$$

$$P_{cheat,fine}(q) = p_{cheat}q + p_{fine}q, q \in [0, q_c] \quad (7.20)$$

With the in-kind sanction, the sanction produces a smaller QOS for the same price. For a given price, p , aiming at a given QOS q , the resulting quality q' with sanction is:

$$q' = q - q_{jail}q = q(1 - q_{jail}). \quad (7.21)$$

This gives the new maximum value of the quality. To normalize the slope of the line correctly, the resulting line has a new effective price, p' which is:

$$p'q' = pq \quad (7.22)$$

$$p'q(1 - q_{jail}) = pq \quad (7.23)$$

$$p' = p/(1 - q_{jail}) \quad (7.24)$$

The new price functions are therefore:

$$P_{user,jail}(q) = p_{user}q, q \in [0, \infty) \quad (7.25)$$

$$P_{honest,jail}(q) = p_{honest} \left(\frac{q}{1 - \delta q_{jail}} \right), q \in [0, q_h(1 - \delta q_{jail})] \quad (7.26)$$

$$P_{cheat,jail}(q) = p_{cheat} \left(\frac{q}{1 - q_{jail}} \right), q \in [0, q_c(1 - q_{jail})] \quad (7.27)$$

Note that the fine price can be anything, $p_{fine} \in [0, \infty)$, but the in-kind sanction only makes sense if limited, $q_{jail} \in [0, 1]$. This is because there is no reason to make the sanction worse than choosing to operate only at the base level.

Deciding whether jails or fines are the better sanction will depend on which is effective at deterring cheating and which has a smaller overhead. Overhead is defined as the extra utility an honest user loses because of the sanction. Formally, define the overhead as

$$O_{fine} = \frac{P_{honest,fine} - P_{honest}}{P_{honest}} \quad (7.28)$$

$$= \delta \frac{p_{fine}}{p_{honest}} \quad (7.29)$$

$$O_{jail} = \frac{P_{honest,jail} - P_{honest}}{P_{honest}} \quad (7.30)$$

$$= \frac{\delta q_{jail}}{1 - \delta q_{jail}} \quad (7.31)$$

7.3.2 Without uncertainty

This part will first find the required sanction to deter cheating and then examine the overhead. The problem falls into three cases, depending on the relationship between the honest and cheating lines and their endpoints:

Case 1: $q_c \leq q_h$ This case is shown in Fig. 7.5, with red lines showing the required fine and in-kind sanction. Both must move the cheating line to be greater than the honest line. So the required sanction for the fine case is:

$$P_{honest,fine}(q) < P_{cheat,fine} \quad (7.32)$$

$$p_{honest}q + \delta p_{fine}q < p_{cheat}q + p_{fine}q \quad (7.33)$$

$$p_{fine} > \frac{p_{honest} - p_{cheat}}{1 - \delta} \quad (7.34)$$

For the in-kind case, we have:

$$P_{honest,jail}(q) < P_{cheat,jail} \quad (7.35)$$

$$p_{honest} \left(\frac{q}{1 - \delta q_{jail}} \right) < p_{cheat} \left(\frac{q}{1 - q_{jail}} \right) \quad (7.36)$$

$$q_{jail} > \frac{p_{honest} - p_{cheat}}{p_{honest} - \delta p_{cheat}} \quad (7.37)$$




Figure 7.5: **Case 1:** This shows the options for sanctions given the depicted relationship between lines. The in-kind sanction can push the cheating line back to be worse than the honest line in QOS for the same price. Alternatively, the fine sanction can push the cheating line up so that any given QOS costs more to provide through cheating than honestly.

Plugging into the overhead definition, they are found to be exactly the same:

$$O_{fine} = O_{jail} = \left(\frac{\delta}{1 - \delta} \right) \left(\frac{p_{honest} - p_{cheat}}{p_{honest}} \right) \quad (7.38)$$

This makes intuitive sense because both sanctions are moving the cheating (and honest) QOS-price tradeoff to the same resulting line.

Case 2: $q_c \geq q_h$ and $P_{honest}(q_h) \leq P_{cheat}(q_c)$

This case is shown in Figs. 7.6 and 7.7. Again, the red lines show the required movement of the sanction. Because it cannot do anything to change the QOS, the only option for the fine is to move the cheating line up to the user line. The in-kind sanction, however, *can* change the maximum cheating QOS, and therefore, it needs to move the cheating line to the minimum of two places. Either the in-kind sanction can make the highest cheating QOS the same as the honest case (which is sufficient because $P_{honest}(q_h) < P_{cheat}(q_c)$), or it can move the line to be slightly higher than the user line.

The required fine sanction is:

$$P_{user, fine}(q) < P_{cheat, fine}(q) \quad (7.39)$$

$$p_{user}q < p_{cheat}q + p_{fine}q \quad (7.40)$$

$$p_{fine} > p_{user} - p_{cheat} \quad (7.41)$$

which results in overhead

$$O_{fine} = \delta \frac{p_{user} - p_{cheat}}{p_{honest}} \quad (7.42)$$

And the required in-kind sanction is either moving the maximum QOS to the honest case:

$$q_c(1 - q_{jail}) < q_h(1 - \delta q_{jail}) \quad (7.43)$$

$$q_{jail} > \frac{q_c - q_h}{q_c - \delta q_h} \quad (7.44)$$




Figure 7.6: **Case 2, version 1:** The options for the sanction given this configuration either uses jails to push the cheating line so that it cannot provide a better QOS than the honest line. Or, fines can make the price of that high level of service so large that the user is unwilling to pay it.




Figure 7.7: **Case 2, version 2:** Given this configuration of lines, the only choice is to use either sanction to push the cheating line to be so expensive for a given QOS that the user is unwilling to pay it.




Figure 7.8: **Case 3:** Given this configuration, the in-kind sanction can push the cheating line back to the honest line, or the fine sanction can make cheating too expensive for the user to willing pay.

or, it can move the cheating line to the user line:

$$P_{user,jail}(q) < P_{cheat,jail}(q) \quad (7.45)$$

$$p_{user}q < p_{cheat} \left(\frac{q}{1 - q_{jail}} \right) \quad (7.46)$$

$$q_{jail} > \frac{p_{user} - p_{cheat}}{p_{user}} \quad (7.47)$$

In general, the smaller of these two is needed:

$$q_{jail} > \min \left(\frac{q_c - q_h}{q_c - \delta q_h}, \frac{p_{user} - p_{cheat}}{p_{user}} \right) \quad (7.48)$$

which results in overhead

$$O_{jail} = \min \left(\delta \frac{q_c - q_h}{q_c - \delta q_h}, \delta \frac{p_{user} - p_{cheat}}{p_{user} - \delta(p_{user} - p_{cheat})} \right) \quad (7.49)$$

The particular parameters will determine which of these overheads is worse. The better overhead will be discussed after all cases are complete.

Case 3: $q_c \geq q_h$ and $P_{honest}(q_h) > P_{cheat}(q_c)$

This case is shown in Fig. 7.8. For this case, the sanction requirements are clear: the in-kind sanction just needs to move the cheating line to the honest line, while the monetary sanction must move the cheating line up to the user line.

The required sanctions and overheads are repeated here, but are the same as parts of the previous cases. For the fines case:

$$P_{user,fine}(q) < P_{cheat,fine}(q) \quad (7.50)$$

$$p_{user}q < p_{cheat}q + p_{fine}q \quad (7.51)$$

$$p_{fine} > p_{user} - p_{cheat} \quad (7.52)$$

which results in overhead

$$O_{fine} = \delta \frac{p_{user} - p_{cheat}}{p_{honest}} \quad (7.53)$$




Figure 7.9: A guide to the parameter regions when jails or fines have better overhead. The line represents the points where the overheads are equivalent. To the left of the line, jails have better overhead. To the right of the line, fines have better overhead. δ is capturing the wrongful conviction, with higher δ implying a higher wrongful conviction rate.

The in-kind case requires:

$$P_{\text{honest}, \text{jail}}(q) < P_{\text{cheat}, \text{jail}} \quad (7.54)$$

$$p_{\text{honest}} \left(\frac{q}{1 - \delta q_{\text{jail}}} \right) < p_{\text{cheat}} \left(\frac{q}{1 - q_{\text{jail}}} \right) \quad (7.55)$$

$$q_{\text{jail}} > \frac{p_{\text{honest}} - p_{\text{cheat}}}{p_{\text{honest}} - \delta p_{\text{cheat}}} \quad (7.56)$$

which results in overhead

$$O_{\text{jail}} = \left(\frac{\delta}{1 - \delta} \right) \left(\frac{p_{\text{honest}} - p_{\text{cheat}}}{p_{\text{honest}}} \right) \quad (7.57)$$

When are fines or in-kind punishments better? Fig. 7.10 gives the answer, with Fig. 7.9 providing a guide. Each line represents the points for which the overhead is the same for jails and fines, given a δ . To the left of the line, in-kind sanctions produce better overheads. To the right of the line, fines have smaller overheads. Notice that as δ gets smaller, a larger range of parameters have in-kind sanctions as the better option. Also note that this plot only shows the $q_c > q_h$ case. If this is not true, the sanctions have equivalent overheads.

It was assumed that these lines were known. But in reality, there will be some uncertainty. This will be covered in the next part.

7.3.3 With uncertainty

Uncertainty is modeled by allowing the parameters to be completely unknown within specified ranges. These ranges can be interpreted as mapping out $1 - \gamma$ of the support of the actual




Figure 7.10: Regions where jails or fines have better overhead, for different δ . Higher δ implies higher wrongful conviction. The regions are as in Fig. 7.9. As δ gets smaller, jails are the better choice for a wider range of parameters.

distribution of the parameter. So the sanctions will fail in only γ of the possible realizations. The parameters and ranges of interest are:

$$p_{\text{user}} \in [p_{\text{user},\min}, p_{\text{user},\max}] \quad (7.58)$$

$$p_{\text{honest}} \in [p_{\text{honest},\min}, p_{\text{honest},\max}] \quad (7.59)$$

$$p_{\text{cheat}} \in [p_{\text{cheat},\min}, p_{\text{cheat},\max}] \quad (7.60)$$

$$q_h \in [q_{h,\min}, q_{h,\max}] \quad (7.61)$$

$$q_c \in [q_{c,\min}, q_{c,\max}] \quad (7.62)$$

The effect of uncertainty will be explored in three stages: only the prices are unknown, only the maximum QOS parameters are unknown, and then what happens when they are all unknown.

Only the QOS is known

Let the uncertainty be restricted to

$$p_{\text{user}} \in [p_{\text{user},\min}, p_{\text{user},\max}] \quad (7.63)$$

$$p_{\text{honest}} \in [p_{\text{honest},\min}, p_{\text{honest},\max}] \quad (7.64)$$

$$p_{\text{cheat}} \in [p_{\text{cheat},\min}, p_{\text{cheat},\max}] \quad (7.65)$$

with q_h and q_c fixed.

Because the cases without uncertainty were determined by q_h and q_c , the with uncertainty problem has the same cases. Furthermore, the goal now is to guarantee that the whole range of cheating prices is worse than anything in the range of honest or user prices depending on the case. So, in general, the required sanctions are the same as they were without uncertainty with the

following substitution:

$$\begin{aligned} p_{user} &\rightarrow p_{user,max} \\ p_{honest} &\rightarrow p_{honest,max} \\ p_{cheat} &\rightarrow p_{cheat,min} \end{aligned} \tag{7.66}$$

The overheads are assessed on the actual realization of the prices with the required sanctions. The solutions and overhead are therefore:

Case 1: $q_c \leq q_h$

$$p_{fine} > \frac{p_{honest,max} - p_{cheat,min}}{1 - \delta} \tag{7.67}$$

$$q_{jail} > \frac{p_{honest,max} - p_{cheat,min}}{p_{honest,max} - \delta p_{cheat,min}} \tag{7.68}$$

with overhead

$$O_{fine} = \left(\frac{\delta}{1 - \delta} \right) \left(\frac{p_{honest,max} - p_{cheat,min}}{p_{honest}} \right) \tag{7.69}$$

$$O_{jail} = \left(\frac{\delta}{1 - \delta} \right) \left(\frac{p_{honest,max} - p_{cheat,min}}{p_{honest,max}} \right) \tag{7.70}$$

The in-kind sanction has the better overhead.

Case 2: $q_c \geq q_h$ and $p_{honest,max} q_h < p_{cheat,min} q_c$

$$p_{fine} > p_{user,max} - p_{cheat,min} \tag{7.71}$$

$$q_{jail} > \min \left(\frac{q_c - q_h}{q_c - \delta q_h}, \frac{p_{user,max} - p_{cheat,min}}{p_{user,max}} \right) \tag{7.72}$$

with overhead

$$O_{fine} = \delta \frac{p_{user,max} - p_{cheat,min}}{p_{honest}} \tag{7.73}$$

$$O_{jail} = \min \left(\delta \frac{q_c - q_h}{q_c - \delta q_h}, \delta \frac{p_{user,max} - p_{cheat,min}}{p_{user,max} - \delta(p_{user,max} - p_{cheat,min})} \right) \tag{7.74}$$

Case 3: $q_c \geq q_h$ and $p_{honest,max} q_h > p_{cheat,min} q_c$

$$p_{fine} > p_{user,max} - p_{cheat,min} \tag{7.75}$$

$$q_{jail} > \frac{p_{honest,max} - p_{cheat,min}}{p_{honest,max} - \delta p_{cheat,min}} \tag{7.76}$$

with overhead

$$O_{fine} = \delta \frac{p_{user,max} - p_{cheat,min}}{p_{honest}} \tag{7.77}$$

$$O_{jail} = \left(\frac{\delta}{1 - \delta} \right) \left(\frac{p_{honest,max} - p_{cheat,min}}{p_{honest,max}} \right) \tag{7.78}$$

The range of parameter values that produce better overheads for jails or fines are also shown in Fig. 7.10, with the appropriate parameter substitution.

More interesting is what happens in the limit of increasing uncertainty.

Let $\{p_{user,min}, p_{honest,min}, p_{cheat,min}\} \rightarrow 0$ and $\{p_{user,max}, p_{honest,max}, p_{cheat,max}\} \rightarrow \infty$.

Theorem 35. *For the fine sanction, regardless of q_c, q_h , $p_{fine} \rightarrow \infty$ and therefore $O_{fine} \rightarrow \infty$*

Proof. Making the required substitutions in Eq. (7.67) for the fines cases without uncertainty, and letting $\{p_{user,max}, p_{honest,max}, p_{cheat,max}\} \rightarrow \infty$ gives the result. \square

On the other hand, the in-kind sanction can still handle the situation:

Theorem 36. *With the in-kind sanction, $\{p_{user,min}, p_{honest,min}, p_{cheat,min}\} \rightarrow 0$, and $\{p_{user,max}, p_{honest,max}, p_{cheat,max}\} \rightarrow \infty$, the required sanction to deter cheating is*

$$q_{jail} = 1 \quad (7.79)$$

which produces overhead

$$O_{jail} = \frac{\delta}{1 - \delta} \quad (7.80)$$

Proof. In order to deter cheating, the sanction must make cheating a worse option than operating at the base level, which is at (0,0). This is done by making the maximum achievable QOS with cheating equal to zero:

$$q_c(1 - q_{jail}) = 0 \quad (7.81)$$

$$q_{jail} = 1 \quad (7.82)$$

Plugging this into the definition of overhead gives the second part of the result. \square

If we have no information about the prices, and fines can only affect prices, then fines cannot deter cheating. However, because we know the maximum QOS and in-kind sanctions operate by changing QOS, the in-kind sanction is still effective.

This is similar to the situation in Chapters 2 to 5 of this thesis: bounds on the technical, QOS cost function of the secondaries is known. However, little is known about how this QOS translates into a price. So, operating in fines is not possible, but the sanction can always make the QOS worse with cheating than it is in the home band.

Only the prices are known

Let the uncertainty be restricted to

$$q_h \in [q_{h,min}, q_{h,max}] \quad (7.83)$$

$$q_c \in [q_{c,min}, q_{c,max}] \quad (7.84)$$

with all of the price parameters known. With no uncertainty, the maximum QOS parameters determined which case the solution came from, and a combination of maximum QOS and price parameters determined the solution. With uncertainty in the maximum QOS, the solutions can

still be split into similar cases, with similar solutions. But like the last uncertainty set, one must substitute the worst case of the QOS uncertainty set.

So, with uncertainty on the maximum QOS, the solution is to make the following substitutions:

$$q_h \rightarrow q_{h,min} \quad (7.85)$$

$$q_c \rightarrow q_{c,max} \quad (7.86)$$

and then use the deterministic solutions, with overhead assessed on the actual realizations:

Case 1: $q_{c,max} \leq q_{h,min}$

$$p_{fine} > \frac{p_{honest} - p_{cheat}}{1 - \delta} \quad (7.87)$$

$$q_{jail} > \frac{p_{honest} - p_{cheat}}{p_{honest} - \delta p_{cheat}} \quad (7.88)$$

with overhead

$$O_{fine} = O_{jail} = \left(\frac{\delta}{1 - \delta} \right) \left(\frac{p_{honest} - p_{cheat}}{p_{honest}} \right) \quad (7.89)$$

Case 2: $q_{c,max} \geq q_{h,min}$ and $p_{honest}q_{h,min} < p_{cheat}q_{c,max}$

$$p_{fine} > p_{user} - p_{cheat} \quad (7.90)$$

$$q_{jail} > \min \left(\frac{q_{c,max} - q_{h,min}}{q_{c,max} - \delta q_{h,min}}, \frac{p_{user} - p_{cheat}}{p_{user}} \right) \quad (7.91)$$

with overhead

$$O_{fine} = \delta \frac{p_{user} - p_{cheat}}{p_{honest}} \quad (7.92)$$

$$O_{jail} = \min \left(\delta \frac{q_{c,max} - q_{h,min}}{q_{c,max} - \delta q_{h,min}}, \delta \frac{p_{user} - p_{cheat}}{p_{user} - \delta(p_{user} - p_{cheat})} \right) \quad (7.93)$$

Case 3: $q_{c,max} \geq q_{h,min}$ and $p_{honest}q_{h,min} > p_{cheat}q_{c,max}$

$$p_{fine} > p_{user} - p_{cheat} \quad (7.94)$$

$$q_{jail} > \frac{p_{honest} - p_{cheat}}{p_{honest} - \delta p_{cheat}} \quad (7.95)$$

with overhead

$$O_{fine} = \delta \frac{p_{user} - p_{cheat}}{p_{honest}} \quad (7.96)$$

$$O_{jail} = \left(\frac{\delta}{1 - \delta} \right) \left(\frac{p_{honest} - p_{cheat}}{p_{honest}} \right) \quad (7.97)$$

Again, it is of interest to see what happens when there is no information about the maximum QOS. Let $q_{h,min}, q_{c,min} \rightarrow 0$ and $q_{h,max}, q_{c,max} \rightarrow \infty$. With the substitutions and taking the limit on uncertainty, the problem falls into only Case 2. So, the solution is to use fines and in-kind sanctions to move the cheating line to worse than the user line:

$$p_{fine} > p_{user} - p_{cheat} \quad (7.98)$$

$$q_{jail} > \frac{p_{user} - p_{cheat}}{p_{user}} \quad (7.99)$$

which results in overhead

$$O_{fine} = \delta \frac{p_{user} - p_{cheat}}{p_{honest}} \quad (7.100)$$

$$O_{jail} = \delta \frac{p_{user} - p_{cheat}}{p_{user} - \delta(p_{user} - p_{cheat})} \quad (7.101)$$

where the worse overhead depends on the values of the parameters. The regions where fines are the better option are again in Fig. 7.10 where the parameters are their maximum (or minimum) values according to the substitution in Eq. (7.83).

Uncertainty on all parameters

Let the uncertainty set be:

$$p_{user} \in [p_{user,min}, p_{user,max}] \quad (7.102)$$

$$p_{honest} \in [p_{honest,min}, p_{honest,max}] \quad (7.103)$$

$$p_{cheat} \in [p_{cheat,min}, p_{cheat,max}] \quad (7.104)$$

$$q_h \in [q_{h,min}, q_{h,max}] \quad (7.105)$$

$$q_c \in [q_{c,min}, q_{c,max}] \quad (7.106)$$

As a combination of the above two situations, the solution is obviously again substitution of all parameters with their worst case possibility:

$$p_{user} \rightarrow p_{user,max} \quad (7.107)$$

$$p_{honest} \rightarrow p_{honest,max}$$

$$p_{cheat} \rightarrow p_{cheat,min}$$

$$q_h \rightarrow q_{h,min}$$

$$q_c \rightarrow q_{c,max}$$

and then using the solution to the deterministic problem, with overheads assessed on the actual realizations:

Case 1: $q_{c,max} \leq q_{h,min}$

$$p_{fine} > \frac{p_{honest,max} - p_{cheat,min}}{1 - \delta} \quad (7.108)$$

$$q_{jail} > \frac{p_{honest,max} - p_{cheat,min}}{p_{honest,max} - \delta p_{cheat,min}} \quad (7.109)$$

with overhead

$$O_{fine} = \left(\frac{\delta}{1 - \delta} \right) \left(\frac{p_{honest,max} - p_{cheat,min}}{p_{honest}} \right) \quad (7.110)$$

$$O_{jail} = \left(\frac{\delta}{1 - \delta} \right) \left(\frac{p_{honest,max} - p_{cheat,min}}{p_{honest,max}} \right) \quad (7.111)$$

Case 2: $q_{c,max} \geq q_{h,min}$ and $p_{honest,max}q_{h,min} < p_{cheat,min}q_{c,max}$

$$p_{fine} > p_{user,max} - p_{cheat,min} \quad (7.112)$$

$$q_{jail} > \min \left(\frac{q_{c,max} - q_{h,min}}{q_{c,max} - \delta q_{h,min}}, \frac{p_{user,max} - p_{cheat,min}}{p_{user,max}} \right) \quad (7.113)$$

with overhead

$$O_{fine} = \delta \frac{p_{user,max} - p_{cheat,min}}{p_{honest}} \quad (7.114)$$

$$O_{jail} = \min \left(\delta \frac{q_{c,max} - q_{h,min}}{q_{c,max} - \delta q_{h,min}}, \delta \frac{p_{user,max} - p_{cheat,min}}{p_{user,max} - \delta(p_{user,max} - p_{cheat,min})} \right) \quad (7.115)$$

Case 3: $q_{c,max} \geq q_{h,min}$ and $p_{honest,max}q_{h,min} > p_{cheat,min}q_{c,max}$

$$p_{fine} > p_{user,max} - p_{cheat,min} \quad (7.116)$$

$$q_{jail} > \frac{p_{honest,max} - p_{cheat,min}}{p_{honest,max} - \delta p_{cheat,min}} \quad (7.117)$$

with overhead

$$O_{fine} = \delta \frac{p_{user,max} - p_{cheat,min}}{p_{honest}} \quad (7.118)$$

$$O_{jail} = \left(\frac{\delta}{1 - \delta} \right) \left(\frac{p_{honest,max} - p_{cheat,min}}{p_{honest,max}} \right) \quad (7.119)$$

When there is no information about the parameters (all minima go to 0, and all maxima go to ∞), the fine penalty must again go to infinity. However, because the sanction scales linearly with the QOS, the in-kind sanction can again force the cheating line to zero for *any* maximum cheating QOS. Therefore, it still has a solution

$$q_{jail} = 1 \quad (7.120)$$

which produces overhead

$$O_{jail} = \frac{\delta}{1 - \delta}. \quad (7.121)$$

This indicates that the in-kind sanction for this model is very strong, and so this model must be used with care to make sure the application of the sanction scales as assumed here. Presumably, in the real world, only a sanction with infinite overhead (like an infinite fine, a death sentence, or complete loss of service) should be effective for all possibilities if nothing is known about the parameters of the problem. However, note that as the sanctions get higher, it is impossible to implement an approximately infinite fine. It is possible to implement an approximately infinite jail sentence.

7.4 Concluding remarks

7.4.1 What we have learned so far

This chapter introduced overhead through wrongful convictions as the right way to quantitatively address the differences between fines and jail. It introduced two toy models to illustrate the problem in the context of spectrum sharing.

The first model emphasized the problem of uncertainty. For spectrum, and likely in other situations as well, the utility gained from a crime can be converted into in-kind and monetary terms. While the in-kind utility is easily measurable for radios in terms of quality of service, the monetary conversion can be very difficult. This added uncertainty makes a monetary sanction much less desirable.

The second toy showed that fines and in-kind punishments can operate on the choice to cheat in fundamentally different ways. For example, if one can provide a quality of service while cheating that is not attainable while honest, it may be worthwhile to cheat regardless of the fine. This is particularly important for the spectrum sharing context. Regulators *must* guarantee protection for primary systems. If the secondary can circumvent the rules simply by being willing to pay, the enforcement system has failed.

Only a small range of parameters in the second model indicated fines as preferable to jails from an overhead perspective.

From ease of implementation, the guarantee of deterrence, and the uncertainty and resulting overhead of estimating fines, this chapter has shown that for cognitive radio, in-kind sanctions are the better option.

7.4.2 Future work

This analysis is only in the beginning stages. More work is required to flesh out and truly analyze the general model to get a comprehensive concept of when fines or jails would be most appropriate.

This approach of evaluating enforcement strategies from an overhead perspective is different from the traditional law and economics approach. This seems to be because wrongful conviction is not treated as a parameter that can be controlled or measured in any way. Indeed, the control allowed by a technical identity system is one of the most useful features of the spectrum sharing problem. However, we believe this perspective has use beyond this particular context. This should be explored.

Chapter 8

Concluding remarks

This thesis concludes with one last comparison. Consider the Internet. It has many of the same elements as the problem of spectrum regulation: there is a multitude of different services and end users trying to accomplish different things. There is a limited resource to be shared – the wired backbone. There is a desire to allow freedom in connecting a new device or adding a new service.

The Internet was able to become a hot-bed of innovation because the protocols guarding entry only controlled the base problem: connectivity. IP provided a common language of connectivity on which everything else could be built.

In spectrum regulation, the desire for innovation amongst diverse services is the same, but the base regulatory problem is different. Different companies do not necessarily want their networks to directly communicate and different companies certainly do not want to have to jointly design their networks. The natural mode of interaction is *not* a common connectivity protocol, it is interference. The problem of regulation is therefore about controlling interference. The current approach to regulation does this by forcing companies to prove that their device is not capable of causing too much interference.

The philosophical shift proposed in this thesis is to directly regulate the interference itself instead of the technology that produces it. This thesis shows that spectrum jails can accomplish this goal. They are effective at incentivizing radios to follow sharing rules, they allow the introduction of new and complex sharing technologies, and they can be implemented through technical solutions already required to make TV whitespace databases secure.

New and complex spectrum sharing technologies can be introduced because spectrum jails enable light-handed regulations. The regulations require only that secondaries are certified to follow the operation tokens issued by a database. The database, then, gives tokens that permit the secondary to find and recover spectrum holes. The secondary can do this by implementing its own sensing protocol, or by getting information from other sensing or database services.

If the secondary causes harmful interference, the database will issue it a jail token instead of an operation token at the next request. The secondary will then have to turn off its radios and burn energy for the length of the token. This added functionality only really requires the database to identify secondaries causing interference. But this will already be required – even in the TV whitespaces, the database is in the right position to detect and turn off malfunctioning secondaries.

Notice that this certification requirement of following the database tokens can be applied to *any* secondary device, regardless of its sharing technology. This means that new technologies can be

easily certified. However, it also means that devices may be certified that cannot appropriately find spectrum holes. For this reason, spectrum sharing must be thought of as a band-expander. Devices that cannot appropriately find primaries should not be trying to share spectrum. If certification lets them through, they must have somewhere to operate where they will not be a burden on the enforcement system. They must have at least an unlicensed band to operate in.

Within this context, the following section reviews what this thesis has accomplished.

8.1 What has been accomplished

Prescription for trust through spectrum jails

Chapters 2 through 4 flesh out a secondary. It has three dimensions to its quality of service, average delay, average energy usage, and time-dependent packet constraints. By addressing the first two, this thesis has shown how to develop in-kind sanctions along the dimension of the utility so that the sanction is actually painful for the radio.

The secondary also has some cost to respect the priority of the primary, which it is rationally weighing against the cost of going to jail for interference. So, if the primary is very rarely transmitting, the secondary will rationally choose to ignore the primary.

Whether it goes to jail is determined primarily by what level of harm the secondary causes, whether it does so intentionally or not. It can also be sent to jail wrongfully.

In order to guarantee that this secondary will not cheat, regardless of its cost of sensing, the regulator must make decisions on when the primary will be protected. The regulator must set a minimum protected activity level for the primary. It must also decide how much harm is too much harm. The primary must actually use its band and must accept some small amount of harm. Then, the regulator must guarantee a particular size and quality of home band (or unlicensed band), either by having plentiful opportunities, or monitoring the quality and changing the sanction appropriately. With this information, a sanction can be set at – either certification time or at the time scale of the database – which will work for all secondaries at runtime.

Performance improvement with better technology

The performance is measured at runtime. Performance includes how well secondaries can recover spectrum holes and how well the primary is protected. The primary has a guaranteed amount of protection, but the secondary may cheat if the primary is not transmitting enough.

The performance depends both on the size of the sanction and the current quality of technology. If the secondary has a high sensing cost, it might cheat for some low value of primary utilization. If the regulator's identity system has a high wrongful conviction rate, the secondary will spend a lot of time in jail, perhaps choosing to just operate in the home band. If the secondary cannot help but cause lots of interference, it will again spend a lot of time in jail and may choose to leave the band entirely.

Even if the sanction stays the same, improving these technologies will improve performance. In fact, as sensing cost, wrongful conviction, and honest harm all go to zero (perfect technology), there will be no performance cost of implementing jail.

Requirements for identity

Making spectrum jails a reality is obviously dependent on being able to identify secondaries causing interference. Although this thesis does not address the question of how to implement identity, the results here do flesh out the problem. Before this thesis, the only understanding of identity was that it is required [75].

With the results in this thesis, it is known that for trust to be achieved, the identity system must be able to catch secondaries causing interference. It does *not* need to have low wrongful conviction. That needs to come only to improve performance. So, the first deployment of spectrum jails may include turning off all secondary devices in the area when any interference is observed. Later iterations will narrow those in jail to a smaller and smaller set. But the initial roll-out does not have to be perfect.

Trusting the primary

Although the current generation of primaries cannot do anything to protect themselves from secondary interference, future generations of primaries could be designed to participate in the enforcement system. Intuitively, the primary is in the perfect position to report interference because it has the best information about when harmful interference has occurred. But can the primary be trusted to report interference responsibly?

Chapter 5 has shown that the primary has incentive to report interference. The bigger question is what the primary will do in response to a secondary that it cannot easily coexist with. By setting the cost of reporting correctly, the regulator can incentivize the primary in the same way it did the secondary. For the primary, the regulator's action can determine whether the primary will wrongfully report interference or actively hire a "band-sitter" to help it defend its band.

Databases and evolvability

The results for the role of the primary and the specification of the identity problem indicate changes to the jail system as technology evolves. Because spectrum jails can be implemented through databases, this evolution is actually quite natural.

In the current style of FCC regulations, every new band required a new set of rules. So, the evolution of rules was naturally generational, with each band representing a different generation. Databases, or active spectrum managers, do not have to conform to this pattern. Multiple generations of devices and rules can coexist in the same band at the same time. Then, versions of the rules can be phased out as the generation of devices they apply to fall out of use.

A framework for multiple secondaries

With multiple secondaries and aggregate interference, it is not yet even clear what the desired secondary behavior is, so it is not yet possible to fully answer whether spectrum jails can incentivize this behavior. So, Chapter 6 gives a framework of basic results to begin the process of addressing multiple secondaries.

Stochastic geometry tools are used to understand how the distribution of aggregate interference from randomly placed secondaries changes with the no-talk radius. A phase shift occurs between "near" and "far" field when the distribution is dominated by the placement of one interferer

vs having contributions from the placement of many interferers. Far field can be well approximated by a Gaussian, allowing regulators to make rules in these situations based on simpler calculations.

Given this information, how should secondary transmission be controlled? This thesis investigates three possibilities, using information that the database already has about secondary placement. If the database cannot use any local information, the only option is a fixed no-talk radius that must be conservative enough to deal with placement uncertainty. This can only protect a primary to a certain probabilistic fidelity. At the same time, secondaries close to the primary receiver are not able to recover spectrum holes because of the very conservative rule.

Much better average performance can be obtained by changing the secondary transmit power or the no-talk radius in response to the actual realization of secondary placement. The primary will then always be protected against poor secondary placement, and the secondary service will only be degraded when it must be.

Finally, a simple spatial queuing model was introduced to show how the Markov jail model can be extended to include multiple secondaries. This shows promise of being able to extend trust even when aggregate interference is considered.

A note on fines

This thesis is built with an in-kind punishment mechanism. But this is not the traditional sanction. Fines are traditionally used because they do not destroy any resources; they simply move wealth around. Although jails are attractive from an implementation standpoint, are there other reasons to use in-kind vs monetary sanctions?

The law and economics literature addresses this problem of choosing imprisonment or fines for their deterrence ability with people. Chapter 7 adds to this discussion by introducing an engineering perspective. If the destruction-of-resource aspect is ignored, jails and fines can be compared on two levels: whether the sanction works, and how much overhead does it produce? The first is similar to the law and economics literature, but the second is a unique perspective.

The overhead comes from wrongful conviction. Traditional law and economics acknowledges wrongful conviction as bad, but does not use it as a quantitative tool. In the context of spectrum, without the emotional ramifications of addressing humans, wrongful conviction and overhead are performance parameters that can be optimized over the type of sanction.

With this perspective, it is easier to see the fundamental differences between jails and fines, particularly for radios. The first comparison is in uncertainty – while it is possible to understand the quality of service available while cheating, it is much harder to estimate the monetary value of a cheating transmission. So, a monetary sanction will need to be higher relative to an in-kind sanction to get the same deterrent effect. This leads directly to higher overhead.

This thesis also shows an example of how a fine can be thought of as a price for a higher quality of service. If cheating allows a quality of service not attainable with honest use, and the user is willing to pay a premium rate for that higher service, there is no way to deter cheating. If regulators need to prioritize protection of primaries, in-kind sanctions are the only viable option.

Comments on policy

The introductory chapter introduced many policy perspectives on the future of spectrum regulation. This thesis can contribute to this discussion in many ways.

Policy experts are concerned with what the future will look like – whether spectrum use should be governed by open access, managed commons, real-time markets, etc. Regardless of the chosen future, radios must be trusted to follow those rules. Spectrum jails can be applied in any of these systems to make sure that system will work as intended.

Policy literature also debates who should be responsible for dealing with spectrum disputes, traditional courts or spectrum tribunals. This thesis shows that they are missing a third option: tribunals implemented in technology. Spectrum jails have the benefits of traditional courts in that they use technologies that are already going to be required. But, they also have the benefit of being designed specifically for spectrum and therefore better able to address spectrum regulation needs.

Defining enforceable usage rights

Finally, much thought has been devoted to understanding what usage rights should look like for spectrum. This thesis gives direction to this discussion because usage rights must be defensible. The lack of quantitative study in the current literature makes it difficult to understand what defensible even means.

With spectrum jails, the question of enforceable usage rights can be answered. The usage right must actually be exercised: the primary must actually use its allocation or its transmissions cannot be protected. Also, the usage right includes a certain amount of harm the primary will have to accept from secondary users. Even with the current regulatory system, the purpose of building usage boxes in time/space/frequency with appropriate guard bands was to determine exactly how much a transmission can spill into another user's allocation. With spectrum jails, this same philosophy can be extended to a shared paradigm by directly regulating the amount of harm any secondary network can cause to any primary network.

Finally, usage rights for future generations of primaries may include some responsibility. PCAST suggests this concept through receiver standards that make a primary less vulnerable to interference. Spectrum jails introduces a required reporting cost primaries must pay in order to defend themselves by reporting interference.

8.2 Where to go from here: moving spectrum jails closer to reality

This thesis begins the research field of trust and enforcement for cognitive radio. It can therefore only address the beginning of the problem, and there is much work now to be done. Each chapter includes a future work section specific to the topic of that chapter. As there is so much future work to do, this conclusion chapter will not include those specifics. It will instead focus on the broader issues that need to be addressed to move spectrum jails closer to reality.

Requirements for trust

For the cases considered in this thesis, trust could be guaranteed regardless of secondary technology. But, there are cases not included in this thesis. So, to truly guarantee trust, work needs to be done to extend spectrum jails to include all kinds of secondaries. This includes utility functions that have time dependencies. It also includes completing the treatment of multiple secondaries to appropriately manage aggregate interference.

There is also significant work to be done to secure tokens and decisions. For example, the secondary must follow orders to turn off and burn energy. This operation must be built into the device in a way that either can be verified or cannot be tampered with at runtime. Also, this thesis assumes that the secondary will be able to decide to operate just in the home band where it will not be subject to jail. How can this decision be trusted if the choices result in just a different waveform produced in software? Especially if the home band changes with the database token, how can this be secured?

Finally, an identity system must be built. For trust, this includes just a catching ability, but what is the right method of discovering interference and connecting it to a particular radio? Perhaps the right approach includes a system whereby radios can develop alibis to prove that they did not cause the interference in question.

Improving performance

Improving performance really includes all kinds of technical improvements to better find spectrum holes. But particularly for enforcement, performance gains will come with improved wrongful conviction rate. New ideas about how to build adaptive jails will also allow performance improvement in the future.

Other kinds of enforcement

This thesis addresses only spectrum jails, which represents one possibility combining *a priori* certification and *a posteriori* policing. While we believe this is a good enforcement mechanism for radios, it is certainly not the only option. What other methods are there?

Perhaps the enforcement mechanism could go beyond just rationality, tapping into the long term relationships between different companies. Perhaps human traits like altruism could be built into devices, or some kind of network effect could be created within cognitive radio.

This thesis has provided a baseline and metrics for comparison. Now, new ideas are needed to test whether this is actually the best option.

Policy and law

Spectrum sharing with databases in TV whitespaces represents the first time the FCC has allowed spectrum regulatory decisions to be made by technology. This first step includes only the modest (and well specified) decision of band assignment based on location. This thesis has argued that to really exploit the potential of cognitive radio, the scope of regulatory decisions made by technology must expand dramatically.

The first expansion should be allowing databases to use secondary location information to improve primary protection and secondary performance in the TV bands. This decision can be well specified so that the database is actually just propagating a context-dependent rule instead of subjectively making decisions. Philosophically, this expanded capability is very similar to what is currently allowed.

Turning the database into a spectrum manager with enforcement capability is a philosophically much larger change. The FCC would have to either manage the databases itself or deputize private companies to act as spectrum police.

Although all of the decisions can again be well specified by the FCC (by certifying an identity system, etc.), legally moving enforcement activity into technology seems like a difficult problem. Much work will need to be done in policy and law to really embrace automated enforcement.¹

However, the benefits of *a posteriori* enforcement are potentially huge. Certification can be stream-lined. *Any* new technology can be trusted, regardless of its solution to finding spectrum holes. Regulatory overhead can be significantly reduced, especially as technology improves. Harmful interference can be directly regulated, instead of implicitly controlled through specific technology requirements. The only way to get these benefits is to automate the policing and enforcement activities. Addressing the corresponding legal and policy hurdles is essential.

¹Perhaps the right place to start is through the work investigating rights for robots [158, 159].

Bibliography

- [1] J. Mitola, “Cognitive Radio: an integrated agent architecture for software defined radio,” Ph.D. Thesis, KTH Royal Inst. of Tech., Stockholm, Sweden, 2000.
- [2] http://fjallfoss.fcc.gov/edocs_public/attachmatch/FCC-07-132A1.pdf.
- [3] “Cisco visual networking index: Global mobile data traffic forecast update, 20122017,” http://www.cisco.com/en/US/solutions/collateral/ns341/ns525/ns537/ns705/ns827/white_paper_c11-520862.html.
- [4] “Connecting America: The National Broadband Plan,” <http://www.broadband.gov/plan/>.
- [5] M. A. McHenry and K. Steadman, “Spectrum occupancy measurements, location 1 of 6: Riverbend park, Great Falls, Virginia,” Shared Spectrum Company, Tech. Rep., 2005. [Online]. Available: http://www.sharespectrum.com/inc/content/measurements/nsf/1_NSF_Riverbend_Park_Report.pdf
- [6] S. Mishra and A. Sahai, “How much white space has the FCC opened up,” *IEEE Communication Letters*, 2010.
- [7] D. Willkomm, S. Machiraju, J. Bolot, and A. Wolisz, “Primary users in cellular networks: A large-scale measurement study,” in *New frontiers in dynamic spectrum access networks, 2008. DySPAN 2008. 3rd IEEE symposium on*. IEEE, 2008, pp. 1–11.
- [8] D. Tse and P. Viswanath, *Fundamentals of Wireless Communication*. Cambridge University Press, May 2005.
- [9] K. Yee, “Numerical solution of initial boundary value problems involving Maxwell’s equations in isotropic media,” *Antennas and Propagation, IEEE Transactions on*, vol. 14, no. 3, pp. 302–307, 1966.
- [10] D. Hatfield and P. Weiser, “Toward Property Rights in Spectrum: The Difficult Policy Choices Ahead,” *CATO Institute*, no. 575, Aug. 2006.
- [11] A. V. Oppenheim, A. S. Willsky, and S. H. Nawab, *Signals and systems*. Prentice-Hall Englewood Cliffs, NJ, 1983, vol. 2.
- [12] B. Hofmann-Wellenhof, H. Lichtenegger, and J. Collins, “Global positioning system. theory and practice.” *Global Positioning System. Theory and practice.*, by Hofmann-Wellenhof, B.; Lichtenegger, H.; Collins, J.. Springer, Wien (Austria), 1993, 347 p., ISBN 3-211-82477-4, Price DM 79.00. ISBN 0-387-82477-4 (USA)., vol. 1, 1993.

- [13] R. H. Coase, "The Federal Communications Commission," *Journal of Law and Economics*, vol. 2, pp. 1–40, Oct. 1959.
- [14] T. W. Hazlett, "Optimal Abolition of FCC Spectrum Allocation," *Journal of Economic Perspectives*, vol. 22, no. 1, pp. 103–128, 2008.
- [15] K. Woyach, "Crime and Punishment for Cognitive Radios," Masters Thesis, University of California, Berkeley, 2008.
- [16] NTIA, "U.S. Frequency Allocations," <http://www.ntia.doc.gov/osmhome/allochrt.pdf>.
- [17] <http://www.fcc.gov/complaints>.
- [18] <http://transition.fcc.gov/eb/sed/ulo.html>.
- [19] J. L. King and J. West, "Ma Bell's orphan: US cellular telephony, 1947–1996," *Telecommunications Policy*, vol. 26, no. 34, pp. 189 – 203, 2002, [jce:titleMaking the mobile revolution: standards and other legaciesjce:title](#).
- [20] G. M. Rebeiz, *RF MEMS, Theory, Design and Technology*. John Wiley & Sons, 2003.
- [21] "In the matter of unlicensed operation in the tv broadcast bands: Second report and order and memorandum opinion and order," http://hraunfoss.fcc.gov/edocs_public/attachmatch/FCC-08-260A1.pdf, Federal Communications Commission, Tech. Rep. 08-260, Nov. 2008.
- [22] "Spectrum policy task force report," Federal Communications Commission, Tech. Rep. 02-135, Nov. 2002. [Online]. Available: http://hraunfoss.fcc.gov/edocs_public/attachmatch/DOC-228542A1.pdf
- [23] <http://adaptrum.com/>.
- [24] "Third memorandum opinion and order," http://transition.fcc.gov/Daily_Releases/Daily_Business/2012/db0405/FCC-12-36A1.pdf, 2012.
- [25] T. M. Cover and J. A. Thomas, *Elements of Information Theory 2nd Edition*. Wiley-Interscience, 2006.
- [26] A. El Gamal and Y.-H. Kim, *Network information theory*. Cambridge University Press, 2011.
- [27] A. Goldsmith, S. A. Jafar, N. Jindal, and S. Vishwanath, "Capacity limits of MIMO channels," *Selected Areas in Communications, IEEE Journal on*, vol. 21, no. 5, pp. 684–702, 2003.
- [28] V. Cadambe and S. A. Jafar, "Interference alignment and degrees of freedom of the K-user interference channel," *IEEE Transactions on Information Theory*, vol. 54, no. 8, pp. 3425–3441, Aug. 2008.
- [29] P. Grover, K. Woyach, and A. Sahai, "Towards a communication-theoretic understanding of system-level power consumption," *Selected Areas in Communications, IEEE Journal on*, vol. 29, no. 8, pp. 1744–1755, 2011.

- [30] S. M. Mishra, A. Sahai, and R. W. Broderson, "Cooperative sensing among Cognitive Radios," *ICC*, June 6-10 2006.
- [31] K. Harrison and A. Sahai, "Potential collapse of whitespaces and the prospect for a universal power rule," in *New Frontiers in Dynamic Spectrum Access Networks (DySPAN), 2011 IEEE Symposium on*. IEEE, 2011, pp. 316–327.
- [32] ———, "Seeing the bigger picture: context-aware regulations," in *Dynamic Spectrum Access Networks (DYSPAN), 2012 IEEE International Symposium on*. IEEE, 2012, pp. 21–32.
- [33] R. Murty, R. Chandra, T. Moscibroda, and P. Bahl, "Senseless: A database-driven white spaces network," *Mobile Computing, IEEE Transactions on*, vol. 11, no. 2, pp. 189–203, 2012.
- [34] A. Stirling, "Exploiting hybrid models for spectrum access," *Proceedings of the Fifth IEEE International Symposium on New Frontiers in Dynamic Spectrum Access Networks*, Oct. 2011.
- [35] DARPA, XG, "Working Group DARPA XG Policy Language Framework," vol. 1, 2004.
- [36] G. Denker, D. Elenius, R. Senanayake, M.-O. Stehr, and D. Wilkins, "A policy engine for spectrum sharing," in *New Frontiers in Dynamic Spectrum Access Networks, 2007. DySPAN 2007. 2nd IEEE International Symposium on*, 2007, pp. 55–65.
- [37] D. Wilkins, G. Denker, M.-O. Stehr, D. Elenius, R. Senanayake, and C. Talcott, "Policy-based cognitive radios," *Wireless Communications, IEEE*, vol. 14, no. 4, pp. 41–46, 2007.
- [38] F. Perich, "Policy-based network management for NeXt generation spectrum access control," in *Proceedings of the Second IEEE International Symposium on New Frontiers in Dynamic Spectrum Access Networks*, Dublin, Ireland, Apr. 2007.
- [39] A. R. Young, N. J. Kaminski, A. Fayed, and C. W. Bostian, "CSERE (Cognitive System Enabling Radio Evolution): A modular and user-friendly cognitive engine," in *Dynamic Spectrum Access Networks (DYSPAN), 2012 IEEE International Symposium on*. IEEE, 2012, pp. 59–67.
- [40] M. M. Kokar, D. Hillman, S. Li, B. Fette, P. Marshall, M. Cummings, T. Martin, and J. Strassner, "Towards a unified policy language for future communication networks: a process," in *New Frontiers in Dynamic Spectrum Access Networks, 2008. DySPAN 2008. 3rd IEEE Symposium on*. IEEE, 2008, pp. 1–10.
- [41] B. Bahrak, J.-M. Park, and H. Wu, "Ontology-based spectrum access policies for policy-based cognitive radios," in *Dynamic Spectrum Access Networks (DYSPAN), 2012 IEEE International Symposium on*. IEEE, 2012, pp. 489–500.
- [42] J. Riihijarvi, M. Petrova, V. Atanasovski, and L. Gavrilovska, "Extending policy languages with utility and prioritization knowledge: the capri approach," in *New Frontiers in Dynamic Spectrum, 2010 IEEE Symposium on*. IEEE, 2010, pp. 1–9.

- [43] B. Bahrak, A. Deshpande, M. Whitaker, and J.-M. Park, "BRESAP: a policy reasoner for processing spectrum access policies represented by binary decision diagrams," in *New Frontiers in Dynamic Spectrum, 2010 IEEE Symposium on*. IEEE, 2010, pp. 1–12.
- [44] D. Elenius, G. Denker, M.-O. Stehr, R. Senanayake, C. Talcott, and D. Wilkins, "Coral-policy language and reasoning techniques for spectrum policies," in *Policies for Distributed Systems and Networks, 2007. POLICY'07. Eighth IEEE International Workshop on*. IEEE, 2007, pp. 261–265.
- [45] A. M. Turing, "On computable numbers, with an application to the Entscheidungsproblem," *Proceedings of the London mathematical society*, vol. 42, no. 2, pp. 230–265, 1936.
- [46] K. Gödel, "Über formal unentscheidbare Sätze der Principia Mathematica und verwandter Systeme," *I. Monatshefte für Mathematik und Physik*, vol. 38, pp. 173–198, 1931.
- [47] T. W. Hazlett, "The Rationality of U.S. Regulation of the Broadcast Spectrum," *Journal of Law and Economics*, vol. 33, no. 1, pp. 133–175, Apr 1990.
- [48] "Comments of 37 Concerned Economists. Promoting Efficient use of Spectrum Through Elimination of Barriers to the Development of Secondary Markets," *Before the Federal Communications Commission*, no. WT Docket 00-230, Feb. 2001.
- [49] A. S. de Vany, R. D. Eckert, C. J. Meyers, D. J. O'Hara, and R. C. Scott, "A Property System for Market Allocation of the Electromagnetic Spectrum: A Legal-Economic-Engineering Study," *Stanford Law Review*, vol. 21, no. 6, pp. 1499–1561, Jun. 1969.
- [50] H. A. Shelanski and P. W. Huber, "Administrative Creation of Property Rights to Radio Spectrum," *Journal of Law and Economics*, vol. 41, no. 2, pp. 581–607, Oct 1998.
- [51] D. W. Webbink, "Radio licenses and frequency spectrum use property rights," *Communications and the Law*, no. 3, 1987.
- [52] J. McMillan, "Why Auction Spectrum," *Telecommunications Policy*, vol. 19, no. 3, pp. 191–199, Apr 1995.
- [53] S. M. Benjamin, "Spectrum Abundance and the Choice Between Private and Public Control," *New York University Law Review*, vol. 78, no. 6, pp. 2007–2102, Dec 2003.
- [54] E. Noam, "Spectrum Auctions: Yesterday's Heresy, Today's Orthodoxy, Tomorrow's Anachronism. Taking the Next Step to Open Spectrum Access," *Journal of Law and Economics*, vol. 41, no. 2, pp. 765–790, 1998.
- [55] K. Werbach, "Supercommons: Toward a Unified Theory of Wireless Communication," *Texas Law Review*, vol. 82, no. 4, pp. 863–973, March 2004.
- [56] S. Buck, "Replacing Spectrum Auctions with a Spectrum Commons," *Stanford Technology Law Review*, 2002.
- [57] Y. Benkler, "Overcoming Agoraphobia: Building the Commons of the Digitally Networked Environment," *Harvard Journal of Law and Technology*, vol. 11, pp. 287–400, Winter 1998.

- [58] R. J. Matheson, "The electrospace model as a frequency management tool," *Proceedings of the ISART Conference*, Mar 2003.
- [59] T. W. Hazlett, "A Law and Economics Approach to Spectrum Property Rights: A Response to Weiser and Hatfield," *George Mason Law Review*, vol. 15, no. 4, Summer 2008.
- [60] C. R. Sunstein, "Incompletely theorized agreements," *Harvard Law Review*, vol. 108, no. 7, pp. 1733 – 1772, May 1995.
- [61] P. J. Weiser and D. N. Hatfield, "Policing the Spectrum Commons," *Fordham Law Review*, vol. 74, no. 2, pp. 663–694, 2005.
- [62] E. Kwerel and J. Williams, "A Proposal for a Rapid Transition to Market Allocation of Spectrum," *FCC Office of Strategic Planning Working Paper Series*, no. 38, Nov 2002.
- [63] C. Bazelon, "Licensed or Unlicensed: The Economic Considerations in Incremental Spectrum Allocations," in *Proceedings of the Third IEEE International Symposium on New Frontiers in Dynamic Spectrum Access Networks*, Chicago, IL, Oct. 2008.
- [64] M. M. Bykowsky, K. R. Carter, M. A. Olson, and W. W. Sharkey, "Enhancing Spectrum's Value Through Market-informed Congestion Etiquettes," *FCC Office of Strategic Planning Working Paper Series*, no. 41, Feb 2008.
- [65] M. M. Bykowsky, M. A. Olson, and W. W. Sharkey, "Modeling the Efficiency of Spectrum Designated to Licensed Service and Unlicensed Operations," *FCC Office of Strategic Planning Working Paper Series*, no. 42, Feb 2008.
- [66] —, "A Market-based Approach to Establishing Licensing Rules: Licensed Versus Unlicensed Use of Spectrum," *FCC Office of Strategic Planning Working Paper Series*, no. 43, Feb 2008.
- [67] P. J. Weiser and D. Hatfield, "Spectrum Policy Reform and the Next Frontier of Property Rights," *George Mason Law Review*, vol. 60, no. 3, 2008.
- [68] A. Sahai, K. Woyach, K. Harrison, H. Palaiyanur, and R. Tandra, "Towards a "theory of spectrum zoning"," *Allerton*, Oct. 2009.
- [69] H. Palaiyanur, K. A. Woyach, R. Tandra, and A. Sahai, "Spectrum zoning as robust optimization," *Proceedings of the Fourth IEEE International Symposium on New Frontiers in Dynamic Spectrum Access Networks*, Apr. 2010.
- [70] E. Goodman, "Spectrum Rights in the Telecosm to Come," *San Diego Law Review*, vol. 41, no. 269, 2004.
- [71] B. Lennett, "The Lobby that Cried Wolf: NAB Campaign Against Using TV White Space Follows a Familiar Script," *New America Foundation Wireless Future Program Issue Brief*, no. 23, Oct 2008.
- [72] J. C. Whitaker, *DTV: The revolution in digital video*. McGraw-Hill Professional, 2001.

- [73] K. Harrison, S. M. Mishra, and A. Sahai, "How much white-space capacity is there?" in *New Frontiers in Dynamic Spectrum, 2010 IEEE Symposium on*. IEEE, 2010, pp. 1–10.
- [74] R. Tandra and A. Sahai, "SNR walls for signal detection," *IEEE Journal of Selected Topics in Signal Processing*, vol. 2, no. 1, pp. 4–17, Feb. 2008.
- [75] "Second Report and Order and Memorandum Opinion and Order: Unlicensed Operation in the TV Broadcast Bands," http://hraunfoss.fcc.gov/edocs_public/attachmatch/FCC-08-260A1.pdf, November 2008.
- [76] O. Grondalen, M. Lahteenoja, and P. Gronsund, "Evaluation of business cases for a cognitive radio network based on wireless sensor network," in *New Frontiers in Dynamic Spectrum Access Networks (DySPAN), 2011 IEEE Symposium on*. IEEE, 2011, pp. 242–253.
- [77] <http://www.whitehouse.gov/the-press-office/statement-president-national-broadband-plan>.
- [78] "Expanding the economic and innovation opportunities of spectrum through incentive auctions," <http://www.fcc.gov/document/broadcast-television-spectrum-incentive-auction-nprm>, 2012.
- [79] "Report to the President realizing the full potential of government-held spectrum to spur economic growth," http://www.whitehouse.gov/sites/default/files/microsites/ostp/pcast_spectrum_report_final_july_20_2012.pdf.
- [80] P. Marshall, "Dynamic spectrum access as a mechanism for transition to interference tolerant systems," *Proceedings of the Fourth IEEE International Symposium on New Frontiers in Dynamic Spectrum Access Networks*, Apr. 2010.
- [81] "FCC. Comment deadlines established regarding the GPS-LightSquared technical working group," <http://www.fcc.gov/document/comment-deadlines-established-regarding-gps-lightsquared-technical-working-group-report>.
- [82] Receivers and Standards Working Group, "Interference Limits Policy: The use of harm claim thresholds to improve the interference tolerance of wireless systems," FCC Technological Advisory Council, Tech. Rep., Feb 2013.
- [83] L. Doyle, J. McMenamy, and T. K. Forde, "Regulating for carrier aggregation & getting spectrum management right for the longer term," in *Dynamic Spectrum Access Networks (DYSPAN), 2012 IEEE International Symposium on*. IEEE, 2012, pp. 10–20.
- [84] E. Dahlman, S. Parkvall, J. Skold, and P. Beming, *3G evolution: HSPA and LTE for mobile broadband*. Access Online via Elsevier, 2010.
- [85] L. Lessig, *Code v2.0*. Basic Books, 2006.
- [86] N. Isaacs, "Barrier Activities and the Courts: A Study in Anti-Competitive Law," *Law and Contemporary Problems*, vol. 8, no. 2, pp. 382 – 390, 1941.
- [87] M. Weiss, P. Krishnamurthy, L. E. Doyle, and K. Pelechrinis, "When is electromagnetic spectrum fungible?" in *Dynamic Spectrum Access Networks (DYSPAN), 2012 IEEE International Symposium on*. IEEE, 2012, pp. 349–357.

- [88] A. Selim, B. Ozgul, and L. Doyle, "Efficient sidelobe suppression for ofdm systems with peak-to-average power ratio reduction," in *Dynamic Spectrum Access Networks (DYSPAN), 2012 IEEE International Symposium on*. IEEE, 2012, pp. 510–516.
- [89] A. Stirling, "Exploiting hybrid models for spectrum access: Building on the capabilities of geolocation databases," in *New Frontiers in Dynamic Spectrum Access Networks (DySPAN), 2011 IEEE Symposium on*. IEEE, 2011, pp. 47–55.
- [90] R. Tandra, S. M. Mishra, and A. Sahai, "What is a spectrum hole and what does it take to recognize one?" *Proceedings of the IEEE*, Jan. 2009.
- [91] R. Tandra, A. Sahai, and V. V. Veeravalli, "Space-time metrics for spectrum sensing," in *New Frontiers in Dynamic Spectrum, 2010 IEEE Symposium on*. IEEE, 2010, pp. 1–12.
- [92] D. W. Webink, "Impact of the UHF Promotion: The All-Channel Television Receiver Law," *Law & Contemp. Probs.*, vol. 34, p. 535, 1969.
- [93] <http://www.fcc.gov/digital-television>.
- [94] A. Sahai, "Anytime information theory," Ph.D. dissertation, Massachusetts Institute of Technology, 2001.
- [95] T. Hazlett and M. L. Spitzer, "Advanced wireless technologies and public policy," *Southern California Law Review*, 2004.
- [96] <http://www.ieee802.org/11/>.
- [97] C. Fowler, J. Entzminger, and J. Corum, "Assessment of ultra-wideband (UWB) technology," *IEEE Aerospace and Electronic Systems Magazine*, vol. 5, no. 11, pp. 45–49, Nov. 1990.
- [98] V. V. Patel, C. Cao, N. Hovakimyan, K. A. Wise, and E. Lavretsky, "L1 adaptive controller for tailless unstable aircraft in the presence of unknown actuator failures," *International Journal of Control*, vol. 82, no. 4, pp. 705–720, 2009.
- [99] M. K. Sparrow, *Imposing duties: government's changing approach to compliance*. Greenwood Publishing Group, 1994.
- [100] K. Levchenko, A. Pitsillidis, N. Chachra, B. Enright, M. Félegyházi, C. Grier, T. Halvorson, C. Kanich, C. Kreibich, H. Liu *et al.*, "Click trajectories: End-to-end analysis of the spam value chain," in *Security and Privacy (SP), 2011 IEEE Symposium on*. IEEE, 2011, pp. 431–446.
- [101] <http://www.ieee802.org/22/>.
- [102] D. Ariely, *Predictably irrational, revised and expanded edition: The hidden forces that shape our decisions*. HarperCollins, 2009.
- [103] R. Etkin, A. Parekh, and D. Tse, "Spectrum sharing for unlicensed bands," vol. 25, no. 3. IEEE, 2007, pp. 517–528.

- [104] G. Costa, L. Garcia, A. Cattoni, K. Pedersen, and P. Mogensen, "Dynamic spectrum sharing in femtocells: A comparison of selfish versus altruistic strategies," in *Vehicular Technology Conference (VTC Fall), 2011 IEEE*, sept. 2011, pp. 1 –5.
- [105] J. Suris, L. Dasilva, Z. Han, A. Mackenzie, and R. Komali, "Asymptotic optimality for distributed spectrum sharing using bargaining solutions," *Wireless Communications, IEEE Transactions on*, vol. 8, no. 10, pp. 5225 –5237, october 2009.
- [106] C. Beccaria, *An Essay on Crimes and Punishments*. W.C. Little, Albany, 1872.
- [107] C. Montesquieu, *The Spirit of the Laws*. University of California Press, Berkeley, 1977.
- [108] J. Bentham, *An introduction to the principles of morals and legislation*, in: *The Utilitarians*. Anchor Books, Garden City, NY, 1973.
- [109] R. A. Posner, *Economic Analysis of Law*, 3rd ed. Little, Brown, and Company, 1986.
- [110] A. M. Polinsky and S. Shavell, "The theory of public enforcement of law," *Journal of Economic Literature*, vol. 38, pp. 45 – 76, Mar. 2000.
- [111] L. Kaplow and S. Shavell, "Economic analysis of law," *NBER Working Paper*, no. 6960, Feb. 1999.
- [112] A. M. Polinsky, "The optimal use of fines and imprisonment when wealth is unobservable," *Journal of Public Economics*, vol. 90, no. 4-5, pp. 823 – 835, 2006.
- [113] W. E. Woodson and D. W. Conover, *Human engineering guide for equipment designers*. University of California Pr, 1964.
- [114] C. C. Chu, "Are fines more efficient than imprisonment?" *Journal of Public Economics*, vol. 51, no. 3, pp. 391 – 413, Jul. 1993.
- [115] D. R. Wong, "Stigma: A more efficient alternative to fines in deterring corporate misconduct," *California Criminal Law Review*, vol. 3, no. 3, Oct. 2000.
- [116] U. Gneezy and A. Rustichini, "A fine is a price," *Journal of Legal Studies*, vol. 29, no. 1, Jan. 2009.
- [117] M. Weiss, S. Delaere, and W. Lehr, "Sensing as a Service: An Exploration into Practical Implementations of DSA," in *New Frontiers in Dynamic Spectrum, IEEE Symposium on*, Singapore, Apr. 2010.
- [118] S. Mishra, A. Sahai, and R. Broderson, "Cooperative sensing among cognitive radios," *IEEE ICC*, pp. 1658 – 1663, Jun. 2006.
- [119] K. Woyach, A. Sahai, G. Atia, and V. Saligrama, "Crime and punishment for cognitive radios," *Allerton*, 2008.
- [120] A. Sahai, K. Woyach, G. Atia, and V. Saligrama, "A technical framework for light-handed regulation of cognitive radios," *IEEE Communications Magazine*, pp. 96 – 102, Mar. 2009.

- [121] A. Sahai, S. M. Mishra, R. Tandra, and K. A. Woyach, "Cognitive radios for spectrum sharing," *IEEE Signal Processing Magazine*, Jan. 2009.
- [122] G. Atia, A. Sahai, and V. Saligrama, "Spectrum enforcement and liability assignment in cognitive radio systems," in *Proceedings of the Third IEEE International Symposium on New Frontiers in Dynamic Spectrum Access Networks*, Chicago, IL, Oct. 2008.
- [123] K. Ellis and N. Serinken, "Characteristics of radio transmitter fingerprints," *Radio Science*, vol. 36, no. 4, pp. 585–597, 2001.
- [124] P. Moulin and R. Koetter, "A framework for the design of good watermark identification codes," in *Electronic Imaging 2006*. International Society for Optics and Photonics, 2006, pp. 60 721H–60 721H.
- [125] R. Ahlswede and G. Dueck, "Identification via channels," *Information Theory, IEEE Transactions on*, vol. 35, no. 1, pp. 15–29, 1989.
- [126] T. S. Han and S. Verdu, "New results in the theory of identification via channels," *Information Theory, IEEE Transactions on*, vol. 38, no. 1, pp. 14–25, 1992.
- [127] K. Woyach, K. Harrison, G. Ranade, and A. Sahai, "Comments on unknown channels," in *Information Theory Workshop (ITW), 2012 IEEE*. IEEE, 2012, pp. 172–176.
- [128] N. Xu, S. Rangwala, K. K. Chintalapudi, D. Ganesan, A. Broad, R. Govindan, and D. Estrin, "A wireless sensor network for structural monitoring," in *Proceedings of the 2nd international conference on Embedded networked sensor systems*. ACM, 2004, pp. 13–24.
- [129] G. Werner-Allen, J. Johnson, M. Ruiz, J. Lees, and M. Welsh, "Monitoring volcanic eruptions with a wireless sensor network," in *Wireless Sensor Networks, 2005. Proceedings of the Second European Workshop on*. IEEE, 2005, pp. 108–120.
- [130] B. Warneke, M. Last, B. Liebowitz, and K. Pister, "Smart dust: communicating with a cubic-millimeter computer," *Computer*, p. 44, 2001.
- [131] D. Cabric, M. Chen, D. Sobel, S. Wang, J. Yang, and R. Brodersen, "Novel radio architectures for UWB, 60 GHz, and cognitive wireless systems," *EURASIP Journal on Wireless Communications and Networking*, vol. 2006, no. 2, p. 22, 2006.
- [132] Z. Ji and K. J. R. Liu, "Dynamic spectrum sharing: a game theoretical overview," *IEEE Communications Magazine*, pp. 88 – 94, May 2007.
- [133] B. Wang, Y. Wu, and K. R. Liu, "Game theory for cognitive radio networks: An overview," *Computer Networks*, vol. 54, no. 14, pp. 2537 – 2561, 2010. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S1389128610001064>
- [134] K. Woyach and A. Sahai, "Should primaries be considered victims or police?" *Proceedings of the Sixth IEEE International Symposium on New Frontiers in Dynamic Spectrum Access Networks*, Oct. 2012.

- [135] J. W. DeCew, "The scope of privacy in law and ethics," *Law and Philosophy*, vol. 5, no. 2, pp. 145 – 173, Aug. 1986.
- [136] J. H. Ely, "The wages of crying wolf: A comment on Roe v. Wade," *The Yale Law Journal*, vol. 82, no. 5, pp. 920 – 949, Apr. 1973.
- [137] D. Satapathy and J. Peha, "Performance of unlicensed devices with a spectrum etiquette," in *Global Telecommunications Conference, 1997. GLOBECOM '97., IEEE*, vol. 1, nov 1997, pp. 414 –418 vol.1.
- [138] S. Adlakha, R. Johari, and A. J. Goldsmith, "Competition in wireless systems via bayesian interference games," *CoRR*, vol. abs/0709.0516, 2007.
- [139] H. von Stackelberg, *Market Structure and Equilibrium: 1st Edition Translation into English*, Bazin, Urch and Hill. Springer, 2011.
- [140] N. Hoven and A. Sahai, "Power scaling for cognitive radio," in *Wireless networks, communications and mobile computing, 2005 International Conference on*, vol. 1. IEEE, 2005, pp. 250–255.
- [141] M. Vu, N. Devroye, and V. Tarokh, "The primary exclusive region in cognitive networks," *5th IEEE Consumer Communications and Networking Conference, (CNCC) 2008*, pp. 1014 – 1019, Jan. 2008.
- [142] F. Baccelli and B. Blaszczyzyn, "Stochastic geometry and wireless networks," URL for Volume I: <http://hal.inria.fr/inria-00403039> and for Volume II: <http://hal.inria.fr/inria-00403040>, 2009. To appear in *Foundation and Trend in Networking*, NOW publisher.
- [143] M. Haenggi and R. K. Ganti, "Interference in large wireless networks," *Foundations and Trends in Networking*, no. 2, pp. 127–248, 2009.
- [144] T. V. Nguyen and F. Baccelli, "A probabilistic model of carrier sensing based cognitive radio," in *New Frontiers in Dynamic Spectrum (DySPAN), IEEE Symposium on*, Singapore, April 2010.
- [145] K. Woyach, P. Grover, and A. Sahai, "Near vs. Far Field: Interference Aggregation in TV Whitespaces," in *Global Telecommunications Conference (GLOBECOM 2011), 2011 IEEE*. IEEE, 2011, pp. 1–5.
- [146] Y. Selen and J. Kronander, "Optimizing power limits for white space devices under a probability constraint on aggregated interference," in *Dynamic Spectrum Access Networks (DYSPAN), 2012 IEEE International Symposium on*. IEEE, 2012, pp. 201–211.
- [147] J. Abate and W. Whitt, "The fourier-series method for inverting transforms of probability distributions," *Queueing systems*, vol. 10, no. 1-2, pp. 5–87, 1992.
- [148] wikipedia, "Central Limit Theorem," http://en.wikipedia.org/wiki/Central_limit_theorem.
- [149] R. Ganti, F. Baccelli, and J. Andrews, "Series Expansion for Interference in Wireless Networks," *Arxiv preprint arXiv:1101.3824, Submitted to IEEE Trans. Inform. Theory*, 2011.

- [150] “Second memorandum opinion and order,” Sep. 2010. [Online]. Available: http://www.fcc.gov/Daily_Releases/Daily_Business/2010/db1025/FCC-10-174A1.pdf
- [151] “TV white spaces: A consultation on white space device requirements,” Nov. 2012. [Online]. Available: <http://stakeholders.ofcom.org.uk/binaries/consultations/whitespaces/summary/condoc.pdf>
- [152] L. Kleinrock, “Queueing systems. volume 1: Theory,” 1975.
- [153] R. A. Posner, “Optimal sentences for white-collar criminals,” *Am. Crim. L. Rev.*, vol. 17, p. 409, 1979.
- [154] J. John R. Lott, “Should the wealthy be able to “buy justice”?” *The Journal of Political Economy*, vol. 95, no. 6, pp. 1307 – 1316, Dec. 1987.
- [155] L. Kaplow, “The optimal probability and magnitude of fines for acts that definitely are undesirable,” *National Bureau of Economic Research Working Paper Series*, no. 3008, Jun. 1989.
- [156] I. P. L. Png, “Optimal subsidies and damages in the presence of judicial error,” *International Review of Law and Economics*, vol. 6, no. 1, pp. 101 – 105, Jun. 1986.
- [157] A. Volokh, “N guilty men,” *Univ. of PA Law Rev.*, no. 1, pp. 173–216, 1997.
- [158] P. M. Asaro, “Robots and Responsibility from a Legal Perspective,” *Proceedings of the IEEE International Conference on Robotics and Automation*, Apr. 2007.
- [159] L. B. Solum, “Legal Personhood for Artificial Intelligences,” *North Carolina Law Review*, pp. 1231–1287, Apr. 1992.