Early Resolution of Uncertainty in Public Policy: Applications to Financial Regulation, the U.S. Supreme Court, and International Taxation

A DISSERTATION

SUBMITTED TO THE FACULTY OF THE GRADUATE SCHOOL

OF THE UNIVERSITY OF MINNESOTA

BY

ANDREW SNYDER DUST

IN PARTIAL FULFILLMENT OF THE REQUIREMENTS

FOR THE DEGREE OF

DOCTOR OF PHILOSOPHY

JAN WERNER, ADVISOR

May, 2013

Acknowledgements

I thank Jan Werner for his support and guidance over the last five years. He has been a terrific teacher, advisor, and mentor. His comments have greatly improved this dissertation and made me a better researcher.

Beth Allen provided wonderful career advice and innumerable improvements in my research. I am grateful to David Rahman for giving me an opportunity to learn about contract theory as his research assistant. I want to thank Guatam Ray for his work on my committee and his fascinating class on the theory of the firm. I also want to thank Kim-Sau Chung for his early support and suggestions on Chapter 3.

Peter Orazem, David Frankel, and Charles Zheng all graciously invested their time in me as an undergraduate student. I have been well served by listening to their advice.

Lastly, my deep gratitude goes to my family. Without their support this dissertation would not exist.

Dedication

To Emily, for everything.

Abstract

This dissertation examines the trade-off between the additional benefit of waiting for better information and the additional cost of delaying a decision. The chapters below focus on government policies designed to address this trade-off and show when better social outcomes can be achieved if early decisions–even those based on scarce and unreliable information–are allowed. Chapter 1 shows that conditioning banker bonuses on early loan repayment forecasts can reduce the amount of risk taken by a bank only if bonuses can be "clawed" back in the event of the loan's default. Chapter 2 shows that allowing a court to intervene early in the legislative process can result in more constitutional laws than forcing the court to rule in just judicial review proceedings. And in Chapter 3, the IRS's advance pricing agreement (APA) program is studied. The APA program can reduce a multinational's tax uncertainty by allowing the firm and IRS to agree on future tax payments. Emphasis is placed in this dissertation on how these government policies affect the public's incentives and when these policies are socially beneficial.

Contents

Acknowledgements **i**

Dedication **ii**

Abstract **iii**

List of Tables **viii**

List of Figures **ix**

1 Introduction **1**

2 Reducing Risk-Taking: Dodd-Frank Clawbacks vs. Deferred Compensation **7**

2.1 Introduction and Motivation 7

2.2 A Legislative History of Clawbacks 11

2.3 Related Literature 14

2.4 Basic Model 16

2.4.1 Information 18

2.4.2 Contracts 19

2.4.3 First Best 22

2.4.4 Contractual Constraints 22

2.5 Threshold Incentive Compatible Contracts 26

2.5.1 TIC-Benchmark Contracts 27

2.5.2 TIC-Deferred Compensation Contracts 28

2.5.3 TIC-Clawback Contracts 30

2.6 Optimal TIC Cutoffs z 32

2.6.1 Clawback Limits 35

2.6.2 Low Effort 36

2.6.3 Are Risk-Reducing Contracts Optimal? 40

2.7 Securitization 41

2.8 Conclusion 45

3 Advisory Opinions in the Courts: Overruling Frankfurter's Objection 47

3.1 Introduction 47

3.1.1 Advisory Opinions in the Judiciary 49

3.1.2 Advisory Opinions in State Courts 52

3.1.3 Abstract Judicial Review in Europe 54

3.2 Related Literature 55

3.2.1 Judicial Bias & Fact Discretion 57

3.3 Model 61

3.3.1 Payoffs 64

3.3.2 Fact Discretion 65

3.3.3 Court Payoffs 66

3.3.4 Equilibrium 68

3.3.5 Beliefs 73

3.3.6 Single Reports 77

3.3.7 Comparison of Judicial Systems 78

3.3.8 Model without Fact Discretion 83

3.3.9 Fact Discretion Uncertainty 84

3.4 Ability is Private and Information Imperfect 87

3.4.1 When are Courts Necessary? 93

3.5 Conclusion 95

4 Advance Pricing Agreements and Tax Enforcement 97

4.1 Introduction and Motivation 97

4.2 Related Literature 103

4.3 Model 106

4.3.1 Audit Game Only 106

4.3.2 Advance Pricing Agreements Only 114

4.3.3 Voluntary Advance Pricing Agreements 116

4.4 Alternative APA Settlements: Take-it-or-leave-it Offers 119

4.5 Conclusion 126

5 Bibliography **128**

Appendix A. Proofs from Chapter 2 **138**

Appendix B. Proofs from Chapter 3 **142**

Appendix C. Additional Material from Chapter 4 **152**

List of Tables

3.1 U.S. State Supreme Courts with Advisory Opinions 53

3.2 European Abstract Judicial Review 55

List of Figures

2.1 Timing of the Contracting Game 19

2.2 Distribution of μ with e_L and e_H. 38

2.3 Principal's expected payoff from high and low effort 39

3.1 Timing of the Court's Reporting Protocol 69

3.2 Thresholds used by the court 72

3.3 Social Welfare 82

3.4 $E[V_{12}(r_1, r_2, \theta)]$ as σ^2 increases 86

3.5 $E[V_2(r_2, \theta)]$ as σ^2 increases 87

4.1 Number of APA Applications, 1991-2007 100

4.2 Advance Pricing Agreement Game 107

4.3 Audit Subgame Equilibrium 113

4.4 Comparison of Social Welfare 119

4.5 APA Subgame with Take-it-or-leave-it offers 120

4.6 Firm's APA Choice 124

C.1 Firm's Payoff from Random IRS Auditing 155

Chapter 1

Introduction

When a decision is made can be just as important as the decision itself. If better information arrives over time, a decision maker can benefit by delaying her decision until all relevant information is available. In a frictionless environment, there would be no cost to delay. However, in many interesting contexts, delaying a decision to wait for better information can be costly. This dissertation examines the trade-off between the benefit of waiting for better information and the costs incurred by delay. The chapters below focus on government policies designed to address this trade-off and show when better social outcomes can be achieved if early decisions–even those based on scarce and unreliable information–are allowed. The objective of this dissertation is to characterize how these policies affect the public's incentives and when these policies can lead to improved social welfare.

Chapter 1 studies when to pay bankers to reduce their risk-taking incentives. The chapter compares the efficiency of clawback and deferred compensation contracts in reducing risk-taking. Such an analysis is important since the Dodd-Frank Wall Street Reform and Consumer Protection Act of 2010 mandated all publicly-traded companies implement clawback contracts. Given the government is willing to regulate the types of contracts allowed in firms, I answer the following question: *If the government wants to reduce risk-taking by dictating the type of contract offered by firms, what contract will do the best job?*

The chapter describes the agency problem between a financial firm's management and a banker. In light of the recent financial crisis, the setting I have in mind is a large financial institution that originates and possibly securitizes loans. The banker can work to find a suitable loan candidate for the firm. The more effort put forth by the banker, the greater the chance he finds a suitable candidate. The banker then privately observes the potential loan's risk and decides to originate the loan or not. Using incentive compensation, the firm's management must motivate the correct effort choice and approval decision by the banker.

The clawback and deferred compensation contracts both resolve the agency issues of private information and moral hazard, but at different costs to the principal. Deferred compensation forces all incentive compensation to be delayed until the loan is repaid or defaults. The primary benefit of deferred compensation is that management is able to observe the loan's final outcome before any compensation is paid to the agent. But this better information comes at a cost due to the banker's discounting. Pushing back the date of compensation means management must pay the banker more to maintain the banker's incentives. Clawback contracts offer a way around the inefficiency created when the principal is forced to defer compensation. The clawback contract allows the firm to gain information by observing the loan's outcome before paying the banker, but also avoids the additional costs incurred when compensation is delayed.

When management can reclaim the entirety of the banker's previously paid compensation, the clawback contract is optimal for management and society. The clawback contract is the cheapest contract to implement for the principal and induces less risky loans to be made than deferred compensation (and hence is optimal for society). However, clawbacks are optimal only if management can reclaim a sufficient amount of previously paid compensation. Introducing limits on the amount reclaimed may reduce the efficiency of the clawback contract to the extent that the deferred compensation contract is optimal for society and management. If management is allowed to securitize part of an originated loan, the firm uses securitization as a way to reduce risk–the firm securitizes a larger fraction of the loan when riskier loans are made.

Chapter 2 studies when a court such as the U.S. Supreme Court should rule on the constitutionality of a law. In many U.S. state supreme courts and European constitutional courts, the court can review the law in advisory opinion proceedings immediately before or after the law's enactment. However, the U.S. Supreme Court usually rules on a law's constitutionality only in judicial review proceedings after the law has been implemented. This chapter shows when society can benefit from the court ruling on the constitutionality of a law immediately after a law is enacted.

Society chooses the judicial system that maximizes the likelihood the court makes the correct ruling. The precision of the court's information about the law's constitutionality increases over time. While society prefers the court to make the correct ruling and hence benefits from delayed judgment, the court becomes biased toward a particular ruling as it learns more about the law's consequences. Society then faces the following trade-off: Should the court make an unbiased decision based on imprecise information, or a potentially biased decision based on precise information?

In this chapter, the court's bias is not politically driven, but results from the court's reading of the case's salient details. Such bias is termed "fact discretion" and captures the following notion (see Gennaioli and Schleifer (2008)):

> Such expression need not be conscious or unethical. Judges may unconsciously interpret the evidence, or disregard some inconvenient truths, through the lens of their experiences, beliefs, or ideologies or perhaps even something as mundane as attitudes toward specific litigants or lawyers.

Fact discretion may distort the court's judicial review ruling and lead to decreased social welfare. Advisory opinion rulings eliminate the threat of fact discretion as there are no facts or litigants before the court–the court just has access to the text of the law. A prominent critic of advisory opinions, Felix Frankfurter, argued the lack of information in advisory opinion proceedings meant such opinions would produce sub-optimal outcomes:

> Legislation is largely empirical . . . the history of modern legislation amply

> proves that facts may often be established in support of measures after enactment, although not in existence previously ... In the attitude of court and counsel, in the availability of facts which underlie litigation, there is a wide gulf between opinions in advance of legislation or executive action, and decisions in litigation after such proposals are embodied into law or carried into execution (Frankfurter (1930))

I show that despite the U.S. Supreme Court's reluctance to use advisory opinions, advisory opinions can be the optimal judicial system for society to adopt.

Chapter 3 studies the role of advance pricing agreements (APAs) in reducing corporate tax uncertainty. Due to the ambiguity in transfer pricing standards among countries, it is possible for a multinational and tax authority to disagree about the firm's tax bill despite the firm's best effort to comply with local tax standards. PricewaterhouseCoopers (2012) notes

> Since there is no absolute rule for determining the right transfer price for any kind of international transaction with associated enterprises, whether it involves tangibles, intangibles, services, financing or cost allocation/sharing arrangements, there is huge potential for disagreement as to whether the correct amount of taxable income has been reported in a particular jurisdiction.

To resolve the uncertainty, a recent development in avoiding international tax disputes is the advance pricing agreement (APA) program which allows tax authorities and multinationals to agree on the appropriate transfer pricing methodology in advance. The APA program started in the United States in 1991 and has since expanded to many industrialized and developing countries.

Opposed to adversarial ex-post audits of a multinational's transfer pricing methodology, APAs are cooperative in nature. APAs offer certainty to the firm–by engaging with the IRS before the firm implements its transfer pricing methodology, the firm eliminates the possibility of penalties and unnecessary audit costs. However, APAs can be costly: since the APA is conducted before the firm's transactions take place and is based on preliminary information, the IRS and firm may determine the firm owes a higher tax payment than is actually the case. The firm's trade-off by participating in the APA is potentially paying higher taxes through the APA versus costly ex-post audits.

The main result of this chapter is that the IRS may benefit from reducing its bargaining power if the APA program is voluntary. This is because when the IRS has high bargaining power, it is able to extract large tax concessions from the firm. The firm can avoid these tax payments by forgoing the APA in favor of ex-post auditing. Since it is cheaper for the IRS to complete an APA than audit ex-post, the IRS can maximize social welfare by reducing its negotiating position enough to induce the firm's participation in the APA program. This result explains the widely-held belief that the IRS is lenient on firms that do request APAs.

Chapter 2

Reducing Risk-Taking: Dodd-Frank Clawbacks vs. Deferred Compensation

2.1 Introduction and Motivation

A major criticism of large banks following the 2007 Financial Crisis was that banker incentive compensation led to excessive risk-taking. In response, the Wall Street Reform and Consumer Protection Act of 2010 (henceforth Dodd-Frank) required publicly-traded firms to use clawback provisions–the right to reclaim previously awarded compensation–in executive contracts. However, little public debate took place before Dodd-Frank was enacted about the desirability of clawbacks as a means of reducing risk:

> The clawback provisions that were passed in the final Dodd-Frank bill were only included in previous Senate drafts of the legislation, not in House drafts. There is no mention of clawbacks in Senate hearings on the Dodd-Frank bill. (Sharp (2012))

Given the willingness of the U.S. Government to mandate contractual provisions of public companies, this chapter considers the following question: *If the government wants to reduce risk-taking by dictating the type of contract offered by firms, what contract will do the best job?* Specifically, I compare clawback contracts with another commonly used risk-reducing contract: deferred compensation.

Many firms voluntarily adopted deferred compensation and clawback contracts. For instance, in 2013, Morgan Stanley "deferred the entire annual bonuses of thousands of highly-paid employees" and "...the proportion of bonus-eligible employees' compensation that was deferred rose to 75% in 2011 from 40% in 2009"[1] . Anecdotal evidence suggests that firms are also using clawbacks to reclaim past compensation–Deutsche Bank reclaimed $53.5 million from a trader involved in rigging the Libor rate[2] . And after J.P. Morgan's "London Whale" trading loss in 2012, the amount the firm "...clawed back from each person represent[ed] about two years of total annual compensation."[3]

Banks such as UBS, Goldman, and Credit Suisse have indicated clawback use but

[1] *Deferred Pay Draws Fed's Scrutiny*, Aaron Lucchetti and Michael Rapoport, Wall Street Journal, February 15, 2013.

[2] *'Claw' Claims Trader Bonus*, Laura Stevens, Wall Street Journal, January 25, 2013.

[3] *Whale Clawbacks About Two Years of Compensation*, Dan Fitzpatrick, Wall Street Journal, July 13, 2012.

little is known about the frequency in which clawbacks are enforced at large financial firms. However, in 2005 only 3 percent of Fortune 100 firms had adopted clawback provisions–by 2010 80 percent had[4] .

This chapter describes an agency problem between a financial firm's management and a banker. In light of the recent financial crisis, the setting I have in mind is a large financial institution that originates and possibly securitizes loans. The banker can work to find a low risk loan candidate for the firm. The more effort put forth by the banker, the greater the chance he will find a low risk loan. The banker then privately observes the potential loan's riskiness and decides to approve or reject the loan.

When designing incentives, management faces a trade-off between motivating the banker's effort and ensuring his private information is used to make the desired approval decision. If the banker rejects a potential loan, no information is revealed about the loan's final outcome–it is not possible to learn if the loan would have defaulted or been repaid. So in the event the banker rejects the loan, management will pay a "rejection" wage to the banker that is independent of the loan's outcome. In practice the size of the rejection bonus will be less than the bonus paid in the event of loan approval and repayment, but more than the bonus received if the loan is approved and defaults.

While the rejection wage aligns management's and the banker's incentives in regard to the type of loan to be approved, it has a deleterious effect on the banker's incentive to exert effort. The higher the rejection wage is set, the greater the banker's benefit from

[4] See Dehaan, Hodge, and Shevlin (2011) who cite Equilar (2009) and Equilar (2011)

not exerting effort and uniformly rejecting loans. To ensure effort is exerted to find a low risk loan candidate, the bonus when the loan is approved and repaid must be increased to counteract the incentive to shirk and reject loans. The banker's private information about the riskiness of the loan allows him to collect a "rent" from management.

I consider three contracts: a standard "benchmark" contract, the deferred compensation contract, and the clawback contract. All of these compensation contracts can resolve the agency conflicts of private information and moral hazard, but at different costs to the principal. The benchmark contract pays the banker based on an imprecise repayment forecast before the loan's repayment is observed. Deferring compensation is a commonly used compensation practice–it forces all bonuses to be delayed until the outcome of the loan is realized. The primary benefit of deferred compensation is that management is able to observe the loan's repayment before incentive compensation is paid to the agent. But this information comes at a cost due to banker discounting. Pushing back the date of compensation means management must pay the banker more to account for the banker's discounting. When the banker discounts heavily, the gain from delaying compensation to observe the loan's repayment is outweighed by the additional cost of incentives.

Lastly, clawback contracts allow for transfers between management and banker based on all available information including preliminary loan repayment forecasts and the final repayment or default. This gives the management an option-like contract that allows her to gain more information by waiting until the loan outcome is realized and paying the banker early to avoid extra discounting costs. This leads to bonuses being paid out based on repayment forecasts and being reclaimed by management only if the loan defaults.

The primary difference between the clawback and deferred compensation contract is who holds the compensation until the loan's repayment is observed. With deferred compensation, management holds the money until the loan's outcome is realized. With the clawback contract, the banker is given the money to hold (and put in his bank account and potentially spend) but may be required to return the money at a future date.

2.2 A Legislative History of Clawbacks

Since 2002, the U.S. Government has taken a greater role in regulating private contracts in the financial industry. The government's first intervention was in the wake of corporate accounting scandals in the early twenty first century (i.e. Tyco, Enron, and WorldCom). The Public Accounting Reform and Investor Protection Act (commonly referred to as Sarbanes-Oxley or SOX) did not mandate the adoption of clawback provisions in private contracts, but gave the SEC the power to reclaim past incentive compensation in the event of misconduct (i.e. fraud). Section 304 of SOX states:

> If an issuer is required to prepare an accounting restatement due to the material noncompliance of the issuer, as a result of misconduct, with any

> financial reporting requirement under the securities laws, the chief executive officer and chief financial officer of the issuer shall reimburse the issuer for–(1) any bonus or other incentive-based or equity–based compensation received by that person from the issuer during the 12-month period following the first public issuance. . .

The SEC was the only body with enforcement power for clawbacks–boards of directors, investors, and shareholders did not have the right to reclaim previously paid compensation[5] . Although SOX was passed in 2002, no clawbacks were enforced until 2007 when the former CEO of UnitedHealth Group William McGuire was forced to repay $600 million in incentive compensation. McGuire backdated options to boost his compensation from 1994-2005, making the clawback a result of a SEC investigation (Bowe and White (2007)).

The Emergency Economic Stabilization Act (ESSA) of 2008 was signed into law on October 3, 2008 following the collapse of Lehman Brothers and Bear Stearns. Within ESSA, the Troubled Asset Relief Program (TARP) allowed the U.S. Treasury to purchase distressed assets from banks and other financial institutions. Among the requirements for accepting TARP assistance were executive pay limits including clawbacks. Section 111(B) of ESSA required

> . . . a provision for the recovery by the financial institution of any bonus or

[5] Nothing prevented firms from voluntarily writing clawback contracts. But in the absence of such contracts, only the SEC could legally clawback incentive compensation based on misconduct.

> incentive compensation paid to a senior executive officer based on statements of earnings, gains, or other criteria that are later proven to be materially inaccurate. . .

Unlike SOX, the TARP clawback did not require misconduct to instigate a clawback–no proof was required that the CEO did anything illegal to enforce the clawback.

The focus of this chapter is the Dodd-Frank Act signed into law in July 2010 which made clawbacks in publicly-traded companies mandatory. Section 954 of Dodd-Frank requires that

> . . . in the event that the issuer is required to prepare an accounting restatement due to the material noncompliance of the issuer with any financial reporting requirement under the securities laws, the issuer will recover from any current or former executive officer of the issuer who received incentive-based compensation (including stock options awarded as compensation) during the 3-year period preceding the date on which the issuer is required to prepare an accounting restatement, based on the erroneous data, in excess of what would have been paid to the executive officer under the accounting restatement.

The Dodd-Frank clawback requirement goes further than Sarbanes-Oxley and TARP in mandating the use of clawback contracts for all publicly-traded firms. The penalty for non-compliance is delisting–removal of the company from the stock exchange. The Dodd-Frank clawback is also a no-fault contract–the clawback can be exercised regardless of the reason for an earnings restatement. Troy Paredes, a SEC Commissioner, argued that an executive at a firm with otherwise strong corporate governance "... may nonetheless have to pay back a considerable portion of his or her compensation if the company has to restate because of an accounting error (Paredes (2012))."

2.3 Related Literature

This chapter is related to a number of papers in economics and finance. The two closest papers to mine are Levine and Smith (2010) and Inderst and Pfeil (2011), both of which consider deferred compensation contracts.

Levine and Smith (2010) build an earnings management model in which the agent can increase the likelihood of a high earnings statement by exerting costly effort. The authors study when earnings management will occur and how deferring compensation affects the agent's incentives. Unlike this chapter, Levine and Smith (2010) do not focus on risk-taking incentives or on different risk-reducing contracts. Inderst and Pfeil (2011) build a model that jointly accounts for a firm's compensation practices and securitization decision. I use Inderst and Pfeil (2011) as a starting point for my model, but they do not consider clawbacks.

My model is a multi-task agency problem as the agent must exert effort and use his private information as the principal desires. The ex-ante moral-hazard problem and interim private information is reminiscent of Levitt and Snyder (1997). Levitt and Snyder (1997) show how the possibility of project cancelation can affect the agent's incentives. Papers that allow for risk-shifting by the agent include DeMarzo, Livdan, and Tchistyi (2011), Palomino and Prat (2003), Hellwig (2009), and Biais and Casamatta (1999). In these models, the agent controls via unobserved effort the distribution of output. In Palomino and Prat (2003), the agent has limited liability and can misreport the realized return. The authors show the optimal contract is a bonus contract. Relatedly, Malcomson (2011) provides conditions in which a risk neutral agent's limited liability reduces risk taking by rewarding high-probability events.

This chapter is also related to the literature on CEO incentives. A large literature has developed in part to determine what compensation practices reduce CEO incentives to take risk (see Murphy (1999) for a review of the executive compensation literature). Jensen and Meckling (1976) introduced the agency conflict between equity and debt holders of a leveraged firm. Hellwig (2009) and Biais and Casamatta (1999) study the case when an agent must exert effort and can privately choose among risky projects.

Lastly, this chapter fits into the accounting literature on earnings management. Dutta and Gigler (2002) study the relationship between earnings management and voluntary earnings forecasts. The authors show that voluntary forecasts may prevent the agent from misrepresenting information and that shareholders may not find it optimal to prohibit earnings management. Levine and Smith (2010) similarly show that earnings management may be optimal from the principal's standpoint when the costs of discouraging it are large.

Since mandated and voluntary clawbacks are a recent phenomenon, there is still a paucity of evidence favoring the effectiveness of one compensation system over the other. A few recent empirical studies are related to this chapter. First, using a large sample Chen, Greene, and Owers (2012) find that the incidence of clawback provisions is related to managerial risk aversion, the noisiness of internal accounting information, and inversely related to firm risk. The authors find clawback provisions are correlated with higher CEO compensation and higher CEO pay-performance sensitivity. Addy, Chu, and Yoder (2009) show that voluntarily adopting a clawback contract increases with the number of previous earnings restatements. Babenko, Bennett, Bizjak, and Coles (2012) find that voluntary adoption of clawback provisions are more likely when there has been prior corporate mis-governance at the firm and when earnings management is harder to detect. The authors also find the board's composition matters–more independent boards increase the likelihood a firm voluntarily adopts the clawback contract.

2.4 Basic Model

The model has two players, a risk-neutral principal and agent. I interpret the agent as a loan officer at a bank or financial firm and the principal as the firm's management or board of directors. The model applies more generally to a CEO, banker, or equity trader in charge of potentially risky activities. A loan's repayment is denoted $R \in \{R_h, R_l\}$ where the loan is either repaid in full or defaults and $R_h > R_l$. The probability of loan repayment depends on the loan's type, $\theta \in \{G, B\}$ which is unobserved to both players. A θ loan generates repayment R_h with probability $\gamma_\theta = Pr(R_h|\theta) \in [0, 1]$ and defaults with probability $1 - \gamma_\theta$. $\theta = G$ loans are repaid with a higher probability than $\theta = B$ loans, i.e. $\gamma_G > \gamma_B$. The probability the loan is $\theta = G$ is $\mu \in [0, 1]$ where conditional on μ, the loan is repaid with probability $Pr(R_h|\mu) = \mu\gamma_G + (1 - \mu)\gamma_B$.

The agent can exert high effort e_H or low effort e_L. The agent's cost of exerting high effort is $c_H > 0$ and the cost of low effort is $c_L > 0$ where $\Delta c = c_H - c_L \geq 0$. The agent's effort is unobserved by the principal so it is non-contractible.

The random variable μ is distributed according to distribution F with continuous support on the interval $[0, 1]$. At the lower bound where $\mu = 0$, the probability of loan repayment is γ_B while at $\mu = 1$, the loan is repaid with probability γ_G. If the agent chooses e_H, μ is a realization from the distribution $F_H(\mu)$, and if the agent chooses e_L, μ is realized from distribution $F_L(\mu)$. High effort leads to higher realizations of μ, i.e., $F_H(\mu)$ first order stochastically dominates $F_L(\mu)$, $F_H(\mu) < F_L(\mu)$ for $\mu \in [0, 1]$. After observing μ the agent can approve or reject the loan.

Before proceeding, it is useful to state two assumptions on the p.d.f. f that are used to ensure the principal's problem outlined below has an interior solution. Assumption 1 implies the principal's payoff is concave in μ (and hence a solution exists) and Assumption 2 implies the solution is interior.

Assumption 1 *The density $f_e(\mu)$ is concave for all $\mu \in [0,1]$ and $e \in \{e_H, e_L\}$.*

Assumption 2 *The density is strictly positive $f_e(\mu) > 0$ for all $\mu \in [0,1]$ and $e \in \{e_H, e_L\}$.*

If a loan is approved, the initial capital required to make the loan is $I > 0$. Given θ, a loan's expected repayment is $\tilde{R}_\theta = \gamma_\theta R_h + (1-\gamma_\theta)R_l$. Good loans have higher expected repayments, $\tilde{R}_G > \tilde{R}_B$, and the principal's repayment net investment is $\eta_\theta = \tilde{R}_\theta - I$. Conditional on $\mu \in [0,1]$ the expected repayment net investment is $\eta(\mu) = \mu\eta_G + (1-\mu)\eta_B$. It is assumed a loan has positive expected repayment net investment when μ is sufficiently large:

Assumption 3 $\eta(0) < 0 < \eta(1)$

Assumption 3 implies the principal wishes to reject loans with low μ. If a loan defaults, the principal loses initial capital I and receives a payoff of R_l. In many cases, it is natural to assume that $R_l = 0$. However, R_l could be positive if some loan payments were made before default or negative if additional costs are required in the event of default. I make no assumption on the sign or magnitude of R_l other than $R_h > R_l$.

2.4.1 Information

If the agent approves the loan a noisy signal about the loan's repayment is observed by the principal and agent. The signal is verifiable and contractible. The observed interim signal $s \in \{h, l\}$, is interpreted as a "high" or "low" loan repayment forecast. The interim signal is informative, $q = Pr(h|R_h) = Pr(l|R_l) \in (1/2, 1]$. The probability the public signal is high, $s = h$, given θ is

$$\rho_\theta = Pr(s = h|\theta) = q\gamma_\theta + (1 - q)(1 - \gamma_\theta) \tag{2.1}$$

The timing of the contracting game is shown in Figure 2.1.

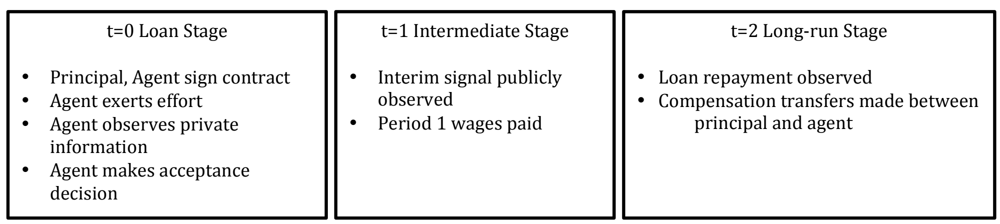

Figure 2.1: Timing of the Contracting Game

The agent discounts $t = 2$ compensation by $\delta \in [0, 1]$, while the principal does not discount. This is natural as the principal can be thought of as a financial firm (with access to financial markets, borrowing, etc.) and the agent is a single financially-constrained individual.

2.4.2 Contracts

The objective of this chapter is to understand how government regulation of compensation contracts (i.e. Dodd-Frank) affect the principal's incentive to originate risky loans. Three contracting "regimes" indexed by i are considered: a benchmark contracting regime $i = b$, deferred compensation contracting regime $i = d$, and clawback contracting regime $i = c$. A contracting regime is a restriction on the nature of payments made between the principal and agent. In this section, the contractual payments available to the principal when incentivizing the agent are defined for each type of contracting regime.

In the benchmark contract, all payments to the agent must take place at $t = 1$ (transfers cannot take place at $t = 2$). This implies all bonuses are paid conditional only on the imprecise repayment forecast $s \in \{h, l\}$. Denote the $t = 1$ bonuses as $w_1(s)$ for $s \in \{h, l\}$. If the agent rejects the loan, he is paid a rejection bonus denoted $\bar{w}$. Let W_b denote the payments allowed in the benchmark contracting regime,

$$W_b = (\bar{w}, w_1(h), w_1(l)) \tag{2.2}$$

The benchmark contract is reminiscent of compensation practices for many firms pre-financial crisis–executives, traders, and bankers were paid bonuses before the outcome of their risky activity was observed by the firm.

In contrast to the benchmark contract, the deferred compensation contract requires the agent's bonuses to be delayed until the outcome of the loan $R \in \{R_h, R_l\}$ is observed at $t = 2$. Denote the bonuses conditional on (R_h, R_l) as $w_2(R_h)$ and $w_2(R_l)$ respectively. The agent receives $\bar{w}$ if he rejects the loan. Let W_d denote the contractual payments available to the principal when the deferred compensation contract is mandated,

$$W_d = (\bar{w}, w_2(R_h), w_2(R_l)) \tag{2.3}$$

The principal could condition bonuses $w_2(R_h), w_2(R_l)$ on the interim signal $s \in \{h, l\}$ in addition to $R \in \{R_h, R_l\}$. However, since payments can only be made at $t = 2$ (payments based on the interim signal are set mandated to be zero, $w_1(s) = 0$ for $s \in \{h, l\}$) and since R is a sufficient statistic of s, the principal does not gain from using s in her deferred compensation contract. Hence, I will assume the deferred compensation contract is independent of the interim signal.

Finally, in the clawback contract transfers are made at $t = 1$ and $t = 2$. At $t = 1$, the principal pays the agent bonuses $w_1(h)$ or $w_1(l)$, but reserves the right to make additional transfers at $t = 2$. In period two, transfer $r(s, R)$ is made between the agent and principal where $r(s, R)$ depends on the history of public signals $s \in \{h, l\}$ and $R \in \{R_h, R_l\}$. I adopt the convention that $r(s, R)$ is negative if it is a transfer from the agent to the principal and otherwise positive. Ex-ante, the principal must specify transfers for all possible histories: $r(h, R_h)$, $r(h, R_l)$, $r(l, R_h)$, and $r(l, R_l)$. If the agent rejects the loan, he is paid $\bar{w}$. The payments available to the principal in a clawback contracting regime are:

$$W_c = (\bar{w}, w_1(h), w_1(l), r(h, R_h), r(h, R_l), r(l, R_h), r(l, R_l)) \tag{2.4}$$

In practice, due to the agent's discounting of $t = 2$ compensation, the principal will want to use "up-front" payments to motivate the agent's effort and approval decision and use period two transfers to reclaim compensation only if the loan defaults[6] .

[6] Below, I restrict attention to the case where $r(s, R) \leq 0$. If $r(s, R)$ could be positive, the clawback would encompass the deferred compensation contract–any outcome achievable via deferred compensation would also be achievable via clawback.

2.4.3 First Best

Before stating the principal-agent contracting problem, first suppose the principal did not face any agency conflict–if effort and μ were publicly observed. The principal could implement her first-best policy and allow the loan to be approved only if its expected return was non-negative,

$$\mu\eta_G + (1-\mu)\eta_B \geq 0$$

The principal's first-best policy would be to approve the loan if $\mu > z^*$ and reject if $\mu < z^*$ where

$$z^* \equiv \frac{-\eta_B}{\eta_G - \eta_B}$$

By Assumption 3, $\eta_B < 0$ and $\eta_G - \eta_B > 0$, so $1 > z^* > 0$.

2.4.4 Contractual Constraints

A contract consists of payments W_i, a prescribed loan acceptance policy $z_i \in [0,1]$, and a prescribed effort level $e \in \{e_H, e_L\}$. The principal's objective is to maximize her expected return on the loan minus compensation paid to the agent. The principal solves this maximization problem in three steps:

1. Given contracting regime i, the principal determines the payments (selected from W_i) that maximize her payoff for each level of effort $e \in \{e_L, e_H\}$ and arbitrary loan acceptance policy $z_i \in [0,1]$.

2. Given the optimal contractual payments for each level of effort and z_i, the principal

solves for the optimal loan approval cutoff z_i^*.

3. The principal selects the effort level that maximizes her payoff given the contracting regime i and optimal cutoff z_i^*.

This section describes the constraints that must be satisfied given contracting regime i when the principal wants the agent to exert high effort and implement the desired approval cutoff z.

Conditional on observing μ and approving a loan, the agent receives expected payoff

$$\mu\pi_G^i + (1-\mu)\pi_B^i$$

where $\pi_G^i > 0$ and $\pi_B^i > 0$ are the agent's expected payoffs from approving $\theta = G$ and $\theta = B$ loans respectively. These payoffs are dependent on the contracting regime and explained in detail in Section 2.5.

By Assumption 3, the principal wishes to avoid a loan if μ is sufficiently low such that $\eta(\mu) < 0$. Since the principal's expected repayment $\eta(\mu)$ is continuous and increasing in μ, by Assumption 3, it is optimal for the principal to use a cut-off $z \in [0, 1]$ such that a loan with $\mu > z$ is accepted and a loan with $\mu < z$ is rejected.

Given arbitrary cutoff z, the principal chooses payments W_i to incentivize the agent to i) accept the loan if $\mu > z$ and ii) reject the loan if $\mu < z$[7] . The first condition requires the payoff from rejecting the loan $\bar{w}$, to be less than the agent's expected payoff

[7] If the agent were allowed to choose the cutoff z, he would set $z = 0$ and approve all loans. This follows from his expected payoff conditional on μ being positive, $\mu\pi_G^i + (1-\mu)\pi_B^i \geq 0$.

from approving the loan when $\mu > z$,

$$\bar{w} \leq \mu\pi_G^i + (1-\mu)\pi_B^i \quad \text{if } \mu > z$$

The second condition requires the payoff from rejecting the loan $\bar{w}$ to be greater than the expected payoff from approving the loan when $\mu < z$,

$$\bar{w} \geq \mu\pi_G^i + (1-\mu)\pi_B^i \quad \text{if } \mu < z$$

By continuity of the agent's expected payoff when approving a loan, the cutoff chosen by the principal z, must satisfy

$$\bar{w} = z\pi_G^i + (1-z)\pi_B^i \tag{2.5}$$

Equation (2.5) implies the principal's cutoff z will depend on the contracting regime $i \in \{b, d, c\}$, but for now I suppress this notation for convenience. At cutoff z, the agent is indifferent between approving and rejecting the loan. Following the previous literature, I assume when the agent is indifferent between approving and rejecting a loan, he does as the principal prefers. However, since F is a continuous distribution, the agent is indifferent between approving and rejecting a loan with probability zero.

The agent's expected payoff from choosing high effort and following the principal's cutoff z is

$$U_H^i(z) = \int_0^z \bar{w} f_H(\mu) d\mu + \int_z^1 \left[\mu\pi_G^i + (1-\mu)\pi_B^i\right] f_H(\mu) d\mu \tag{2.6}$$

The contract requires that if the agent observes $\mu \in [0, z)$, he reject the loan and receive payment $\bar{w}$. The agent's expected payoff from rejecting a loan is given by the first integral. The second integral is the agent's compensation when he approves a loan with $\mu \in (z, 1]$.

If the principal wants to elicit high effort from the agent, incentive compatibility requires the agent's payoff from high effort to be greater than his payoff from low effort, i.e.

$$U_H^i(z) - U_L^i(z) \geq \Delta c \tag{2.7}$$

The last constraint the contract must satisfy is limited liability. The nature of the agent's limited liability constraint will depend on the contracting regime. For benchmark and deferred compensation contracts, limited liability requires all payments to be non-negative,

$$\bar{w}, w_1(h), w_1(l), w_2(R_h), w_2(R_l) \geq 0 \tag{2.8}$$

The clawback limited liability constraint is more complicated and is discussed in detail in Section 2.5.3.

Definition 1 *A Threshold Incentive-Compatible (TIC) Contract is a contract satisfying constraints (2.5), (2.7), and (2.8).*

In other words, a TIC contract is a contract in which the agent is willing to exert high effort and implement the principal's loan acceptance policy $z \in [0, 1]$. Given contracting regime i, the principal's expected payoff from a TIC contract for $z \in [0, 1]$ is

$$\int_z^1 \left[\mu\eta_G + (1-\mu)\eta_B\right] f_H(\mu)d\mu - K_H^i(z) \tag{2.9}$$

The terms in the bracket represent the principal's expected return conditional on a loan being approved with $\mu > z$. If the loan is rejected, the principal receives a benefit of zero. The expected compensation cost in contracting regime $i \in \{b, d, c\}$ when high effort is elicited and cutoff z is selected is denoted $K_H^i(z)$. The principal's problem given contracting regime $i \in \{b, d, c\}$ is to maximize her payoff (2.9) using a TIC contract.

2.5 Threshold Incentive Compatible Contracts

In this section, the TIC contract for each contracting regime $i \in \{b, d, c\}$ is derived. Despite the differences in contracting regimes, a number of results apply to all TIC contracts as a result of the underlying agency problem. Due to the agent's hidden effort, the principal must provide incentives for high effort by setting the bonus paid in the event of a success (i.e. R_h in the deferred compensation contract and $s = h$ if the benchmark contract) sufficiently large. The contract must also motivate the agent to use his private information as the principal desires and reject any loan with $\mu < z$. This implies the wage paid when a loan is rejected $\bar{w}$, must be positive. If $\bar{w} = 0$, the agent would never reject a loan as the payoff from approving a loan is non-negative due to limited liability, $\mu\pi_G^i + (1 - \mu)\pi_B^i \geq 0$.

Two effects occur as the rejection wage $\bar{w}$ increases: First, the agent becomes more likely to reject a loan, and second, effort incentives become muted. The first effect is clear–as the agent's payoff from rejecting a loan increases, the agent is more willing to do so. The second effect results from the agent's trade-off between exerting high effort and rejecting the loan. If $\bar{w}$ is set equal to zero, the agent has an incentive to approve the loan for all $\mu \in [0, 1]$ due to limited liability. But as $\bar{w}$ increases, the agent's payoff from exerting low effort and rejecting the loan increases relative to the payoff from exerting high effort. To restore incentives for high effort, the principal must set $w_1(h)$ higher than she would in the absence of private information.

The rejection wage $\bar{w}$ is increasing in the principal's cutoff z for all contracting regimes $i \in \{b, d, c\}$. As z increases, the agent is asked to reject higher quality loans, so $\bar{w}$ must also increase to maintain the agent's incentives.

2.5.1 TIC-Benchmark Contracts

In the benchmark contract, all payments to the agent are made at $t = 1$ conditional on either the interim signal $s \in \{h, l\}$ or the loan's rejection. The possible contractual payments for principal to use are W_b given by (2.2). The agent's payoff from approving a θ loan is

$$\pi_\theta^b = \rho_\theta w_1(h) + (1 - \rho_\theta) w_1(l)$$

The first term is the agent's payoff when the interim signal is high $s = h$, and the agent is paid $w_1(h)$. The second term is the agent's payoff when a low interim signal is observed and is paid $w_1(l)$. The TIC-benchmark contract for arbitrary cutoff $z \in [0, 1]$ is given by Proposition 2.5.1.

Proposition 2.5.1 *The TIC-benchmark contract for arbitrary cutoff* $z \in [0,1]$ *is*

$$w_1(h) = \frac{\Delta c}{(\rho_G - \rho_B)\left(E_H[\max(\mu, z)] - E_L[\max(\mu, z)]\right)} \tag{2.10}$$

$$\bar{w} = \frac{\Delta c}{E_H[\max(\mu, z)] - E_L[\max(\mu, z)]}\left(z + \frac{\rho_B}{\rho_G - \rho_B}\right) \tag{2.11}$$

and $w_1(l) = 0$. *The principal's expected compensation cost is*

$$K_H^b(z) = \frac{\Delta c}{E_H[\max(\mu, z)] - E_L[\max(\mu, z)]}\left(E_H[\max(\mu, z)] + \frac{\rho_B}{\rho_G - \rho_B}\right) \tag{2.12}$$

Given any possible cutoff $z \in [0,1]$ the TIC-benchmark contract incentivizes the agent to exert high effort and use cutoff z. If interim signal $s = h$ is observed, the agent is paid a bonus $w_1(h)$ which is a function of the cutoff z, the relative probability of observing $s = h$, $\rho_G - \rho_B$, the relative cost of high effort Δc, and the relative likelihood of the loan being approved, $E_H[\max(\mu, z)] - E_L[\max(\mu, z)]$. If interim signal $s = l$ is observed, it is optimal for the principal to set $w_1(l) = 0$. Setting $w_1(l) > 0$ increases the principal's expected compensation cost but does not affect the agent's incentives to exert high effort, so the principal will optimally choose $w_1(l) = 0$.

2.5.2 TIC-Deferred Compensation Contracts

If the principal is required to use a deferred compensation contract to incentivize the agent, the principal is limited to using payments W_d given by (2.3). The agent's expected payoff from accepting a θ loan with deferred compensation is

$$\pi_\theta^d = \delta\left[\gamma_\theta w_2(R_h) + (1 - \gamma_\theta) w_2(R_l)\right] \tag{2.13}$$

The first term inside the bracket is the agent's payoff when the loan is repaid $w_2(R_h)$, and the second term is the agent's payoff when the loan defaults, $w_2(R_l)$. Delaying bonuses $w_2(R_h)$ and $w_2(R_l)$ until $t = 2$ ensures the agent is paid conditional on the best information available, but also forces the principal to incur additional wage costs due to the agent's discounting $\delta \in [0, 1]$ of $t = 2$ payments. The TIC-deferred compensation contract for arbitrary cutoff $z \in [0, 1]$ is given by Proposition 2.5.2.

Proposition 2.5.2 *The TIC-deferred compensation contract for arbitrary cutoff* $z \in [0, 1]$ *is*

$$w_2(R_h) = \frac{\Delta c}{\delta(\gamma_G - \gamma_B)\left(E_H[\max(\mu, z)] - E_L[\max(\mu, z)]\right)} \tag{2.14}$$

$$\bar{w} = \frac{\Delta c}{E_H[\max(\mu, z)] - E_L[\max(\mu, z)]}\left(z + \frac{\gamma_B}{\gamma_G - \gamma_B}\right) \tag{2.15}$$

and $w_2(R_l) = 0$*. The principal's expected compensation cost is*

$$K_H^d(z) = \frac{\Delta c}{\delta\left(E_H[\max(\mu, z)] - E_L[\max(\mu, z)]\right)}\left(E_H[\max(\mu, z)] + \frac{\gamma_B}{\gamma_G - \gamma_B}\right) \tag{2.16}$$

The bonus $w_2(R_h)$ is increasing in the relative cost of effort Δc, and decreasing in the agent's discount factor δ, the relative probability of loan repayment $\gamma_G - \gamma_B$, and the relative likelihood of the loan being approved, $E_H[\max(\mu, z)] - E_L[\max(\mu, z)]$. When the agent's discount factor decreases, the principal must set $w_2(R_h)$ higher to maintain incentives. This is intuitive in the deferred compensation setting since the principal is forced to use delayed bonuses to elicit effort and the desired cutoff rule. When the agent does not value $t = 2$ compensation, as $\delta \to 0$, it is infinitely expensive to motivate the agent, $w_2(R_h) \to \infty$. To avoid the additional complexity of assuming the principal has bounded wealth, I assume the principal has "deep pockets" and can credibly pay the agent for all parameter values.

2.5.3 TIC-Clawback Contracts

If the principal is mandated to use the clawback contract, she has a variety of contractual payments available to incentivize the agent. The agent's payoff from approving a θ type loan is

$$\pi_\theta^c = \rho_\theta w_1(h) + (1-\rho_\theta)w_1(l) + \delta\left[(1-q)(1-\gamma_\theta)r(h, R_l)\right] \tag{2.17}$$

Since the agent discounts $t = 2$ payments, the principal will pay the agent at $t = 1$ and reclaim compensation only if the loan defaults (it is never optimal for the principal to delay payments to the agent). It is shown below that all other payments are optimally set to zero so they are excluded for exposition.

The clawback contract requires a different limited liability constraint than the one used in the benchmark or deferred compensation contracting problems. This is because the principal can reclaim compensation already paid to the agent in the clawback contract. Hence, in the clawback setting, the total compensation paid to (and reclaimed from) the agent over both periods must be non-negative,

$$w_1(s) + r(s, R) \geq 0 \tag{2.18}$$

for all $s \in \{h, l\}$ and $R \in \{R_h, R_l\}$. Intuitively, this condition implies the principal cannot reclaim more bonus compensation than was initially paid out in the form of $w_1(s)$ bonuses. In practice $r(s,R) < 0$ only when the agent is paid at $w_1(s)$ based on a high interim signal and the loan defaults, i.e. $r(h,R_l) < 0$. All other transfers are set equal to zero, $r(h,R_h) = r(l,R_h) = r(l,R_l) = 0$. To see this, note if $s = l$ is observed, the principal pays the agent $w_1(l) = 0$. Then by the limited liability constraint the principal cannot reclaim positive compensation so $r(l,R_h) = r(l,R_l) = 0$. And since the agent discounts, the principal prefers to pay the agent at $t = 1$, so $r(h,R_h) = 0$, as there is no gain from delaying any payment to the agent. In the contract, the limited liability constraint will bind when history (h,R_l) is observed, i.e., the principal reclaims the entire compensation if the loan defaults, $r(h,R_l) = -w_1(h)$. The TIC-clawback contract for arbitrary cutoff $z \in [0,1]$ is given by Proposition 2.5.3.

Proposition 2.5.3 *The TIC-clawback contract given arbitrary cutoff $z \in [0,1]$ is*

$$w_1(h) = \frac{\Delta c}{[(\rho_G - \rho_B) + \delta(\gamma_G - \gamma_B)]\,(E_H[\max(\mu,z)] - E_L[\max(\mu,z)])} \tag{2.19}$$

$$\bar{w} = \frac{\Delta c}{E_H[\max(\mu,z)] - E_L[\max(\mu,z)]}\left(z + \frac{\rho_B - \delta(1-q)(1-\gamma_B)}{\rho_G - \rho_B + \delta(\gamma_G - \gamma_B)}\right) \tag{2.20}$$

$$r(h,R_l) = -w_1(h) \tag{2.21}$$

and $w_1(l) = r(l,R_l) = r(l,R_h) = r(h,R_h) = 0$. The principal's expected compensation cost is

$$K_H^c(z) = \Delta c\left(\frac{E_H[\max(\mu,z)] + \frac{\rho_B - \delta(1-q)(1-\gamma_B)}{\rho_G - \rho_B + \delta(\gamma_G - \gamma_B)}}{E_H[\max(\mu,z)] - E_L[\max(\mu,z)]}\right) \tag{2.22}$$

2.6 Optimal TIC Cutoffs z

Propositions 2.5.1, 2.5.2, and 2.5.3 characterize the TIC contracts for each contracting regime $i \in \{b, d, c\}$ when the principal uses arbitrary cutoff z. Using Propositions 2.5.1, 2.5.2, and 2.5.3, in the second stage of solving for the principal's optimal contract, the principal chooses cutoff $z_i \in [0, 1]$ to maximize her payoff

$$\int_{z_i}^{1} \left[\mu\eta_G + (1-\mu)\eta_B\right] f_H(\mu) d\mu - K_H^i(z_i) \tag{2.23}$$

where $z_i \in [0, 1]$ is the cutoff chosen under contracting regime i and where $K_H^i(z_i)$ is given by (2.12), (2.16), and (2.22) for $i \in \{b, d, c\}$. By Assumptions 1-3, the principal's problem has an interior solution $z_i^* \in (0, 1)$ for each contracting regime that satisfies

$$-\left[z_i^*\eta_G + (1-z_i^*)\eta_B\right] f_H(z_i^*) = K_H'(z_i^*) \tag{2.24}$$

This condition requires the principal's marginal expected return z_i^* to equal the marginal expected compensation cost at cutoff z_i^*.

Proposition 2.6.1 *The optimal TIC contract's loan acceptance policy z_i^* is strictly less than the first-best, i.e. $z_i^* \in (0, z^*)$.*

For all contracting regimes, the principal will optimally choose a cutoff z_i^* strictly less than her first-best policy z^*. Thus, the principal allows riskier loans to be approved, even if risk-reducing contracts such as deferred compensation and clawbacks are mandated.

To sketch the proof of this result, note this proposition follows from the principal's expected compensation costs being dependent on z. This fact and the next lemma show that for any contract, the principal's optimal cutoff must be less than her first-best policy.

Lemma 2.6.1 *The principal's expected compensation cost* $K_H^i(z)$ *(and marginal cost* $K_H'(z)$*) are increasing in* z*.*

This result implies the right hand side of (2.24) is positive. However, the left hand side is positive only if $z < z^*$ since the first-best cutoff rule z^* guarantees the bracketed terms on the left hand side to be zero. For any $z > z^*$, the left hand side of (2.24) is negative implying no $z > z^*$ is a solution. Hence the cutoff that maximizes the principal's payoff must be smaller than the first-best $z < z^*$.

The main result shows how the contracting regime affects the principal's optimal cutoff z_i^* for each contract $i \in \{b, d, c\}$.

Theorem 2.6.1 *Let Assumptions 1-3 hold. Then there exists* $\hat{\delta}$ *such that*

1. *for* $\delta \in (0, \hat{\delta})$,

$$0 < z_d^* < z_b^* < z_c^* < z^*$$

2. *for* $\delta \in (\hat{\delta}, 1)$,

$$0 < z_b^* < z_d^* < z_c^* < z^*$$

where $\hat{\delta}$ *satisfies*

$$\hat{\delta} \equiv \frac{E_H[\max(\mu, z_b^*)] - E_L[\max(\mu, z_b^*)]}{E_H[\max(\mu, z_d^*)] - E_L[\max(\mu, z_d^*)]} \left(\frac{E_H[\max(\mu, z_d^*)] + \frac{\gamma_B}{\gamma_G - \gamma_B}}{E_H[\max(\mu, z_b^*)] + \frac{\rho_B}{\rho_G - \rho_B}} \right)$$

Part one of Theorem 2.6.1 shows that when the agent discounts heavily $\delta < \hat{\delta}$, deferred compensation results in the most risk being taken by the principal. When $\delta = 0$, $K_H^c = K_H^b$ for all z, the expected compensation of the TIC-clawback contract is identical to the cost of the TIC-benchmark contract and hence the principal's optimal cutoff is identical $z_c^* = z_b^*$. But for all $\delta \in (0, \hat{\delta})$ clawbacks are the cheapest contract for the principal, $z_c^* > z_b^* > z_d^*$.

The ordering of z_i^* changes when the agent does not heavily discount $t = 2$ compensation. If the agent does not discount $\delta = 1$, then $K_H^c = K_H^d$ for all z, so the expected cost of the TIC-clawback contract is equal to the cost of the TIC-deferred compensation contract. In this case, both contracts perform equally well in reducing risk taking. But for $\delta \in (\hat{\delta}, 1)$, the clawback contract induces less risk taking than either the deferred compensation or benchmark contract, $z_c^* > z_d^* > z_b^*$.

Theorem 2.6.1 shows that TIC-clawback contracts are advantageous to the principal and society. The principal's payments to the agent are minimized when using the clawback due to the option to reclaim past compensation if the loan defaults. In addition, the principal can avoid the agent's discounting by incentivizing effort and the optimal cutoff z_i^* using $t = 1$ payments. While allowing more risky loans than the first-best policy (i.e. $z_i^* < z^*$), society benefits the most from clawbacks as the contract induces the least amount of risk, $z_c^* \geq \max\{z_b^*, z_d^*\}$. The next section shows that the clawback contract's optimality changes when limits on the amount of reclaimed compensation are introduced.

2.6.1 Clawback Limits

In practice, there are often limits on how much compensation can be reclaimed by the principal in the event of a clawback. The Dodd-Frank clawback allows the firm to reclaim incentive compensation from the previous three years. In this section, the role of clawback limits are explored. Restricting the amount of money that can be reclaimed increases the principal's expected compensation costs with the clawback contract. Denote $\kappa \in [0,1]$ as the fraction of previously paid compensation that the principal can reclaim (note κ is not a "hard" limit, but a fraction of $t = 1$ compensation). When $\kappa = 1$, the principal can reclaim the entirety of $w_1(s)$, and so Theorem 2.6.1 holds. When $\kappa = 0$, the principal is prohibited from reclaiming any compensation and the TIC-clawback contract is identical to the TIC-benchmark contract. The TIC-clawback contract with clawback limit $\kappa \in [0,1]$ is given by Proposition 2.6.2.

Proposition 2.6.2 *The clawback contract with clawback limit $\kappa \in [0,1]$ that elicits high effort conditional on cutoff z is*

$$w_1(h) = \frac{\Delta c}{\left[(\rho_G - \rho_B) + \kappa\delta(\gamma_G - \gamma_B)\right]\left(E_H[\max(\mu, z)] - E_L[\max(\mu, z)]\right)} \tag{2.25}$$

$$\bar{w} = \frac{\Delta c}{E_H[\max(\mu, z)] - E_L[\max(\mu, z)]}\left(z + \frac{\rho_B - \kappa\delta(1-q)(1-\gamma_B)}{\rho_G - \rho_B + \kappa\delta(\gamma_G - \gamma_B)}\right) \tag{2.26}$$

and $r(h, R_l) = -\kappa w_1(h)$, and $w_1(l) = 0$. The principal's expected compensation cost is

$$K_H^c(z) = \Delta c\left(\frac{E_H[\max(\mu, z)] + \frac{\rho_B - \kappa\delta(1-q)(1-\gamma_B)}{\rho_G - \rho_B + \kappa\delta(\gamma_G - \gamma_B)}}{E_H[\max(\mu, z)] - E_L[\max(\mu, z)]}\right) \tag{2.27}$$

Proposition 2.6.2 is a generalization of Proposition 2.5.3, if $\kappa = 1$, the two are identical. The rejection wage $\bar{w}$, bonus $w_1(h)$, and expected compensation K_H^c are all decreasing in κ. This implies that as the principal becomes more limited in reclaiming $w_1(h)$, as $\kappa \to 0$, her compensation cost K_H^c increases. The next result shows that in the presence of clawback limits, clawbacks may no longer induce the least risky loans.

Theorem 2.6.2 *Let Assumptions 1-3 hold. For $\delta \in [0,1]$, there exists $\kappa^*(\delta) \in [0,1]$ such that if*

1. *$\kappa > \kappa^*(\delta)$, $z_c^* \geq z_d^*$ and*
2. *$\kappa < \kappa^*(\delta)$, $z_d^* \geq z_c^*$.*

Thus far, it has been assumed the clawback limit is set by a third party such as the U.S. Congress or a regulator. If the principal could choose κ, she would optimally set $\kappa = 1$ since K_H^c is decreasing in κ. An immediate implication of Theorem 2.6.2 is that if clawbacks are mandated, the principal should be allowed to reclaim all previously awarded compensation. Restricting the principal's ability to implement the clawback induces the principal to allow more risk taking.

2.6.2 Low Effort

When the principal uses a TIC-clawback contract there is a trade-off in choosing cutoff z: the higher z is set, the more likely the loan is repaid but the larger the expected compensation costs. Since compensation costs can outweigh the benefit from originating the loan, an alternative for the principal is to have the agent exert low effort. The expected compensation cost when low effort is elicited is independent of the chosen loan cutoff z. This is because the agent is paid a flat wage $w = c_L$ to incentivize $e = e_L$. When paid $w = c_L$, the agent still follows the principal's rejection policy independently of contracting regime $i \in \{b, c, d\}$. The principal's expected compensation cost when low effort is elicited is c_L.

Lemma 2.6.2 *If the principal elicits low effort, the agent is paid a flat wage $w = c_L$ and the principal chooses first-best cutoff z^* independently of contracting regime $i \in \{b, d, c\}$.*

The principal's optimal cutoff with low effort is the first-best cutoff, $z^* = \frac{-\eta_B}{\eta_G - \eta_B}$. Since expected compensation $w = c_L$ is independent of the principal's chosen cutoff z, the marginal cost of setting $z = z^*$ is zero. Hence the principal's first order condition with respect to z is the same as in the first-best case.

However, this does not imply the principal will always elicit low effort. To see when eliciting low effort is optimal, it is convenient to specify distributions F_L and F_H in the following manner. Let $F_L(\mu) = \mu$ be the uniform distribution with density $f_L(\mu) = 1$ and let $F_H(\mu) = \mu^a$ for $a > 1$ with density $f_H(\mu) = a\mu^{a-1}$.

Lemma 2.6.3 *Let $F_L(\mu) = \mu$ and $F_H(\mu) = \mu^a$ for $\mu \in [0, 1]$ and $a > 1$. There exists $a_i^* \in (1, \infty)$ such that for any $a > a_i^*$, the optimal TIC contract is used, and for any $a < a_i^*$, the low effort contract is used.*

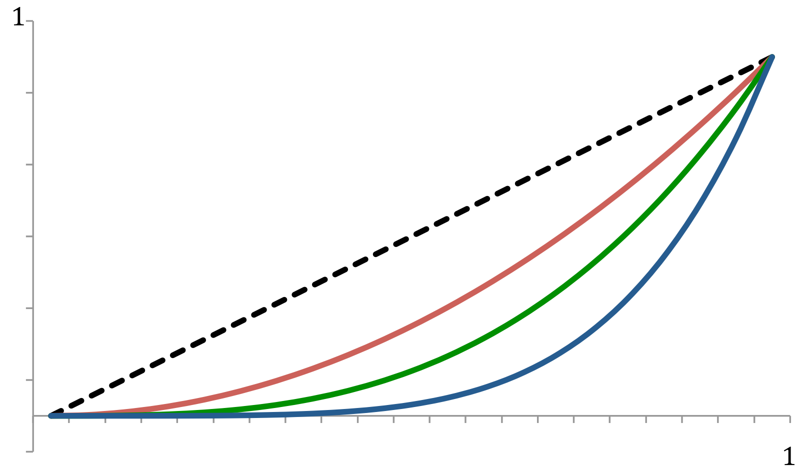

Figure 2.2: The diagonal dotted line is $F_L(\mu) = \mu$. The distribution $F_H(\mu) = \mu^a$ is graphed for (from highest to lowest) $a = 2$, $a = 3$, and $a = 5$.

For a specified cutoff z, the probability $\mu < z$ shrinks to zero as $a \to \infty$. Hence, high effort becomes more important to the principal as a increases. When a is close to one, the probability $\mu < z$, $F_H(z)$, is approximately $F_L(z)$, so the realizations of μ from high and low effort are approximately the same: the gain from eliciting e_H is outweighed by the additional costs of eliciting high effort. Another implication of Lemma 2.6.3 is that the principal's optimal cutoff z_i^* and high effort may be complements instead of substitutes. When effort is more important, as a grows large, it is also optimal for the principal to set a higher cutoff since the probability of $\mu < z$ shrinks to zero.

Figure 2.3 shows when eliciting high effort is optimal from the principal's perspective. What is important to the principal is not the shape of her payoff from eliciting e_H for all z, but her payoff from the optimal cutoff z_i^*. It is clear that when a is sufficiently

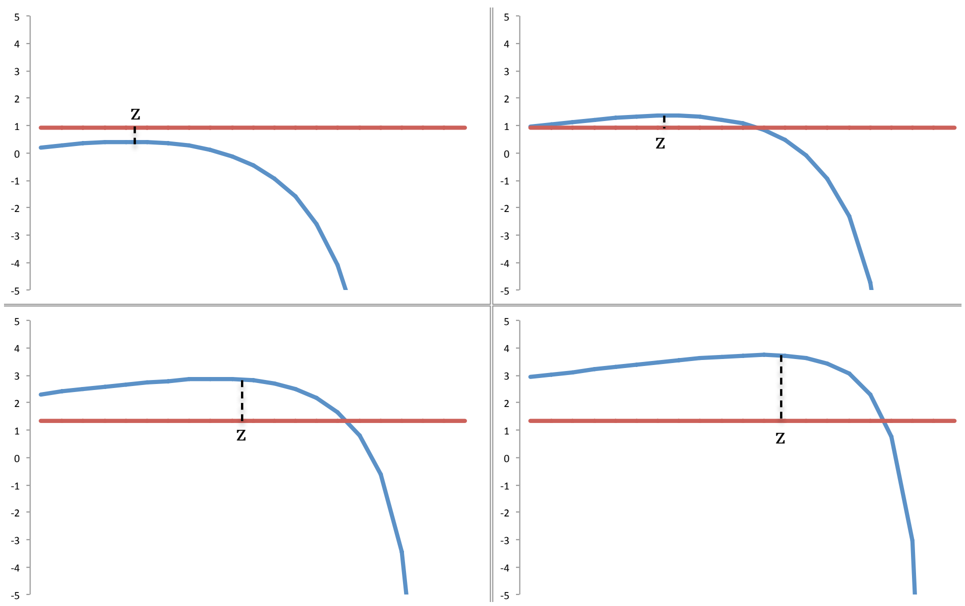

Figure 2.3: The horizontal line is the principal's expected payoff from eliciting e_L and using first-best cutoff z^*. The principal's payoff from eliciting e_H is shown when $a = 2$ (upper left), $a = 3$ (upper right), $a = 4$ (lower left), and $a = 8$ (lower right) for the benchmark contract.

close to one, the low effort contract is optimal. But when a is large, eliciting e_H can be optimal for the principal.

2.6.3 Are Risk-Reducing Contracts Optimal?

In this section, the principal's optimal contract is derived if she were allowed to choose contracting regime $i \in \{b, d, c\}$, cutoff cutoff z, effort level, and clawback limit κ. This analysis is relevant as many firms voluntarily adopted clawback and deferred compensation contracts without government intervention (see the Section 4.2 for a brief overview of empirical work on voluntarily adoption of clawback contracts). The following lemma derives the principal's optimal contract.

Lemma 2.6.4 *Let $F_L(\mu) = \mu$ and $F_H(\mu) = \mu^a$ for $\mu \in [0, 1]$ and $a > 1$. Suppose the principal could choose contract type $i \in \{b, d, c\}$. Then,*

1. *if $a < a_c^*$, low effort is elicited and the first-best cutoff z^* is chosen, and*

2. *if $a > a_c^*$, the optimal TIC-clawback contract is chosen.*

Low effort is elicited only when high effort does not sufficiently increase the expected realization of μ. By Lemma 2.6.3, the principal's expected payoff from eliciting high effort is greater than low effort only if $a > a_i^*$. If low effort is elicited, the principal's cutoff is the first-best z^* independently of the principal's chosen contracting regime $i \in \{b, d, c\}$.

For all $\delta \in (0, 1)$, by Theorem 2.6.1 the TIC-clawback contract is the optimal contract for the principal. Since $K_H^c(z)$ is decreasing in κ, the principal will always choose the right to reclaim the entire $t = 1$ bonus, i.e. $r(h, R_l) = -w_1(h)$.

This result shows that the clawback contract is the optimal contract when the principal can reclaim all previously paid incentive compensation. This result can also explain the acceptance of clawback contracts prior to the passage of Dodd-Frank.

2.7 Securitization

In this section, the principal is given the opportunity to securitize a portion of an originated loan. I show how the contracting regime $i \in \{b, d, c\}$ and principal's optimal cutoff z_i^* affect the principal's securitization decision. Securitization is the practice of selling assets comprised of originated loans to investors. The practice increased in popularity prior to 2007 and has been blamed in part for the 2007 Financial Crisis. The "originate-to-distribute" business model was meant as a way for banks to reduce the risk of their loan portfolio: by originating, pooling, and selling loans, banks could disperse risk to outside investors. In this section, I abstract from many important details of the securitization process (i.e. pooling, tranches, etc.) and focus on the relationship between the level of securitization, the principal's risk cutoff z, and the contracting regime $i \in \{b, d, c\}$.

The principal chooses her securitization fraction $\zeta_i \in [0, 1]$ conditional on the contract offered to the agent (i.e. either a TIC contract or low effort contract). In this way, the characterization of contracts from previous sections still apply and the only decision for the principal is how much of the loan to securitize given her optimal cutoff.

The market is comprised of investors who exogenously demand the firm's securitized loan. Since the focus of this section is on the principal's supply decision and not on characterizing the market's demand, the principal's per-unit benefit from securitization is denoted $\Gamma > 0$. Inderst and Pfeil (2011) suggest Γ results from the market being more willing to pay for the asset than the principal, but any profit-generating motive would suffice.

The market does not observe the principal's risk cutoff z_i, but does observe the contracting regime dictated by the government and the fraction of the loan securitized ζ[8] . Conditional on observing ζ_i, the market's forms belief $\hat{z}_i$ about the principal's cutoff strategy.

Let $R_l = 0$ for simplicity. The principal's expected return conditional on μ is $\tilde{R}(\mu) = [\mu\gamma_G + (1-\mu)\gamma_B]R_h$ and her net return is $\eta(\mu) = [\mu\gamma_G + (1-\mu)\gamma_B]R_h - I$. Investors pay fair value for the securitized loan so the price of the loan conditional on the market's belief $\hat{z}_i$ is

$$p(\hat{z}_i) = (1+\Gamma)E[\tilde{R}(\mu)|\mu > \hat{z}_i]$$

The price paid by investors for the loan is increasing in the principal's cutoff $\hat{z}_i$. This is intuitive as a higher cutoff leads to better loans being made on average, which increases

[8] Dodd-Frank requires firms to publish their compensation scheme annually. However, it can be difficult to discern the exact nature of the contract due to vague language.

the price the market is willing to pay.

The principal's payoff from choosing cutoff z_i and fraction to securitize ζ_i is

$$\int_{z_i}^{1} [\eta(\mu) - \tilde{R}(\mu)\zeta_i + p(\mu)\zeta_i] dF(\mu) - K^i(z_i)$$

When the principal securitizes ζ_i of the loan, her payoff from loan repayment $\eta(\mu)$ is decreased by the fraction sold to the market $\tilde{R}(\mu)\zeta_i$. However, the principal receives $p(\mu)\zeta_i$ by selling ζ_i of the loan. The fraction securitized ζ_i only affects expected compensation $K^i(z)$ through the principal's choice of z_i.

Definition 2 *A rational expectations equilibrium consists of the principal's securitization strategy ζ_i^* and the market's belief $\hat{z}_i(\zeta_i^*)$ such that i) the principal's strategy is optimal given the market's belief $\hat{z}_i$, and ii) the market's belief is consistent with the principal's strategy.*

This definition of equilibrium requires that the market, after observing the fraction of the loan securitized ζ_i, form the correct belief about the principal's cutoff strategy $\hat{z}_i = z_i$.

Adding the securitization decision makes the model more complex, so specifying the distributions F_L, F_H simplifies the analysis. If low effort is elicited, the distribution of μ is the uniform distribution $F_L(\mu) = \mu$ with density $f_L(\mu) = 1$ for $\mu \in [0, 1]$. If high effort is elicited, let $F_H(\mu) = \mu^a$, $a > 1$ with density $f_H(\mu) = a\mu^{a-1}$ for $\mu \in [0, 1]$.

Proposition 2.7.1 *The principal's optimal securitization fraction ζ_i^* is*

1. $\zeta_c^* \leq \zeta_d^*$ *if* $\kappa > \kappa^*(\delta)$

2. $\zeta_c^* \geq \zeta_d^*$ *if* $\kappa < \kappa^*(\delta)$

for $\delta \in [0, 1]$*. If low effort is elicited, the first-best cutoff* z^* *cannot be achieved if* $\zeta^* > 0$*.*

Since ζ_i^* is decreasing in z_i^*, the principal uses securitization as insurance against risky loans. Part one of the proposition shows that the TIC-clawback contract will induce less securitization than the TIC-deferred compensation contract $\zeta_c^* \leq \zeta_d^*$ if the clawback limit is sufficiently large $\kappa > \kappa^*(\delta)$. This follows from part one of Theorem 2.6.2 that shows if $\kappa > \kappa^*(\delta)$, $z_c^* \geq z_d^*$ which implies $\zeta_c^* \leq \zeta_d^*$. Securitizing a low fraction of the loan signals to the market the principal is using a high cutoff z_i^*.

If $\kappa < \kappa^*(\delta)$, by Theorem 2.6.2 $z_c^* \leq z_d^*$. Since securitization is decreasing in z, it must be $\zeta_c > \zeta_d$, the principal securitizes a greater fraction of the loan with a TIC-clawback contract than a TIC-deferred compensation contract.

This proposition highlights the trade-off the principal faces when cutoff z_i^* increases. The larger z_i^* is, the less risk the principal takes on and the higher the loan's expected return. This has two effects: first, the principal receives a higher price for the securitized loan $p(z_i)$, as investors are willing to pay more for higher quality loans. Second, the principal's payoff from keeping the loan also increases. While it is not optimal for the principal to securitize the entire loan or none of the loan, the principal prefers to retain a larger fraction of the loan (ζ_i decreases) as z_i increases.

Suppose the principal can commit to a level of securitization before signing a contract with the agent. In the presence of securitization, the next lemma shows the principal will take more risk.

Lemma 2.7.1 *Let z_i^s be the cutoff used for contracting regime $i \in \{b, d, c\}$ when securitization occurs ($\zeta_i > 0$). Then for all contracts, $z_i^s < z_i^*$.*

This lemma summarizes many of the post-financial crisis criticisms of securitization. The ability to securitize a fraction of the loan allows the principal to off-load the risk from loan defaults onto the financial market. Hence when securitization is possible, the principal is willing to originate more risky loans regardless of contracting regime.

2.8 Conclusion

Motivated by the recent financial crisis and Dodd-Frank legislation, this chapter provides a first attempt at understanding how clawback contracts affect a financial firm's incentive to originate risky loans. When designing incentives, the principal faces a trade-off between motivating the agent's effort and ensuring his private information is used to make the desired approval decision. Conditions are provided in which clawback contracts are optimal for the principal and society. It is shown that limits on the amount reclaimed by the principal only reduces the efficiency of the clawback contract and makes the deferred compensation contract cheaper to the principal and more desirable to society. If the principal is allowed to securitize a fraction of the loan, she will induce the agent to approve more risky loans.

There are a few interesting directions for future research. The first is to understand the role of clawback contracts in the financial labor market. Recent papers along this line include Glode and Lowery (2012), Bond and Axelson (2012), Bond and Glode (2012), and Acharya, Pagano, and Volpin (2012) (who study how competition among firms drives up the compensation of traders). A second line of research worth pursuing is agent risk aversion. Since the clawback contract makes the agent's compensation more risky, risk aversion should play an important role in the optimal contract. Third, the current model is static–the agent makes his effort and approval decision only once. One could imagine a dynamic model in which the agent must manage the firm's riskiness over time. In this case, the contract could be considerably more complicated as information about the loan's return trickles in over time. Papers along this line include Kartik, Halac, and Liu (2012) as well as Biais, Mariotti, Rochet, and Villeneuve (2010) and Myerson (2010).

Chapter 3

Advisory Opinions in the Courts: Overruling Frankfurter's Objection

3.1 Introduction

In many jurisdictions, courts have the right to review executive decisions and legislation. Perhaps the most famous example of this is the U.S. Supreme Court's use of judicial review to rule ex-post on previously enacted laws. But judicial review is not the only point in time in which a court can rule on a law. Certain U.S. states allow courts to judge the constitutionality of a law in an *advisory opinion* ruling before the law is enacted. And in Europe, many constitutional courts use *abstract judicial review* to rule on a law immediately after enactment.

In this chapter, I consider society's problem of choosing a judicial system that maximizes the likelihood the correct ruling is made by the court. The precision of the court's information about the constitutionality of a law increases over time. While society prefers the court make the correct ruling, the court becomes biased toward a particular ruling as it learns more about the law and the law's consequences. Society then faces the trade-off: Should the court make an unbiased decision based on imprecise information, or a biased decision based on precise information? I allow society to choose one of the following observed judicial systems:

1. *Advisory Opinion*: The court makes one binding ruling *before* the law is implemented[1] .

2. *Judicial Review*: The court makes one binding ruling *after* the law is implemented.

3. *Advisory Opinion and Judicial Review*: The court makes a non-binding ruling before the law is enacted and then makes a binding ruling after the law is implemented.

A contribution of this chapter is to apply the notion of "fact discretion" to a principal-expert model. Here the court can choose its interpretation of a law only after it has

[1] In my model advisory opinions and abstract judicial review are both characterized by a lack of concrete information. Advisory opinions should be considered as advice that the legislature may take into account before passing a bill. How advisory opinions and the legislature interact is the subject of Rogers & Vanberg (2002). Abstract judicial review decisions in European courts are almost always binding (once a court rules a law unconstitutional, the court will not see the same law again).

received precise information about how the law was implemented. The court's bias is not politically driven, but is a result of the court's reading of the case. My definition of fact discretion is taken from Gennaioli and Schleifer (2008):

> Such expression need not be conscious or unethical. Judges may unconsciously interpret the evidence, or disregard some inconvenient truths, through the lens of their experiences, beliefs, or ideologies or perhaps even something as mundane as attitudes toward specific litigants or lawyers.

For this chapter I focus on the court's attitude toward specific litigants and on its view of the most salient details of the case–something the court cannot predict before observing the case in judicial review. Because the court does not know who will be litigating and the details the case will depend on, the court cannot form beliefs about future fact discretion before judicial review.

Below, I outline the costs and benefits of advisory opinions as well as their use in judicial systems around the world. Section 2 contains a defense of my main assumption–that the court is biased due to its fact discretion. Section 3 is the main model and Section 4 contains an extension to imperfect information and private ability.

3.1.1 Advisory Opinions in the Judiciary

The U.S. Supreme Court refused to provide advisory opinions in 1793–ten years before it adopted the practice of judicial review in *Marbury v. Madison*. The motivation for the Supreme Court's refusal to give advisory opinions has been subject to debate by legal scholars, but the Court's decision has had a lasting effect on the relationship between the Judicial and other branches of government[2] . The Supreme Court's reluctance to give advisory opinions is more notable when compared to U.S. state supreme courts and international judicial systems. Many European courts and U.S. state supreme courts provide advice on laws, intra-government questions, and other legal cases. These rulings made by a court are not based on concrete facts or details, but on a preliminary reading of laws or statutes. In contrast, the Supreme Court and state supreme courts have adopted the use of judicial review as an acceptable means of reviewing laws.

Legal scholars have long debated the advantages and drawbacks of advisory opinions. Stevens (1959) outlined the major reasons why courts have been reluctant to give advisory opinions. These reasons can be summarized in three arguments:

1. *Reputational Concerns:* Although advisory opinions are not technically binding, the court and others are likely to treat them as such (and subordinate courts are likely to treat them as precedent),

2. *Difference in Information:* There is not enough information to rule on in advisory proceedings, and

3. *Choice of Judicial System:* There is nothing that forces the courts to do so. (Some states specifically mention the right of the legislature or executive to ask for opinions, but many do not and the practice was up to the discretion of the court.)

[2] One argument against advisory opinions is that it gives the Executive Branch too much power over the Judicial Branch and is against the system of "Checks and Balances" envisioned by the Founding Fathers.

Advisory opinions avoid social costs that arise when statutes are struck down by courts through litigation and appeals. As one commentator argued: "an advisory opinion can avoid the harm which is frequently caused by unconstitutional statutes. Such statutes may be in effect for years before they are struck down... Before it is successfully challenged, an unconstitutional statue may discourage legitimate activity; or, conversely, it may encourage reliance which, when the statute is invalidated, will prove to have been ill-founded, thus causing injury to those who have based action upon it." Note (1956)

Scholars who oppose advisory opinions argue that such opinions are based on poor information. In the U.S., the Supreme Court rules in the presence of adversarial legal proceedings, where a plaintiff and defendant produce information useful to the court as evidence. Cases are determined by empirical facts (i.e. a crime was committed or not, a law was implemented in an unconstitutional manner or not, etc.) and not on the law's expected outcomes.

Judging the constitutionality of a law often depends on how the law was enforced and on information not available during advisory proceedings. The primary argument against advisory opinions, due to Felix Frankfurter, is known as Frankfurter's Objection:

> Legislation is largely empirical, based on probabilities... and not on demonstration... [Further] the history of modern legislation amply proves that facts may often be established in support of measures after enactment, although not in existence previously...In the attitude of court and counsel, in the availability of facts which underlie litigation, there is a wide gulf between opinions in

> advance of legislation or executive action, and decisions in litigation after such proposals are embodied into law or carried into execution (Frankfurter (1930))

3.1.2 Advisory Opinions in State Courts

In the United States, nineteen state supreme courts have provided advisory opinions on pending legislation and intra-governmental issues at some point since the U.S. Constitution was adopted. Why the divergence between the federal Supreme Court and state courts? One reason for the divergence is the timing and influences in which the judiciaries were formed. The American Revolutionary War (1776-1783) ended without a national constitution. However, states began writing their new constitutions before the end of the War. States adopted elements of their colonial governing documents under British rule in their new constitutions. Hence the colonies' constitutions were often more *British* than *American*. As cited by Persky (2004-2005), "The advisory opinion itself was a direct adoption of the English practice[3] . The conventioneers had translated "king" to "governor" and "House of Lords" to "senate," and then for parity's sake extended the power to the House of Representatives."

Unlike colonies-turned-states, not all states in Table 3.1 had a history of direct

[3] Notably, the current English judicial system does not explicitly allow for advisory opinions, but has a rich tradition of judicial involvement in legislative affairs. For centuries, English judges would consult the monarch and the House of Lords on legal matters. The practice fell out of favor in the 19th century (Jay (1997)).

British governance. So why did states like Missouri or Nebraska purposely adopt advisory opinions? Persky (2004-2005) writes that "... the adoption of the advisory opinion was merely the result of an absence of a convincing argument against the advisory opinion."

Table 3.1: U.S. State Supreme Courts with Advisory Opinions (Data from Rogers and Vanberg (2002))

State	# of Advisory Opinions (1975-2001)	Avg. #/Year	Year Advisory Opinions Ceased
Alabama	133	5.11	–
Colorado	10	0.63	–
Connecticut	–	–	1867
Deleware	11	0.42	–
Florida	43	1.65	–
Kentucky	–	–	1881
Maine	24	0.92	–
Massachusetts	60	2.30	–
Michigan	12	0.46	–
Missouri	–	–	1874
Nebraska	–	–	1893
New Hampshire	87	3.35	–
New Jersey	–	–	1912
New York	–	–	1845
North Carolina	1	0.04	–
Pennsylvania	–	–	1808
Rhode Island	40	0.65	–
South Dakota	8	0.31	–
Vermont	–	–	1864

States such as Louisiana, Maryland, Minnesota, Mississippi, North Dakota, Ohio, Texas, Tennessee, and Wisconsin allowed for advisory opinions but were never used (Rogers and Vanberg (2002)). However, even in the states that did use advisory opinions, the frequency of opinions has declined over time. Table 3.1 shows that eight states discontinued the practice of advisory opinions altogether–all coming before 1912. Of the states that have continued issuing advisory opinions, there is wide variation in their use.

Using data from 1976-2001, Alabama's Supreme Court has issued an average of more than five opinions per year, while North Carolina has issued just one over the same period.

3.1.3 Abstract Judicial Review in Europe

A practice similar to advisory opinions is "abstract judicial review" and is common among European nations. In abstract judicial review, the court rules on the constitutionality of a law immediately after it is enacted. For example, in France, the Constitutional Council can review a newly passed law before it is implemented if requested by a legislative minority. Once the Council has ruled on the constitutionality of the proposed law, the Council's ruling is final and cannot be challenged. Abstract judicial review differs from advisory opinions in that the court makes a ruling on the law after it has been passed by the legislature but before it is implemented.

In this regard, abstract judicial review and advisory opinions share many of the same features: there are no established facts to decide the constitutionality of the law meaning decisions are made based on poor information such as the law's text. While nations vary in who can instigate an abstract judicial review ruling, all nations in Table 3.2 allow a government body (the Executive, minority party in the legislature, etc.) to demand a ruling on the constitutionality of a law before it is implemented.

Table 3.2: European Abstract Judicial Review (Data from Vanberg (1998))

Country	Court
Austria	Constitutional Court
Bulgaria	Constitutional Court
Czech Republic	Constitutional Court
France	Constitutional Council
Germany	Federal Constitutional Court
Poland	Constitutional Tribunal
Portugal	Constitutional Court
Romania	Constitutional Court
Russia	Constitution Court of the Russian Federation
Slovakia	Constitutional Court
Slovenia	Constitutional Court
Spain	Constitutional Court

3.2 Related Literature

This chapter fits into the literature that discusses the evolution of civil and common law systems including the seminal work of Merryman (1969) and more recently Glaeser and Shleifer (2002). Shin (1998) and Dewatripont and Tirole (1999) show how the adversarial legal processes can increase the amount of information available to a court relative to inquisitorial processes. Maskin and Tirole (2004) compare the quality of decision making under courts, legislatures, and direct democracy.

Advisory opinions have been studied by political scientists, but these papers have focused on the court's role in influencing legislative outcomes. Rogers and Vanberg (2002) show how uncertainty about the court's preferences and the possibility of judicial review forces the legislature to pass only unambiguously constitutional laws. Vanberg (1998) studies the same situation but focuses on how abstract judicial review encourages legislative bargaining. With abstract judicial review, the minority legislative party can request a review of a law passed by the majority. In that model, neither party wants its law struck down by the court and so the majority is more willing to bargain to avoid such a review.

Most closely related to this chapter is the work of Li (2007) and Levy (2005). Li studies the optimal reporting protocol of an expert–having an expert report his signals sequentially can lead to better decision making than reporting simultaneously. Li also focuses on the role of mind changes–when society believes higher ability experts are more likely to change their minds as a result of better information. Unlike Li's model, I generalize the court's (expert's) payoff function to include bias and social welfare. I also consider specific legal institutions which are not the focus of her study.

In the legal literature, reputation is often equivalent with precedent: Landes and Posner (1976) and Miceli and Cosgel (1994) argue that courts follow precedent because they care about not being reversed. Levy (2005) builds a model where a lower court seeks to avoid being overturned by a higher court. The lower court only cares about its reputation (not being overturned) and is unbiased. Reputation in my model is different than these papers: the court cares about how it is perceived by society and not other courts (the "court" in my model is interpreted as the Supreme Court or Constitutional Court and is the court of last resort).

Lastly, this chapter is related to the literature that studies the strategic use of appeals. Shavell (1995) studies how appeals can be used to rectify a lower court's error.

Daughety and Reinganum (2000) and Spitzer and Talley (2000) study the use of "audits" by a high court when reviewing a lower court's ruling.

3.2.1 Judicial Bias & Fact Discretion

A main assumption I make is that in judicial review the court may selectively choose its facts, a phenomenon termed "fact discretion". Fact discretion allows a judge to pick the set of facts he perceives to be most important. Once chosen, these facts are rarely subject to appeal and in large part determine the final legal outcome.

In judicial review the court receives concrete and empirical facts about the law. Based on this information, the court is able to take into account the specific details of the law's implementation and can determine its constitutionality. The court's fact discretion can manifest itself in variety of contexts–a selective reading of "facts", empathy for a litigant, or personal experience. While fact discretion can be thought of in many ways, a natural interpretation is that it represents how the court interprets the facts uncovered through the adversarial legal process. I assume the court cannot predict its interpretation of the facts beforehand and only learns it after observing the concrete facts in judicial review.

In practice, it is not uncommon for justices of the U.S. Supreme Court to read the same briefs and listen to the same oral arguments and still disagree about what evidence is most crucial. Wrightman (2006) discusses the 1985 case *Ake v. Oklahoma*, in which Justices Thurgood Marshall and William Rehnquist wrote opinions that highlighted a different set of facts. These different facts were then used by each justice to justify his differing opinion. In 1979, Glen Burton Ake murdered a minister, his wife, and two children (after attempting to rape one of the children). Ake was denied a psychiatrist in his initial trial but the prosecution was allowed to present its own psychiatrist who concluded Ake was sane at the time of the murders.

Wrightman describes Thurgood Marshall's description of the facts as

> Nothing in [Marshall's] description suggests the possibility that Ake was faking or malingering [his insanity]. We are led to believe that his symptoms were genuine, and the opinion cites the diagnosis, prior to medication, of Ake as a paranoid schizophrenic person. Justice Marshall's view, concurred with the majority of the justices, was that Ake was entitled to the assistance of a state-paid psychiatrist when he claimed, at his trial, insanity.

However, William Rehnquist not only disagreed with the majority's conclusion, but also what facts were important in reaching the opinion. Wrightman describes Rehnquist's description of the facts:

> In contrast to Justice Marshall's opinion (which did not even mention the attempted rape), Justice Rehnquist gave a description of the actual crime, he named the victims, thus personalizing them…Justice Rehnquist concluded his review by saying 'The evidence…would not seem to raise any question of sanity unless one were to adopt the dubious doctrine that no one in his right mind would commit murder.'

The difference in interpretation of the same case highlights an important channel for judicial bias. Gennaioli and Schleifer (2008) term this type of bias "fact discretion"–a selective reading of the cases' facts. Gennaioli and Shleifer give the example of *Garratt v. Dailey*:

> Brian Dailey, a 5-year-old boy, accompanied his mother on a visit to his aunt, Ruth Garratt, in the garden of Garratt's house. The boy allegedly pulled a chair from under his aunt as she started to sit down; she fell and injured herself and subsequently sued Brian. According to the appellate court review of the evidence, "the trial court had accepted boy's statement that he had moved chair and seated himself therein, but when he discovered that plaintiff was about to sit at place where chair had been, attempted to move chair toward plaintiff, and was unable to get it under plaintiff in time" (Garratt v. Dailey, 279 P.2d 1091, 1091 [1955]). Having accepted the boy's view that he was trying to help his aunt rather than hurt her, the trial court ruled for the boy on the grounds that he did not have the purpose–and therefore intent–to harm her.

But after the case was reviewed by a superior court (italics added),

> "…the superior court reviewed the evidence…and entered a finding to the effect that the defendant knew, with substantial certainty, at the time he removed the chair, that the plaintiff would attempt to sit down where the chair had been, since she was in the act of seating herself when he removed

> the chair" (49 Wash. 2d 499, 500 [1956]). *The trial court shifted all the way from the finding that the boy was moving the chair toward the aunt as she was sitting down to the finding that he was pulling it from under her.*

I assume fact discretion is unpredictable to both the court and society. Neither understands how the court will select facts before it is presented with information through the adversarial process[4] . It has been suggested that court justices are political players with an ideology that can be predicted. My model does not allow for this type of "predictable bias". I view this as reasonable for two reasons. First, the court in my model is a collection of different ideologies, personalities, and intra-group dynamics. The court is then an aggregation of each justice's characteristics (while U.S. Supreme Court observers may know what the most conservative or liberal justices are likely to believe, predicting the behavior of moderate justices, or swing votes, is more difficult). Second, fact discretion largely depends on what is most salient to the court when it is presented with the facts. If in the case of *Ake v. Oklahoma*, the state highlighted the personal details of the attempted rape or suffering of the victims, Marshall's opinion might have been different. Predicting what facts will be most important to a court is not possible until its opinion is read aloud.

[4] Accurately predicting how a court will rule on a case is not easy, even for the justices. As Wrightman (2006) notes, there is a large amount of bargaining and negotiating between justices after oral argument but before opinions are published. Sometimes this negotiating changes the outcome of the courts opinion.

3.3 Model

In this section, I describe a simple model of how society chooses its judicial system. There are two players in the game: society, which can be thought of as the ruling (benevolent) elites and the court. I assume the court is a representative decision maker that accounts for the potentially divergent preferences of multiple justices. The court must decide if a law is constitutional or unconstitutional. Let $\theta \in \{G, B\}$ be the constitutionality of the proposed law where $\theta = G$ implies the law is constitutional and $\theta = B$ implies the law is unconstitutional. A constitutional law is beneficial to society if implemented while an unconstitutional law harms society if approved by the court. I assume society is uncertain if the law is constitutional, $Pr(\theta = G) = Pr(\theta = B) = 1/2$, but the court receives private information about θ over time.

Society chooses when and how often the court rules on the law's constitutionality in an effort to ensure the correct decision is made. Society has three options:

1. *Advisory Opinion*: The court makes one binding ruling *before* all information about the law's consequences is received.

2. *Judicial Review*: The court makes one binding ruling *after* all information about the law's consequences is received.

3. *Advisory Opinion & Judicial Review*: The court makes a non-binding ruling based on preliminary information and a binding ruling after more evidence about the law's constitutionality is uncovered.

Society's welfare is directly affected by the court's rulings: in each period $t = 1, 2$, the court makes a public ruling, $r_t \in \{y, n\}$. The court rules in favor of the law if $r_t = y$ (i.e "yes") and rules against the law if $r_t = n$ (i.e. "no"). If only advisory opinion or judicial review is chosen, the court's report is its final ruling which determines society's welfare. If both advisory opinions and judicial review are chosen, the court publicly announces r_1, but its second ruling r_2 determines social welfare. In this case, r_1 has no direct effect on social welfare, but it does influence the court's reporting strategy.

Information

The court sequentially receives two informative signals s_1, s_2 about the law's constitutionality θ. The first signal, $s_1 \in \{s_G, s_B\}$ comes during advisory opinion proceedings while s_2 is received during judicial review. Signal s_1 is interpreted as the court being provided the text of a law before it has been implemented. The court must rule based on preliminary information such as the law's wording or intent. This imprecise information makes it difficult for the court to predict how the law will be implemented and if society will benefit or suffer from the law. If laws were enforced exactly as written, s_1 would perfectly reveal the law's constitutionality θ. However a law written in constitutional terms may be illegally enforced or an unconstitutionally worded bill may be enforced within the confines of the constitution. The facts and considerations that determine the law's constitutionality are only determined after the law has been implemented.

Judicial review provides the court with a more informative signal $s_2 \in \{s_G, s_B\}$

about the law's constitutionality. By reviewing a law after it has been implemented, the court can observe the law's text (i.e. its intent) as well as the empirical facts of its implementation. Michael Stephenson (2011) shows the U.S. Supreme Court's lack of investigative resources forces the Court to rely on interested parties to produce relevant information about θ. The empirical facts and adversarial legal process in judicial review provide the court with better information than it has access to in advisory opinion proceedings about the law's constitutionality.

The court can be either "high" or "low" ability. Let $a = H$ denote high ability courts and $a = L$ denote low ability courts, $a \in A = \{L, H\}$. The court and society have an ex-ante belief $\mu \in [0, 1]$ that the court is high ability, $\mu = Pr(a = H)$. High ability courts are better able to distinguish legal arguments and determine a law's constitutionality at an early date. High ability courts always receive the correct signal in both stages

$$Pr(s_t = s_G|\theta = G, a = H) = Pr(s_t = s_B|\theta = B, a = H) = 1$$

for $t = 1, 2$. Low ability courts receive the correct advisory opinion signal s_1 with probability p_L where

$$p_L = Pr(s_1 = s_G|\theta = G, a = L) = Pr(s_1 = s_B|\theta = B, a = L) \in [1/2, 1)$$

In judicial review, low ability courts are able to determine the constitutionality of the law with certainty,

$$Pr(s_2 = s_G|\theta = G, a = L) = Pr(s_2 = s_B|\theta = B, a = L) = 1$$

Since the court's ability is not known, the probability the court receives the correct signal in advisory opinion proceedings is $\mu + (1-\mu)p_L$. Society's posterior belief the court is high ability is denoted $\hat{\mu}(H|r_1, r_2)$ and is conditional on the court's rulings (r_1, r_2). If the court receives identical signals in the advisory opinion and judicial review stages, $s_1 = s_2$, it is labeled a "consistent" court. If the court receives $s_1 \neq s_2$, it is an "inconsistent" court. Since s_2 is correct with certainty, inconsistent courts receive an incorrect s_1. An inconsistent court knows with certainty it is low ability, $a = L$. However since low ability courts receive the correct s_1 with probability p_L, a consistent court may be high or low ability. The court's belief it is high ability given $s_1 = s_2$ is

$$Pr(a = H|s_1 = s_2) = \frac{\mu}{\mu + (1-\mu)p_L} > \mu.$$

I adopt the notation $r_t = s_t$ when the court's ruling follows its information, i.e., if the court rules $r_t = y$ when $s_t = s_G$ or $r_t = n$ when $s_t = s_B$.

3.3.1 Payoffs

In this section, the payoffs are defined for the judicial system where both advisory opinion and judicial review proceedings (henceforth the "full game") are used. Payoffs accrue to the court and society in the long run. Society's payoff in the full game is conditional on the court's judicial review ruling r_2 and the realized state θ is

$$V_{12}(r_2, \theta) = \begin{cases} W & \text{if } r_2 = y \text{ and } \theta = G \text{ or } r_2 = n \text{ and } \theta = B \\ 0 & \text{otherwise} \end{cases}$$

Naturally, society prefers the correct ruling to be made, $W > 0$.

The court's ex-ante expectation of social welfare is denoted $E[V_{12}(r_1, r_2, \theta)]$ prior to observing θ and is conditional on both reports (r_1, r_2). Only r_2 directly determines $V_{12}(r_2, \theta)$, but society's belief the court made the correct final ruling depends on (r_1, r_2). After observing s_2, the court knows the state θ and hence $V_{12}(r_2, \theta)$ with certainty. Society only observes (r_1, r_2) and its ex-ante expected payoff is

$$E[V_{12}(r_1, r_2, \theta)] = W\left[\hat{\mu}(H|r_1, r_2)Pr(r_t = \theta|H) + (1 - \hat{\mu}(H|r_1, r_2))Pr(r_t = \theta|L)\right] \quad (3.1)$$

The notation $r_t = \theta$ is used to denote the correct final ruling being made, i.e. $r_t = y$ and $\theta = G$ or $r_t = n$ and $\theta = B$ for $t = 1, 12$.

3.3.2 Fact Discretion

Before describing the court's payoffs, it is useful to explain how fact discretion enters the court's reporting problem. The court's fact discretion–or how it views the facts of the case in judicial review–is a random variable x distributed according to a symmetric *c.d.f.*, $F(\cdot)$ with density f, support on the real line, and mean zero.

Depending on the magnitude and sign of x, it is possible for fact discretion to override the court's concern for reputation and social welfare. Fact discretion affects the court's judicial review ruling in the following manner: When $x > 0$ is positive, the court views the facts as favoring $\theta = B$ so it has an incentive to report $r_2 = n$, and when $x < 0$ is negative, the court views the facts as favoring $\theta = G$ and has an incentive to report $r_2 = y$.

Definition 3 *A court "rules in favor" of its fact discretion if i) $x < 0$ and $r_2 = y$ or* *ii) $x > 0$ and $r_2 = n$. A court "rules against" its fact discretion if i) $x < 0$ and $r_2 = n$ or ii) $x > 0$ and $r_2 = y$.*

The court receives payoff $|x|$ if it rules in favor of its fact discretion but receives no benefit from ruling against its fact discretion. The court's payoff from fact discretion is denoted $Q(x, r_2)$ where

$$Q(x, r_2) = \begin{cases} |x| & \text{if } x > 0 \text{ and } r_2 = n \text{ or } x < 0 \text{ and } r_2 = y \\ 0 & \text{if } x > 0 \text{ and } r_2 = y \text{ or } x < 0 \text{ and } r_2 = n \end{cases}$$

Three properties of the fact discretion payoff $Q(x, r_2)$ are worth noting. First, since the expectation of fact discretion x is zero, the expected fact discretion payoff $E[Q(x, r_2)]$ is also zero. This implies fact discretion is not predictable before the court observes x. Second, the court's payoff from fact discretion is independent of information received in judicial review. Signal s_2 informs the court of the state θ, while the court's fact discretion acts as a form of bias potentially distorting the court's report r_2. Third, the magnitude of x is important in determining if the court chooses to follow its fact discretion. A court with stronger fact discretion (larger $|x|$) receives a higher payoff from ruling in favor of its fact discretion than a court with little fact discretion ($|x| \approx 0$).

3.3.3 Court Payoffs

The court's payoffs are received after its final ruling but before society has learned state θ. The risk-neutral court's ex-ante expectation of its payoff is

$$E[U(r_1, r_2; x, \theta)] = \alpha E[V(r_1, r_2, \theta)] + E[Q(x, r_2)] + \lambda \hat{\mu}(H | r_1, r_2) \tag{3.2}$$

The court's expected payoff $E[U(r_1, r_2; x, \theta)]$ is determined by its reporting strategy (r_1, r_2), its future fact discretion x, and the state θ. The court's payoff is characterized by the weight it places on social welfare $\alpha \geq 0$ and reputational concerns $\lambda \geq 0$. The parameter α represents the court's compassion for society: as α increases, the court receives a higher payoff from issuing the correct ruling. Ex-ante, the court cannot predict its fact discretion, so $E[Q(x, r_2)] = 0$ (below, I will suppress the notation (x, r_2) when it creates no confusion and simply use Q). The court only learns x when s_2 is revealed in the judicial review period. The third term on the right hand side is the court's payoff from being perceived as high ability, defined as society's posterior $\hat{\mu}(H|r_1, r_2)$. Both α and λ are interpreted as the *relative* weights compared to fact discretion[5] .

When the court makes a ruling in the advisory opinion stage, the court's ex-ante expected payoff is then

$$\alpha E[V_{12}(r_1, r_2, \theta)] + \lambda \hat{\mu}(H|r_1, r_2)$$

The court has no information about its fact discretion and chooses r_1 based only on its information s_1. In the judicial review stage, the court learns s_2 and its fact discretion x. After observing s_2, the court knows the state θ with certainty and can calculate $V_{12}(r_2, \theta)$. When choosing r_2, the court maximizes

$$E[U(r_1, r_2; x, \theta)] = \alpha V_{12}(r_2, \theta) + Q(x, r_2) + \lambda \hat{\mu}(H|r_1, r_2) \tag{3.3}$$

[5] An alternative formulation would be to have the weights on social welfare, bias, and reputation sum to one. Instead of normalizing these weights, I study the limit behavior as α and λ go to infinity.

3.3.4 Equilibrium

I focus on a perfect Bayesian equilibrium of the game. In equilibrium, the court's second period strategy is conditional on the history of signals received $S = (s_1, s_2)$, period one strategy r_1, and fact discretion x, i.e. $r_2 : S \times \{y, n\} \times \mathbb{R} \to \{y, n\}$. The court's first period strategy r_1 is conditional on its signal s_1, $r_1 : \{s_G, s_B\} \to \{y, n\}$. Society updates its belief $\hat{\mu}^* : \{y, n\} \times \{y, n\} \to [0, 1]$ via Baye's rule,

$$\hat{\mu}(H|r_1, r_2) = \frac{Pr(a = H)Pr(r_1, r_2|a = H)}{Pr(r_1, r_2)}$$

In equilibrium, the court's strategy (r_1, r_2) must be optimal given society's beliefs and society's beliefs must be consistent given the court's strategy. Conditional on the court's equilibrium reporting strategy (r_1, r_2) and beliefs $\hat{\mu}(H|r_1, r_2)$, society chooses the judicial system that maximizes its expected payoff.

The court, not knowing its fact discretion has no incentive to strategically misrepresent period one information. Hence, I focus on truthful revelation equilibrium–the equilibrium where the court truthfully reveals its information in period one, $r_1 = s_1$. I eliminate other informative equilibria that are "mirror" equilibria[6] .

Timing

To summarize, the timing of the game is as follows:

[6] An example of this is the court always lies in the advisory opinion stage $r_1 \neq s_1$ and then uses a mirror reporting strategy in judicial review.

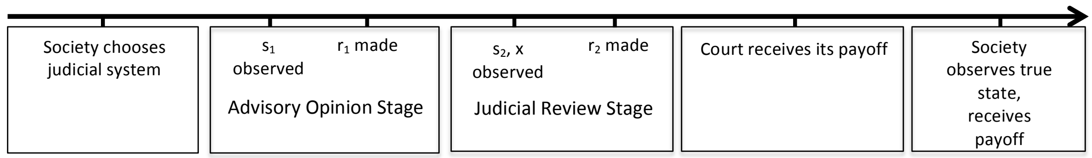

Figure 3.1: The Timing of the Game

Unobserved Fact Discretion

In this section, the court's reporting strategy for the full game is characterized. In judicial review, the court learns i) the state θ and ii) its payoff from fact discretion. Since the court's payoff is continuous and increasing in fact discretion x, the court's optimal reporting strategy is a threshold reporting strategy:

Lemma 3.3.1 *In equilibrium, the court uses a threshold strategy to determine its judicial review report.*

Given Lemma 3.3.1, to find the court's optimal reporting strategy it is sufficient to derive the thresholds used by the court.

Consistent Court Thresholds

First consider a consistent court whose fact discretion agrees with its information ($x < 0$ and $s = s_G$ or $x > 0$ and $s = s_B$). The court then has a dominant strategy to report $r_2 = s_2$. But if the bias disagrees with the court's information, the court will still report truthfully if its payoff from doing so, $\alpha W + \lambda\hat{\mu}(H|r_1 = r_2)$, outweighs the payoff from following its bias $Q + \lambda\hat{\mu}(H|r_1 \neq r_2)$. Then the court will report $r_2 = s_2$ if

$$Q < \bar{Q} \equiv \alpha W + \lambda[\hat{\mu}(H|r_1 = r_2) - \hat{\mu}(H|r_1 \neq r_2)] \tag{3.4}$$

The threshold $\bar{Q}$ is the bias payoff in which the court is indifferent between truthfully reporting its information $r_2 = s_2$, and following its fact discretion $r_2 \neq s_2$.

Since F is symmetric around zero, threshold $\bar{Q}$ holds independently of the state $\theta \in \{G, B\}$. If the court's fact discretion agrees with s_2, a consistent court has a strictly dominant strategy to report $r_2 = s_2$ (and $Q < \bar{Q}$ is trivial). But if fact discretion disagrees with s_2, the court will use $\bar{Q}$ to determine r_2. In effect, fact discretion payoff Q affects the court's reporting decision only when fact discretion *disagrees* with the court's judicial review information. Society's belief that a consistent court will report $r_2 = s_2$ is $F(\bar{Q})$, the probability (3.4) is satisfied.

Inconsistent Court Thresholds

Now consider the case of an inconsistent court whose fact discretion agrees with its judicial review information, s_2. An inconsistent court reports $r_1 = s_1$ but after learning the true state in period two, the court has an incentive to repeat its first report in an effort to appear highly abled. If the court's fact discretion is in favor of s_2, the court's payoff from reporting truthfully, $r_2 = s_2$ is

$$\alpha W + Q + \lambda\hat{\mu}(H|r_1 \neq r_2)$$

By reporting $r_2 = s_2$, the court damages its reputation but receives its fact discretion payoff and higher social welfare. Instead, if the court chooses $r_2 \neq s_2$, it receives a higher reputation from reporting consistently but no payoff from fact discretion or social welfare,

$$\lambda\hat{\mu}(H|r_1 = r_2)$$

The court will rule $r_2 = s_2$ only if

$$\underline{Q} \equiv \alpha W - \lambda[\hat{\mu}(H|r_1 = r_2) - \hat{\mu}(H|r_1 \neq r_2)] > -Q$$

By symmetry of F, the probability an inconsistent court rules $r_2 = s_2$ is $1 - F(-\underline{Q}) = F(\underline{Q})$.

Now consider the case where the inconsistent court's fact discretion is not in favor if its information s_2 (either $x < 0$ and $s_2 = s_B$ or $x > 0$ and $s = s_G$). By reporting $r_2 \neq s_2$ and following its fact discretion, the court receives

$$Q + \lambda\hat{\mu}(H|r_1 = r_2)$$

If the court reports $r_2 = s_2$, it suffers a loss in reputation and forgoes its fact discretion payoff and receives $\alpha W + \lambda\hat{\mu}(H|r_1 \neq r_2)$. The court will truthfully report its second period signal $r_2 = s_2$ only if

$$Q < \underline{Q} \equiv \alpha W - \lambda[\hat{\mu}(H|r_1 = r_2) - \hat{\mu}(H|r_1 \neq r_2)] \tag{3.5}$$

The probability this inequality holds given distribution F is just $F(\underline{Q})$ as before. For an inconsistent court, regardless if fact discretion is in favor of or against s_2, the court rules $r_2 = s_2$ with probability $F(\underline{Q})$.

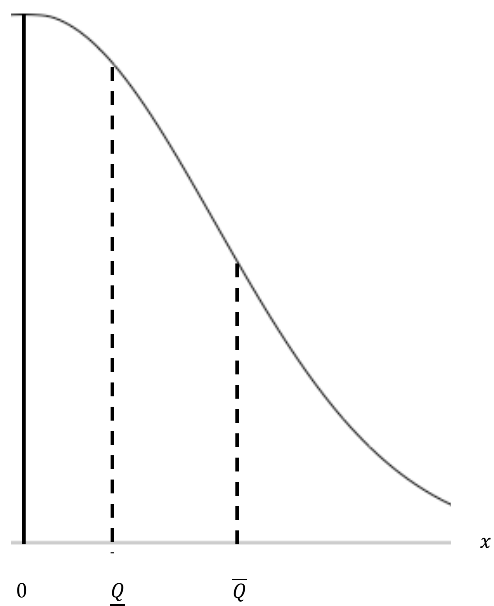

Figure 3.2: Thresholds $\bar{Q}$ and $\underline{Q}$ when $\theta = G$. A consistent court reports $r_2 = y$ only if $Q < \bar{Q}$ while an inconsistent court reports $r_2 = y$ only if $Q < \underline{Q}$.

To fix ideas, consider the following example of a consistent court when the state is $\theta = G$.

Example 3.3.1 (Consistent Courts) *Suppose $\theta = G$. A consistent court receives $s_1 = s_G$ and $s_2 = s_G$ and rules $r_1 = y$ in the advisory opinion stage. When will the court report $r_2 = y$ in judicial review? By reasons provided above, if $x < 0$, the court will always report $r_2 = y$ since its fact discretion supports the court's information.*

But if the court's fact discretion is $x > 0$, the court will report $r_2 = y$ only if x is not too large. When $x > 0$, the court's payoff from $r_2 = n$ is

$$Q(x, n) + \lambda\hat{\mu}(H|r_1 = y, r_2 = n)$$

since ruling $r_2 = n$ reduces social welfare to zero and the court's reputation is damaged by providing conflicting rulings. If instead the court ruled $r_2 = y$, its payoff is

$$\alpha W + \lambda\hat{\mu}(H|r_1 = r_2 = y)$$

Comparing these payoffs, when $x > 0$, the court will report $r_2 = y$ only if

$$Q(x, n) < \bar{Q} \equiv \alpha W + \lambda[\hat{\mu}(H|r_1 = r_2 = y) - \hat{\mu}(H|r_1 = y, r_2 = n)]$$

Where $\bar{Q}$ is the same cutoff describe above in equation (3.4). By symmetry of F, identical analysis is also true for a consistent court when the state in $\theta = B$.

3.3.5 Beliefs

Society updates its belief about the court's ability conditional on reports (r_1, r_2). Given the court's strategies defined by (3.4) and (3.5), society's posterior belief about the court's ability given $r_1 = r_2$ is[7]

$$\hat{\mu}(H|r_1 = r_2) = \frac{\mu F(\bar{Q})}{F(\bar{Q})[\mu + (1-\mu)p_L] + (1 - F(\underline{Q}))(1-\mu)(1-p_L)} \tag{3.7}$$

Similarly, if the court reports $r_1 \neq r_2$, society's belief the court is high ability is

$$\hat{\mu}(H|r_1 \neq r_2) = \frac{\mu(1 - F(\bar{Q}))}{(1 - F(\bar{Q}))[\mu + (1-\mu)p_L] + F(\underline{Q})(1-\mu)(1-p_L)} \tag{3.8}$$

It can be shown that $\hat{\mu}(H|r_1 = r_2) \geq \hat{\mu}(H|r_1 \neq r_2)$ for all parameters values. Now that the court's reporting strategy and society's beliefs have been characterized, it is possible to define the equilibrium for the full game.

Theorem 3.3.2 *A perfect Bayesian equilibrium for the game starting in period two is characterized by the court's thresholds ($\bar{Q}$,$\underline{Q}$) given by (3.4) and (3.5) respectively, and society's beliefs given by (3.7) and (3.8) such that:*

1. *A consistent court's strategy is $r_2 = s_2$ if $Q < \bar{Q}$ and $r_2 \neq s_2$ if $Q > \bar{Q}$, and*
2. *An inconsistent court's strategy is $r_2 = s_2$ if $Q < \underline{Q}$ and $r_2 \neq s_2$ if $Q > \underline{Q}$.*

Theorem 3.3.2 summarizes the court's reporting strategy in judicial review and society's beliefs. An immediate result from the court's thresholds is that consistent courts

[7] Note that

$$\hat{\mu}(H|r_1 = r_2) = Pr(H|s_1 = s_2)Pr(s_1 = s_2|r_1 = r_2)$$

The probability the court is high ability given $s_1 = s_2$ is

$$Pr(H|s_1 = s_2) = \frac{\mu}{\mu + (1-\mu)p_L}$$

The second conditional probability is

$$Pr(s_1 = s_2|r_1 = r_2) \quad = \quad \frac{[\mu + (1-\mu)p_L]F(\bar{Q})}{[\mu + (1-\mu)p_L]F(\bar{Q}) + (1-\mu)(1-p_L)(1-F(\underline{Q}))} \tag{3.6}$$

misrepresent their information for a smaller range of fact discretion than inconsistent courts. This follows from the definitions of $\bar{Q}$ and $\underline{Q}$ given by (3.4) and (3.5) respectively.

Corollary 3.3.1 *In equilibrium, inconsistent courts are more susceptible to fact discretion than consistent courts, i.e.* $\underline{Q} < \bar{Q}$.

By Theorem 3.3.2 and Corollary 3.3.1, the fact discretion payoff Q is always in one of three regions: i) $Q < \underline{Q}$, ii) $Q > \bar{Q}$, or iii) $Q \in (\underline{Q}, \bar{Q})$. In the first case $Q < \underline{Q}$, fact discretion is low enough that neither consistent or inconsistent courts find it optimal to follow its fact discretion so both report $r_2 = s_2$. In case ii) when $Q > \bar{Q}$, both consistent and inconsistent court's fact discretion is large enough to distort their judicial review report so $r_2 \neq s_2$. But if $Q \in (\underline{Q}, \bar{Q})$, each type of court has a different optimal report. Inconsistent courts find Q large enough to deviate from truth telling and report $r_2 \neq s_2$ since $Q > \underline{Q}$. Since a consistent court uses threshold $\bar{Q} > Q$, it reports truthfully, $r_2 = s_2$.

Since the thresholds $(\underline{Q}, \bar{Q})$ determine the court's strategy, knowing how the thresholds behave as parameter value change is important to understand when each judicial system yields the highest expected welfare to society.

Lemma 3.3.2 (Asymptotic behavior of thresholds) *Consider the court's thresholds and society's beliefs defined in Theorem 3.3.2. Then,*

1. *the consistent court's cutoff $\bar{Q}$ increases in λ and α:*

$$\lim_{\lambda \to \infty} \bar{Q} = \infty, \ \lim_{\lambda \to 0} \bar{Q} = \alpha W, \ \lim_{\alpha \to \infty} \bar{Q} = \infty$$

2. the inconsistent court's cutoff $\underline{Q}$ decreases in λ and increases in α:

$$\lim_{\lambda\to\infty}\underline{Q} = -\infty,\ \lim_{\lambda\to 0}\underline{Q} = \alpha W,\ \lim_{\alpha\to\infty}\underline{Q} = \infty$$

The region where consistent and inconsistent courts report differently $(\underline{Q}, \bar{Q})$ expands with reputation. From Lemma 3.3.2, as λ increases, $\bar{Q}$ increases and $\underline{Q}$ decreases. Reputational incentives have the opposite effect on consistent and inconsistent courts. Reputation reinforces the consistent court's incentive to report $r_2 = s_2$ (since $r_1 = r_2$) so $\bar{Q}$ increases meaning it takes larger fact discretion for the court to deviate and rule $r_1 \neq s_2$. Reputational concerns push the inconsistent court to repeat its first ruling. As λ increases, the court's threshold decreases meaning a smaller degree of fact discretion is enough for the court to report $r_2 \neq s_2$.

Given the reporting strategies and beliefs described in Theorem 3.3.2, society's expected welfare from having the court make rulings (r_1, r_2) in the full game is

$$E[V_{12}(r_1, r_2, \theta)] = W \cdot \left[F(\bar{Q})[\mu + (1-\mu)p_L] + F(\underline{Q})(1-\mu)(1-p_L)\right] \tag{3.9}$$

The first term is society's belief the court is consistent $\mu + (1-\mu)p_L$, and consistent courts report $r_2 = s_2$ if $Q < \bar{Q}$ which happens with probability $F(\bar{Q})$. The second term inside the brackets is the payoff from an inconsistent court who reports $r_2 = s_2$ (which occurs with probability $F(\underline{Q})$). In the event the court rules against its judicial review information s_2, society receives a zero payoff.

3.3.6 Single Reports

Now that the full game's reporting strategy, beliefs, and welfare have been characterized, it is straightforward to understand what happens when society elicits only a single report: either the advisory opinion or judicial review judicial system is chosen.

Advisory Opinion

This section considers the case in which all judicial decisions are made in advisory opinion proceedings. Recall that high ability courts always receive the correct signal in the advisory opinion stage but low ability courts only receive the correct signal with probability $p_L > 1/2$. It is not possible for a court to learn its type in the advisory opinion stage as it only views one signal s_1. Similarly, since society cannot update its belief about the court's ability, its belief remains unchanged $\hat{\mu}(H|r_1) = \mu$. This eliminates any strategic reporting incentives by the court to increase its reputation. By assumption, the court is not subject to fact discretion in advisory opinion proceedings, so the court's expected payoff from advisory opinions is $\alpha E[V_1(r_1, \theta)] + \lambda\mu$. The court's optimal strategy is to report truthfully $r_1 = s_1$. From (3.1), society's expected welfare from advisory opinions is:

$$E[V_1(r_1, \theta)] = [\mu + (1 - \mu)p_L] \cdot W \tag{3.10}$$

Society's benefit is simply the payoff from the correct decision W, weighted by the probability the court receives the correct signal $\mu + (1 - \mu)p_L$.

Judicial Review

If the court is asked to rule only once in judicial review, the court observes s_2 and learns the state θ with certainty. The court also learns its fact discretion payoff, $Q(x, r_2)$. To characterize the court's reporting strategy, note the court's judicial review reporting threshold is the same as in the full game with the added constraint $\hat{\mu}(H|r_1) = \mu$. Then thresholds (3.4) and (3.5) reduce to a single cutoff used by all courts in judicial review. The court rules $r_2 = s_2$ only if

$$Q(x, r_2) < \alpha W$$

and $r_2 \neq s_2$ otherwise. This is natural as the court compares the payoff from following its fact discretion with the payoff from maximizing social welfare. The probability a court reports $r_2 = s_2$ is the probability $Q(x, r_2) < \alpha W$, $F(\alpha W)$. Society's expected welfare from judicial review is

$$E[V_2(r_2, \theta)] = W \cdot F(\alpha W) \tag{3.11}$$

3.3.7 Comparison of Judicial Systems

Now that the court's equilibrium reporting behavior and social welfare has been determined in each reporting game, it is possible to find society's optimal judicial system. In this section, it is shown when each judicial system is dominant–when society prefers advisory opinion, judicial review, or both advisory opinion and judicial review proceedings.

The first result characterizes social welfare from the advisory opinion and judicial review single report games. The primary benefit of advisory opinions is the absence of fact discretion that potentially distorts rulings away from society's optimal ruling. However, avoiding fact discretion comes at a cost: the court must make its ruling based on imperfect information. Balancing these two effects, Lemma 3.3.3 shows there is a weight on social welfare α^* such that when the court cares enough about society $\alpha > \alpha^*$, the cost of fact discretion distorting the court's ruling is outweighed by the benefit of better information so judicial review provides higher social welfare than advisory opinions, $E[V_2(r_2, \theta)] > E[V_1(r_1, \theta)]$. The larger α is, the greater the fact discretion payoff $Q(x, r_2)$ must be to distort the court's ruling. If $\alpha < \alpha^*$, the benefit of better information in judicial review is outweighed by the additional welfare loss due to fact discretion, and advisory opinions provide higher social welfare than judicial review.

Lemma 3.3.3 *Advisory opinions provide higher social welfare than judicial review if* $\alpha < \alpha^*$ *where*

$$\alpha^* = \frac{1}{W} F^{-1} \left(\mu + (1 - \mu) p_L\right)$$

Before stating the remaining results, it is useful to understand the asymptotic behavior of the full game's welfare $E[V_{12}(r_1, r_2, \theta)]$, with respect to reputational weight λ. The following corollary describes the relationship between the full game and single report games as λ converges to zero and infinity.

Corollary 3.3.2 *By Lemma 3.3.2,*

1. $\lim_{\lambda \to 0} E[V_{12}(r_1, r_2, \theta)] = E[V_2(r_2, \theta)]$

2. $\lim_{\lambda \to \infty} E[V_{12}(r_1, r_2, \theta)] = E[V_1(r_1, \theta)]$

Part one shows when reputational concerns vanish, social welfare from the full game is identical to welfare from judicial review. Part two shows that $E[V_{12}(r_1, r_2, \theta)]$ converges to advisory opinion welfare when reputation concerns grow large. As reputation concerns increase, the court repeats its first ruling in an effort to appear highly able. In the limit as $\lambda \to \infty$, the court repeats r_1 with probability one leading to the same social welfare in the full game as in advisory opinions. However, Corollary 3.3.2 does not say if $E[V_{12}(r_1, r_2, \theta)]$ converges from above or below in the limit. This is important since society's optimal choice of judicial system crucially depends on how welfare behaves–i.e. if $E[V_{12}(r_1, r_2, \theta)]$ converges to $E[V_2(r_2, \theta)]$ from above, the full game is preferred by society to judicial review when reputational incentives vanish.

The next result shows when the full game is preferred to judicial review. When the court's reputational weight is close to zero, the thresholds used by consistent and inconsistent courts are approximately equal to the threshold used by the court in judicial review only.

Increasing the weight on reputation from zero has two effects. The first effect is society's expected welfare increases because consistent courts become less likely to lie to society. The threshold $\bar{Q}$ increases in λ meaning it is less likely the court's fact discretion will cause it to misrepresent its information, $r_2 \neq s_2$. The second effect of increasing reputation is increased lying by inconsistent courts. An inconsistent court becomes more likely to follow its fact discretion as the equilibrium threshold $\underline{Q}$ decreases in λ. The effect that dominates depends on the size of λ. Since $Pr(s_1 = s_2) > \frac{1}{2}$, a court is more likely to be consistent than inconsistent, so when λ is sufficiently small, the threshold $\bar{Q}$ increases the probability consistent courts report truthfully. This outweighs the costs of increased lying by inconsistent judges as $\underline{Q}$ decreases in λ. For low reputational weight λ, society's expected welfare from the full game is greater than only judicial review, $E[V_{12}(r_1, r_2, \theta)] > E[V_2(r_2, \theta)]$ (See Figure 3.3). As λ continues to increase, the welfare loss from inconsistent judges misrepresenting their information outweighs the gain from decreased lying by consistent judges and $E[V_{12}(r_1, r_2, \theta)] < E[V_2(r_2, \theta)]$.

Proposition 3.3.1 *When the court's weight on reputation is sufficiently small, society's expected welfare from the full game is greater than the welfare from judicial review only. Formally, for $\lambda \in (0, \bar{\lambda})$, $E[V_{12}(r_1, r_2, \theta)] > E[V_2(r_2, \theta)]$. If $\lambda = 0$, $E[V_{12}(r_1, r_2, \theta)] = E[V_2(r_2, \theta)]$.*

The final result shows when the full game dominates advisory opinions. Lemma 3.3.2 shows when reputational concerns grow large $\lambda \to \infty$, the court will repeat its advisory opinion ruling $r_1 = r_2$ in the full game to appear highly abled. Any payoffs derived from reporting truthfully or fact discretion Q is ignored and the court's decision is based solely on s_1. The court's rulings from advisory opinions and the full game will be identical since in the limit $\lambda \to \infty$, $E[V_{12}(r_1, r_2, \theta)] = E[V_1(r_1, \theta)]$.

The shape of $E[V_{12}(r_1, r_2, \theta)]$ is determined by how consistent and inconsistent courts respond to reputational incentives. When λ is small, consistent and inconsistent courts

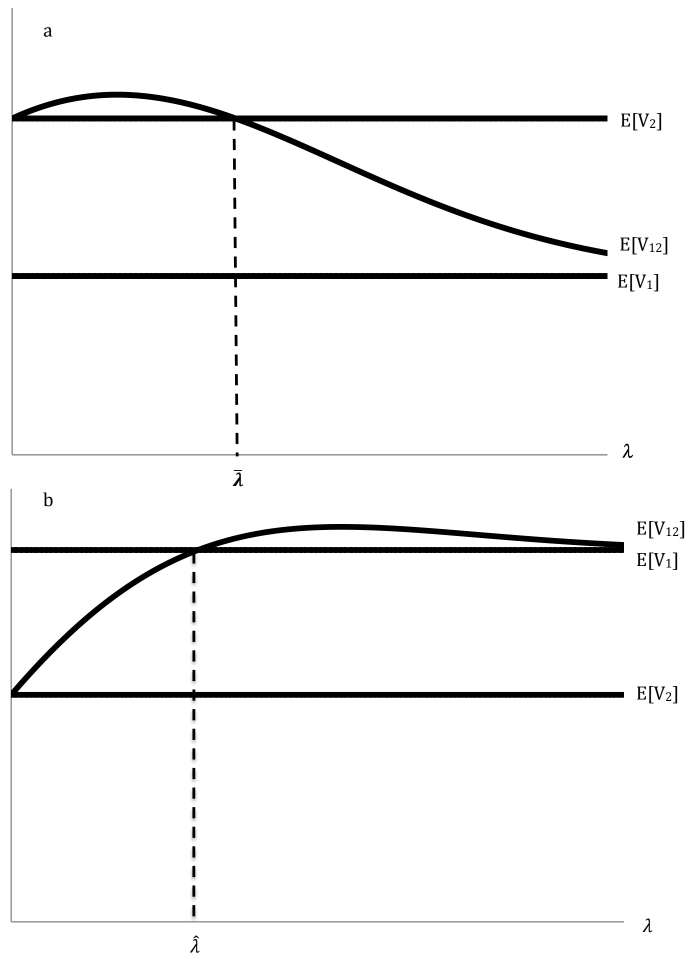

Figure 3.3: Let $x \sim \mathcal{N}(0, \sigma^2)$. Panel a) shows $E[V_{12}(r_1, r_2, \theta)]$ when $E[V_2(r_2, \theta)] > E[V_1(r_1, \theta)]$. Panel b) shows $E[V_{12}(r_1, r_2, \theta)]$ when $E[V_1(r_1, \theta)] > E[V_2(r_2, \theta)]$. In both cases, $E[V_{12}(r_1, r_2, \theta)] \to E[V_2(r_2, \theta)]$ from above as $\lambda \to 0$ and $E[V_{12}(r_1, r_2, \theta)] \to E[V_1(r_1, \theta)]$ from above as $\lambda \to \infty$.

follow the equilibrium threshold strategy defined in Proposition 3.3.2. $F(\bar{Q})$ is the probability a consistent court reports $r_2 = s_2$ and $F(\underline{Q})$ is the probability an inconsistent court reports truthfully $r_2 = s_2$. In the limit as λ increases to infinity, $\lim_{\lambda \to \infty} F(\bar{Q}) = 1$ and $\lim_{\lambda \to \infty} F(\underline{Q}) \to 0$. As λ begins to decrease from infinity, society's gain from consistent courts repeating their first (correct) report outweighs the loss resulting from inconsistent courts either repeating their imprecise information or following their bias. However, the probability of a consistent court repeating its initial report decreases a sufficient amount, the gain from consistent courts repeating their initial reports decreases.

Proposition 3.3.2 *When the court's weight on reputation is sufficiently large, society's expected welfare from the full game is greater than the welfare from advisory opinions only, i.e. for* $\lambda \in (\hat{\lambda}, \infty)$, $E[V_{12}(r_1, r_2, \theta)] > E[V_1(r_1, \theta)]$.

3.3.8 Model without Fact Discretion

Now that the model with fact discretion has been solved, it is useful to consider the case where the court has reputational concerns but does not face fact discretion. Without fact discretion, the court's payoff is

$$E[U(r_1, r_2; \theta)] = \alpha E[V(r_1, r_2, \theta)] + \lambda \hat{\mu}(H|r_1, r_2)$$

The first result shows that in the absence of fact discretion, a consistent court has a dominant strategy to report truthfully $r_1 = r_2$. Doing so simultaneously maximizes society's welfare and the court's reputation.

Lemma 3.3.4 *When the court is not biased, a consistent court will report $r_1 = r_2$.*

The reporting strategy of an inconsistent court depends on how much the court values social welfare. When the court's payoff from making the correct decision αW is large, separation can occur: consistent courts will report $r_1 = r_2$ and inconsistent courts will report $r_1 \neq r_2$. In this region, the gain from maximizing social welfare outweighs an inconsistent court's loss in reputation as society knows with certainty a court reporting $r_1 \neq r_2$ is low ability. When αW is low, by reporting $r_1 = r_2$, an inconsistent court gains from increasing its reputation. Society cannot update its belief since all courts report $r_1 = r_2$ in an effort to appear highly abled. Finally, when social welfare is moderate, consistent courts continue report $r_1 = r_2$ and inconsistent courts are indifferent between reports so they mimic consistent courts with strictly positive probability.

Theorem 3.3.3 *By Lemma 3.3.4, the consistent court will report $r_1 = r_2$ for all parameter values. If $\alpha W \geq \lambda \frac{\mu}{\mu+(1-\mu)p_L}$, an inconsistent court will report $r_1 \neq r_2$. If $\alpha W \leq \lambda\mu$, an inconsistent court will report $r_1 = r_2$. And if $\alpha W \in \left(\lambda\mu, \frac{\lambda\mu}{\mu+(1-\mu)p_L}\right)$ an inconsistent court reports $r_1 = r_2$ with probability $\pi \in (0,1)$ and $r_2 \neq r_2$ with probability $1-\pi$.*

3.3.9 Fact Discretion Uncertainty

Thus far, the court's fact discretion was a realization from a symmetric distribution F with mean zero and support on the real line. In this section, the impact of uncertainty on social welfare is examined by specifying that fact discretion comes from a normal distribution with variance σ^2, $x \sim \mathcal{N}(0, \sigma^2)$. Increasing the court's fact discretion variance σ^2 makes the court more extreme and its rulings less predictable to society. As σ^2 increases, greater mass in the tails of the distribution increases the likelihood fact discretion is large enough to influence the court's reporting strategy. When variance is low, society can expect the court's fact discretion to rarely distort its rulings as the mass of the distribution is close to zero.

Lemma 3.3.5 *Let $x \sim \mathcal{N}(0, \sigma^2)$. As variance σ^2 increases,*

1. *$E[V_1(r_1, \theta)]$ is constant,*
2. *$E[V_2(r_2, \theta)]$ decreases, and*
3. *Denote $\bar{\lambda} > 0$ as the reputational weight where $E[V_{12}(\bar{\lambda})] = E[V_2]$. Then $\bar{\lambda}$ increases as σ^2 increases.*

Part one can be seen from social welfare when only advisory opinions are elicited, $E[V_1(r_1, \theta)] = W[\mu + (1 - \mu)p_L]$. Welfare is independent of distribution F and hence does not depend on variance σ^2. Part two also follows directly from the social welfare from judicial review $E[V_2(r_2, \theta)] = W \cdot F(\alpha W)$. As σ^2 increases, the distribution $F(\cdot)$ decreases (see Figures and 3.4a and 3.5b). This immediately implies that $E[V_2(r_2, \theta)]$ decreases as well. When variance is large, the court's bias is more likely to exceed threshold αW and distort the court's ruling.

When $\sigma^2 \to \infty$, the court's fact discretion becomes extreme and the payoff from judicial review becomes less appealing to society. Since advisory opinions do not suffer from fact discretion distorting the court's ruling, advisory opinions become more appealing to society as $\sigma^2 \to \infty$.

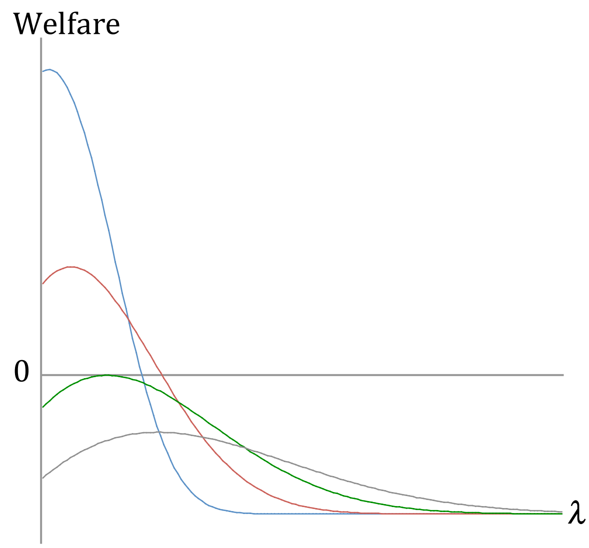

Figure 3.4: $E[V_{12}(r_1, r_2, \theta)]$ as σ^2 increases

The variance affects the full game's social welfare in two ways. First, when $\lambda = 0$, $E[V_{12}(0)] = E[V_2]$ so welfare is decreasing with variance as noted above. Second, the range of λ in which $E[V_{12}(\lambda)] > E[V_2]$ increases. Thus as σ^2 increases, so does the range in which judicial review is dominated by the full game.

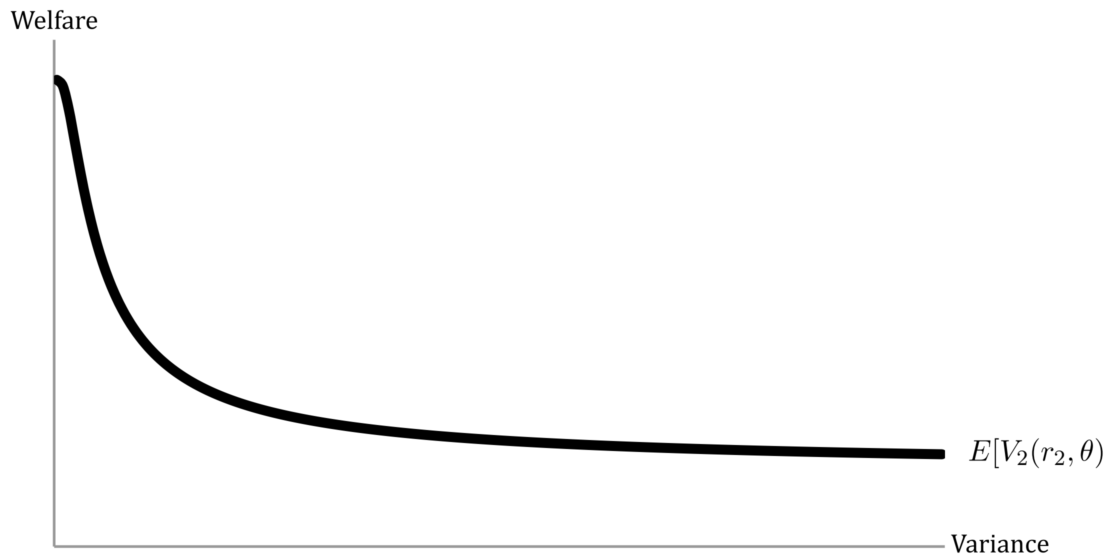

Figure 3.5: $E[V_2(r_2, \theta)]$ as σ^2 increases

3.4 Ability is Private and Information Imperfect

In this section I generalize the model in three directions. First, I consider the possibility the court privately observes its ability. Second, I allow the court's information to be less-than-perfect. And third, I allow for explicit consideration of social costs resulting from incorrect decisions.

Denote $\gamma_1^H \in \left(\frac{1}{2}, 1\right]$ and $\gamma_1^L \in \left[\frac{1}{2}, 1\right]$ as the probability a high and low ability court receive the correct signal in the advisory opinion period, i.e.,

$$\gamma_1^H = Pr(s_1 = s_G | \theta = G, a = H) = Pr(s_1 = s_B | \theta = B, a = H)$$

and

$$\gamma_1^L = Pr(s_1 = s_G|\theta = G, a = L) = Pr(s_1 = s_B|\theta = B, a = L)$$

Since high ability courts are better at discerning legal arguments and predicting the law's consequences, I assume $\gamma_1^H > \gamma_1^L \geq 1/2$. This implies that while the court's information when making an advisory opinion ruling is not perfect, it is informative and is the court's best guess as to the law's constitutionality.

Better information arrives to the court during judicial review. Define γ_2^H and γ_2^L to be the judicial review probabilities of receiving the correct signal s_2 given the court's type,

$$\gamma_2^H = Pr(s_2 = s_G|\theta = G, a = H) = Pr(s_2 = s_B|\theta = B, a = H)$$

and

$$\gamma_2^L = Pr(s_2 = s_G|\theta = G, a = L) = Pr(s_2 = s_B|\theta = B, a = L)$$

High ability courts receive more precise information, $1 \geq \gamma_2^H \geq \gamma_2^L$ and both courts receive better information over time, $\gamma_2^H > \gamma_1^H$ and $\gamma_2^L > \gamma_1^L$.

To generalize society's payoff from the court's rulings, define W_t as society's benefit when the court makes the correct decision (i.e. $r_2 = y$ and $\theta = G$ or $r_2 = n$ and $\theta = B$) and K_t as society's benefit when the court makes the incorrect ruling (i.e. $r_2 = y$ and $\theta = B$ or $r_2 = n$ and $\theta = G$) in $t = 1, 2$. Naturally, society prefers the correct decision to be made, $W_t \geq K_t$ for $t = 1, 2$.

Comparing Systems

Like the case with unknown ability, the court does not have the opportunity for fact discretion in the advisory opinion stage. Society's belief about the court's ability remains unchanged since society's belief is only updated after receiving two rulings, so $\hat{\mu}(H|r_1) = \mu$. It is easy to verify that because the court is not biased or have reputational concerns in advisory opinion proceedings, it never profits from deviating and reporting $r_1 \neq s_1$ (this follows from $W_1 > K_1$). Hence, in the advisory opinion game, the court reports truthfully $r_1 = s_1$. Since both courts report truthfully, the probability the court makes the correct report is

$$Pr(r_1 = \theta|AO) = \mu\gamma_1^H + (1-\mu)\gamma_1^L$$

where the notation $r_t = \theta$ indicates either $r = y$ and $\theta = G$ or $r = n$ and $\theta = B$. Society's ex-ante expected welfare is

$$E[V_1(r_1, \theta)] = W_1[\mu\gamma_1^H + (1-\mu)\gamma_1^L] + K_1[\mu(1-\gamma_1^H) + (1-\mu)(1-\gamma_1^L)] \tag{3.12}$$

and $E[V_1]$ is increasing in μ, γ_1^H, γ_1^L, W_1 and K_1. Using advisory opinions, society benefits when it is more likely the court is highly abled, when the signal precision is increasing, and when payoffs W_1 and K_1 are increasing.

Now consider the case of judicial review. Like advisory opinions, society does not receive enough information to update its belief about the court's ability, so $\hat{\mu}(H|r_2) = \mu$. The court uses threshold strategies to determine its judicial review ruling. Unlike the case with unknown ability, the threshold strategies now depend on the court's ability.

Lemma 3.4.1 *In equilibrium, a court of ability $a \in \{L, H\}$ will report $r_2 = s_2$ if $Q < Q^a$ and $r_2 \neq s_2$ if $Q > Q^a$ where*

$$Q^H = \alpha(W_2 - K_2)(2\gamma_2^H - 1)$$

$$Q^L = \alpha(W_2 - K_2)(2\gamma_2^L - 1)$$

and society's belief about the court's ability is equal to its prior, $\hat{\mu} = \mu$.

A court with ability $a \in \{L, H\}$ will report its second period signal truthfully if $Q < Q^a$. The threshold Q^a is increasing in the benefit from making the correct decision W_1, decreasing in the benefit of a wrong decision K_1, and increasing in the precision of the information received γ_2^a. For any $Q < Q^a$, the court's gain from being truthful outweighs its incentive to follow its fact discretion and misrepresent s_2. But when the payoff from fact discretion is large, $Q > Q^a$ misreporting yields a higher payoff to the court than truthful reporting. If a high and low ability court receive the same signal s_2, the courts will report similarly if $Q > Q^H$ (high and low ability courts both report $r_2 \neq s_2$) and for $Q < Q^L$ (both courts choose $r_2 = s_2$). Only for bias payoff $Q \in (Q^L, Q^H)$ do differently abled courts report differently. High ability courts report $r_2 = s_2$ and low-ability courts report $r_2 \neq s_2$.

If high and low ability courts received the same quality of information, $\gamma_2^H = \gamma_2^L$, then $Q^H = Q^L$, and both types would use the same threshold strategy. Denote society's belief the correct ruling is made in judicial review (that $r_2 = y$ and $\theta = G$ or $r_2 = n$

and $\theta = B$) as

$$Pr(r_2 = \theta | JR) = \sum_{a \in \{L,H\}} Pr(a) \left[\gamma_2^a F(Q^a) + (1 - \gamma_2^a)(1 - F(Q^a))\right] \tag{3.13}$$

Society's expected welfare from judicial review is

$$E[V_2(r_2, \theta)] = W_2 Pr(r_2 = \theta | JR) + K_2 Pr(r_2 \neq \theta | JR)$$

Lastly, I analyze society's decision to require advisory opinion and judicial review rulings. The court's judicial review reporting strategy is conditional on its received signals, its advisory opinion ruling, and fact discretion payoff. A court that has received consistent signals and reports truthfully in the judicial review stage will also make consistent reports and will benefit from higher reputation $\hat{\mu}(H|r_1 = r_2) > \hat{\mu}(H|r_1 \neq r_2)$. As the court's weight on reputation increases, it takes an increasingly large fact discretion to overturn the strategy of $r_2 = s_2$ for a consistent court. Because society values consistent rulings, consistent courts have more to lose by deviating to make inconsistent rulings.

When the court receives inconsistent signals, increasing reputational concerns increases the likelihood the court rules $r_2 \neq s_2$. Since society rewards consistent reports (and the court reports $r_1 = s_1$ in the first period) as λ increases, the court is more likely to repeat its first ruling to take advantage of reputational gains. As in the judicial review only game, the court's reporting strategy depends on the ability of the court and history of signals.

Proposition 3.4.1 *The court's threshold strategy depends on ability $a \in \{L, H\}$, and if the history of signals was consistent or inconsistent $k \in \{c, i\}$. Then there is an equilibrium in which society's posterior belief satisfies $\hat{\mu}(H|r_1 = r_2) \geq \hat{\mu}(H|r_1 \neq r_2)$ and the court uses threshold strategies Q_k^a to determine its judicial review report. If $Q < Q_k^a$ the court reports $r_2 = s_2$, and if $Q > Q_k^a$, the court reports $r_2 \neq s_2$, where*

$$
\begin{aligned}
Q_c^H &\equiv \alpha(2\gamma_2^H - 1)(W_2 - K_2) + \lambda[\hat{\mu}(H|r_1 = r_2) - \hat{\mu}(H|r_1 \neq r_2)] \\
Q_i^H &\equiv \alpha(2\gamma_2^H - 1)(W_2 - K_2) - \lambda[\hat{\mu}(H|r_1 = r_2) - \hat{\mu}(H|r_1 \neq r_2)] \\
Q_c^L &\equiv \alpha(2\gamma_2^L - 1)(W_2 - K_2) + \lambda[\hat{\mu}(H|r_1 = r_2) - \hat{\mu}(H|r_1 \neq r_2)] \\
Q_i^L &\equiv \alpha(2\gamma_2^L - 1)(W_2 - K_2) - \lambda[\hat{\mu}(H|r_1 = r_2) - \hat{\mu}(H|r_1 \neq r_2)]
\end{aligned}
$$

Using equilibrium beliefs and the court's reporting strategy, the probability the court makes the correct ruling in the full game is:

$$Pr(r_2 = \theta|FG) = \sum_{a \in \{L,H\}} Pr(a) \sum_{k \in \{c,n\}} Pr(k|a) \left[\gamma_2^a F(Q_k^a) + (1 - \gamma_2^a)(1 - F(Q_k^a))\right] \tag{3.14}$$

Society's ex-ante expected welfare is

$$E[V_{12}(r_1, r_2, \theta)] = W_2 Pr(r_2 = \theta|FG) + K_2 Pr(r_2 \neq \theta|FG) \tag{3.15}$$

Critics have pointed out numerous flaws in the use of advisory opinions such as reputational effects, poor information, and large costs of incorrect decisions. Despite these arguments, the next proposition shows that advisory opinions can be society's best choice of judicial system, dominating judicial review and the full game. When Proposition 3.4.2 holds, the advisory opinion payoff is high enough that society is willing to forgo better information received in judicial review in favor of an early ruling.

Proposition 3.4.2 *Let $\alpha < \alpha^*$. Advisory Opinions maximize social welfare, $E[V_1(r_1, \theta)] \geq \max\{E[V_{12}(r_1, r_2, \theta)], E[V_2(r_2, \theta)]\}$ if*

$$W_1 \geq \frac{(K_2 - K_1) + Pr(r_2 = \theta|FG)(W_2 - K_2) + Pr(r_1 = \theta|AO)K_1}{Pr(r_1 = \theta|AO)}$$

where $Pr(r_1 = \theta|AO)$ and $Pr(r_2 = \theta|FG)$ are given by (3.12) and (3.14) respectively.

Naturally, the inequality in Proposition 3.4.2 is more likely to hold when the payoff from making the correct decision in advisory opinion proceeding increases, the cost of making an incorrect advisory opinion decreases, the benefit of waiting decreases, and the cost of making an incorrect ruling in judicial review increases.

Propositions 3.3.1 and 3.3.3 also hold in this environment as well if W and K are constant over time. Then with imperfect information and known ability,

Proposition 3.4.3 *The following results are true for the case $W > K = 0$,*

1. *If $\alpha < \alpha^*$, social welfare from advisory opinions is greater than from judicial review.*

2. *If $\lambda \in (0, \bar{\lambda})$, then $E[V_{12}(r_1, r_2, \theta)] > E[V_2(r_2, \theta)]$.*

3. *If $\lambda \in (\hat{\lambda}, \infty)$, then $E[V_{12}(r_1, r_2, \theta)] > E[V_1(r_1, \theta)]$.*

3.4.1 When are Courts Necessary?

Not all legal systems allow the high court the right to review new laws. The analysis has assumed society can ask for a single or multiple rulings, but is silent on why the court is needed in the first place. In other words, why can't society just pass constitutional laws? The assumption society has no information about the constitutionality of the law $Pr(\theta = G) = Pr(\theta = B) = 1/2$ forced it to rely on the court as an "expert" who made the ruling on behalf of the public.

Suppose the probability of a law being constitutional is publicly known to be $\eta = Pr(\theta = G)$ and $1 - \eta = Pr(\theta = B)$ and $\eta \neq \frac{1}{2}$. Society must decide prior to enactment whether the law is constitutional or not. If the law is enacted, with probability η society receives W and with probability $1 - \eta$, society receives K. Then society's expected welfare from not having a court is

$$E[V|\text{No Court}] = \eta W + (1 - \eta)K$$

Proposition 3.4.4 *Society prefers advisory opinions over no court if the probability the correct ruling is made with a court $\mu\gamma_1^H + (1 - \mu)\gamma_1^L$, is greater than the probability the correct ruling is made with no court, η.*

The court's role as a legal expert knowledgeable about society's constitution is a reasonable description of many countries. Founding elites that believe the executive or legislature are sensible and informed about the constitution may believe courts unnecessary. On the other hand, if the founders have a distrust of centralized power and want to provide a check on the executive or legislature, courts can serve that role.

3.5 Conclusion

In this chapter, I consider society's problem of choosing a judicial system and showed when each system was optimal. The court receives better information in judicial review but is susceptible to fact discretion distorting its ruling. By making an advisory opinion, the court has access to poor information, but has no incentive to misrepresent its information. Emphasizing the role of information, fact discretion, and reputation, it was shown when advisory opinions are the optimal judicial system. I show that despite the U.S. Supreme Court's reluctance to use such opinions, advisory opinions can be the optimal judicial system for society to adopt.

The U.S. Supreme Court's refusal to provide advisory opinions in 1793 has not stopped many European and some state courts from doing so. It has been argued that advisory opinions have the major benefit of quick resolution of difficult legal questions. In the United States, it takes many years before a law is considered by the Supreme Court. The cost of implementing laws later deemed unconstitutional can be substantial. Similarly, advisory opinions can prevent the implementation of unconstitutional laws, saving society from enduring bad policies. But there are also a number of arguments against allowing advisory opinions, most notably Frankfurter's Objection. Forcing courts to make a ruling based on imprecise information may lead to poor decision making and subordinate courts (and the justices themselves) may feel compelled to treat such opinions as precedent even when the facts change and the advisory opinion is clearly incorrect. While advisory opinions are technically non-binding, in practice they are treated as such.

Chapter 4

Advance Pricing Agreements and Tax Enforcement

4.1 Introduction and Motivation

Transfer pricing is commonly mentioned by tax authorities and tax professionals as the most important issue in international taxation. This is in part due to the growth in international trade resulting from multinationals–in 2009, intra-firm trade accounted for 30% of imports and 48% of exports in the United States (Lanz and Miroudot (2011)). A multinational's transfer pricing methodology determines where its income and profit are taxed. Variations in tax rates across countries allow firms to strategically choose transfer prices to minimize their tax bill. Realizing the possibility of tax avoidance and facing budget shortfalls, tax authorities around the world have increased enforcement of transfer pricing rules. Ernst & Young (2012) notes

> ... tax authorities continue to add staff devoted to transfer pricing. In a climate of budget freezes for many government agencies, tax authorities appear to have made the cost/benefit calculation to incur additional staffing costs in order to investigate transfer pricing.

Transfer pricing is guided by the "arm's length principle". The arm's length principle dictates that cross-border transactions within a firm should be priced as if the transaction was made by unrelated parties (hence "arm's length"). In general, there is no single methodology that yields the correct transfer price, so the firm is required to use the "best method" available for the specific transaction. However the "best method" and its corresponding transfer price are often ambiguous and often the center of tax disputes between multinationals and tax authorities. The largest tax settlement in the IRS's history resulted from a transfer pricing dispute. In 2006, GlaxoSmithKline paid the IRS "...approximately \$3.4 billion... as part of an agreement to resolve the parties' long-running transfer pricing dispute for the tax years 1989 through 2005 (IRS (2013))." Due to the ambiguity of the arm's length principle, it is possible for the firm and IRS to reach different transfer pricing methodologies despite the firm's best effort to comply with the arm's length principle. PricewaterhouseCoopers (2012) notes

> Since there is no absolute rule for determining the right transfer price for any kind of international transaction with associated enterprises, whether it involves tangibles, intangibles, services, financing or cost allocation/sharing

> arrangements, there is huge potential for disagreement as to whether the correct amount of taxable income has been reported in a particular jurisdiction.

To resolve the uncertainty created by the arm's length principle, a recent development in avoiding international tax disputes is the advance pricing agreement (APA) program. APAs are agreements "... in which governments and taxpayers can agree in advance the appropriate approach to determine the 'arm's length' price to be charged in transactions between related entities (OECD (2012))." The first APA program started in the United States in 1991 and has since expanded to many industrialized and developing countries.

APAs rely on the cooperation between the tax authority and multinational and generally last 3-5 years. From the multinational's perspective, APAs are beneficial due to the reduction of uncertainty about future tax audits, penalties, and litigation. APAs offer certainty–by engaging with the IRS before the firm implements its transfer pricing methodology, the firm eliminates the possibility of ex-post penalties and audit costs. Figure 4.1 shows the number of APAs in the U.S. has been increasing since the APA program was started in 1991.

There are two main costs for a multinational requesting an APA. The first and most direct is financial: the firm must pay an application fee and spend resources preparing documentation for the APA. The second cost is the information revealed to the IRS during the APA process. Givati (2009) argues

> "...in order to obtain an agreement [firms] need to disclose details of their

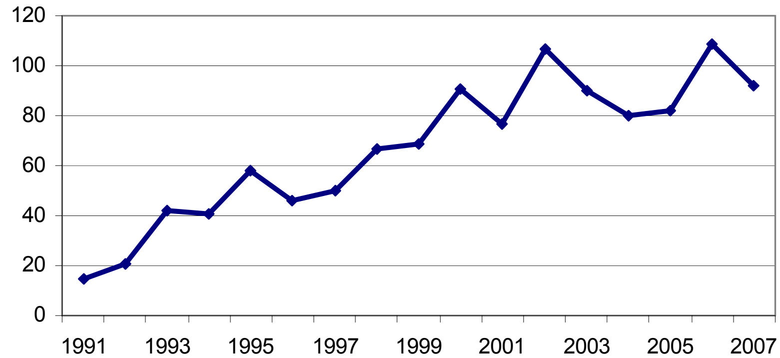

Figure 4.1: Number of APA Applications, 1991-2007, Givati (2009).

> business operations, thus 'red flagging' ambiguous tax issues to be considered by the [Internal Revenue] Service. These tax issues may remain undetected in a regular audit. Furthermore, applying for an advance pricing agreement means that the taxpayer's transfer pricing method will be examined by the APA office, with expert agents who are less likely to make mistakes than regular auditors in the [Internal Revenue] Service's district offices."

Particular industries such as pharmaceuticals, medical devices, semiconductors, auto parts, and financial products have specialized "Issue/Industry Coordination Teams" comprised of economists and an APA manager to handle transfer pricing issues in the specific field.

Despite the growing importance of transfer pricing and advance pricing agreements, there has been little academic research on the subject. The objective of this chapter is to understand when APAs can enhance social welfare and how APAs affect the firm's tax compliance incentives. The model below describes a sequential game between the IRS and a multinational firm. First, neither the firm nor IRS observe what the firm's true transfer pricing methodology should be. This reflects the uncertainty inherent in applying the arm's length principle to firm-specific transactions. Based on observable characteristics, such as the firm's past compliance record, industry, and nature of its international transactions, the firm and IRS share a belief the firm is eligible to take an aggressive transfer pricing methodology. Conditional on this belief, the firm makes its decision to request an APA.

If the firm forgoes the APA, it must choose a transfer pricing methodology to determine its tax liability. I consider two informational environments for the IRS in the audit game. First, if the IRS can observe the firm's transfer pricing methodology (henceforth TPM), the IRS will audit the firm if it chooses an aggressive TPM. Audits are costly to the firm and IRS. If the firm is audited and found to have chosen an aggressive TPM when it was not eligible to do so, the firm must pay the correct tax plus a penalty that can depend on the amount of tax underreported. If a firm is not audited, the firm's reported TPM and tax liability is accepted by the IRS. The second informational environment is that the IRS cannot observe the firm's TPM without first conducting an audit. If the IRS must first invest resources to investigate the firm, the IRS's audit strategy will be to randomly audit the firm (see Appendix C).

If the firm decides to request an APA from the IRS, the IRS will investigate the firm's transactions in question. The IRS's investigation generates public, but imprecise information about the firm's proper TPM. Conditional on this public information, the IRS and firm negotiate a tax payment. I consider two different bargaining regimes in which the tax payments are chosen. In the first bargaining regime, the firm and IRS agree to accept the results of the investigation. Conditional on the information revealed by the IRS's investigation, the firm pays the tax liability of its expected type. The second bargaining game allocates the right to make a take-it-or-leave-it offer to either the firm or IRS. The proposer offers a tax payment that makes the recipient indifferent between accepting and rejecting the offer. If the offer is rejected, the IRS can commit to auditing the firm in the future with certainty.

The firm's decision to request an APA depends on the IRS's audit strategy and the costs of requesting an APA. When the firm and IRS are certain of the firm's type (i.e the firm is either eligible for an aggressive TPM or not with certainty), no auditing takes place in equilibrium. If the firm is likely eligible for the aggressive TPM, it always chooses the aggressive TPM and the IRS does not audit. If the firm is likely eligible for the conservative TPM, the IRS will audit any aggressive TPM, so the firm reports conservatively.

When the IRS and firm are most uncertain about the firm's type, the firm finds it optimal to choose an aggressive TPM while the IRS finds it optimal to audit aggressive TPMs with certainty. This region of disagreement is where APAs are most useful to the firm and the IRS. Wishing to avoid a costly audit, the firm may request an APA to resolve the uncertainty surrounding its appropriate TPM. It is shown that APAs weakly increase social welfare relative to the case in which APAs are prohibited.

An interesting result is that the IRS may increase social welfare by decreasing its bargaining power in the APA negotiation. High bargaining power ensures the IRS can extract a large tax payment from the firm. In response, the firm will optimally choose to forgo the APA even when there is uncertainty about the firm's type. As the IRS's bargaining power decreases (and the firm's increases), the firm's payoff from the APA increases. By voluntarily reducing its bargaining power, the IRS can ensure the firm chooses the APA. This is striking given the arguments against APAs–critics charge the IRS is too accommodating in its treatment of firms in the APA. This chapter suggests the IRS cannot be too aggressive in the APA without discouraging firms from the program. The IRS's payoff from the APA, even if its negotiating power is reduced, is greater than the social welfare from ex-post auditing.

4.2 Related Literature

This chapter is related to a number of papers in accounting, law, and economics. Beck, Davis, and oh Jung (1996) study the role of tax advice in taxpayers' reporting decisions. When individuals are uncertain about their tax liability, asking for help from an accountant can act as a signal that may affect the IRS's audit strategy. Their model incorporates two-sided informational asymmetry: the IRS has less information about the individual's income, and the individual doesn't know how expensive it is for the IRS to audit. A similar paper is Graetz, Reinganum, and Wilde (1986) who build a simple model of firm reporting and IRS auditing. However, they do not consider the possibility of tax advice or uncertainty about the firm's type. Papers that have studied related auditing problems in other contexts include Border and Sobel (1987) and Khalil (1997). For an extensive survey of the tax compliance literature see Andreoni, Erard, and Feinstein (1998).

Tomohara (2004) and Waegenaere, Sansing, and Wielhouwer (2007) address the use of bilateral APAs by multinationals. When firms report profit in multiple countries, there is a risk of double taxation–the same income may be taxed by more than one country. To eliminate the risk of double taxation the firm can apply for a bilateral APA with two tax authorities. Waegenaere, Sansing, and Wielhouwer (2007) characterizes this game, but does not allow for ex-post auditing or for uncertainty about the firm's tax liability. Closest to this chapter is Simone, Sansing, and Seidman (2011) who study a game where the firm has private information about its type. Hyde and Choe (2005) study how keeping two sets of transfer prices–one for tax purposes and one for managerial incentives can be beneficial to a multinational. Holmstron and Tirole (1991) and Anctil and Dutta (1999) study the role of transfer prices in organizational design and managerial incentives.

An empirical literature has emerged analyzing the relationship between a country's corporate tax rate and transfer prices. Clausing (2003) and Bernard, Jensen, and Schott (2006) find evidence to support income shifting, which suggests firms strategically choose their transfer prices to minimize tax payments. Clausing (2003) uses Bureau of Labor

Statistics (BLS) data and finds “a tax rate 1 percent lower in the country of destination/origin is associated with intra-firm export prices that are 1.8 percentage point lower and intra-firm import prices that are 2.0 percentage points higher.” Bernard, Jensen, and Schott (2006) find a “one percentage point reduction in the foreign tax rate increases the difference between arms-length and related-party prices of 0.56-0.66 percent.”

Malik (1993) and Kaplow and Shavell (1994) were the first to study self-reporting of regulatory violations. If the firm commits a violation (i.e. pollutes a river), the firm can either notify the regulator of the violation or not. If the firm does not notify the regulator, the regulator probabilistically audits the firm and if the violation is discovered, the firm is penalized a greater amount than if it had self-reported the violation. These papers show self-reporting is beneficial to society as it reduces the need for ex-post regulator auditing.

The APA program is similar to other voluntary regulatory systems. Hahn (2000), Shekhar and Williams (2004), Gonzalez and Benitez (2009), and Choe and Shekhar (2010) all examine pre-merger notification programs. In the United States many merging firms are required to gain anti-trust approval before the merger can be completed (this would be similar to mandating APAs). In Australia, the firms can voluntarily ask the anti-trust authorities to approve the merger (this is similar to a voluntary APA program).

4.3 Model

There are two players in the game: a risk averse firm and risk neutral IRS. There are two types of firms denoted by $\tau \in \{\tau_E, \tau_N\}$, where τ_E indicates the firm is legally eligible to use an "aggressive" transfer pricing methodology (TPM) and τ_N indicates the firm must use a "conservative" TPM (it is "not eligible"). If the firm is type τ_E, it can use an aggressive TPM to reduce its tax payment to the IRS to T_A. If the firm's type is τ_N, it is not eligible to use the aggressive TPM and must pay the IRS T_C. An aggressive transfer pricing methodology reduces the firm's tax payment to the IRS, i.e. $T_C > T_A > 0$.

The firm's payoff function $u(\cdot)$ is concave and strictly increasing. The firm's objective is to minimize tax payments made to the IRS and its tax compliance costs. The IRS's objective is to maximize social welfare net audit costs.

The firm's type is unknown to the firm and IRS. Both players share a common prior, $\pi \in [0, 1]$, that the firm is type τ_N, where $\pi = Pr(\tau = \tau_N)$. The belief π reflects the firm's industry, past tax payments, history of tax compliance, public financial statements, etc. In this regard, π can be interpreted as the firm's audit class comprised of observably similar firms. The timing of the game is shown in Figure 4.2.

4.3.1 Audit Game Only

To solve the full game shown in Figure 4.2, I begin by considering the audit subgame. Studying this game in isolation is equivalent to studying the case in which APAs are prohibited. In the audit subgame, the firm must choose its TPM based only on π without

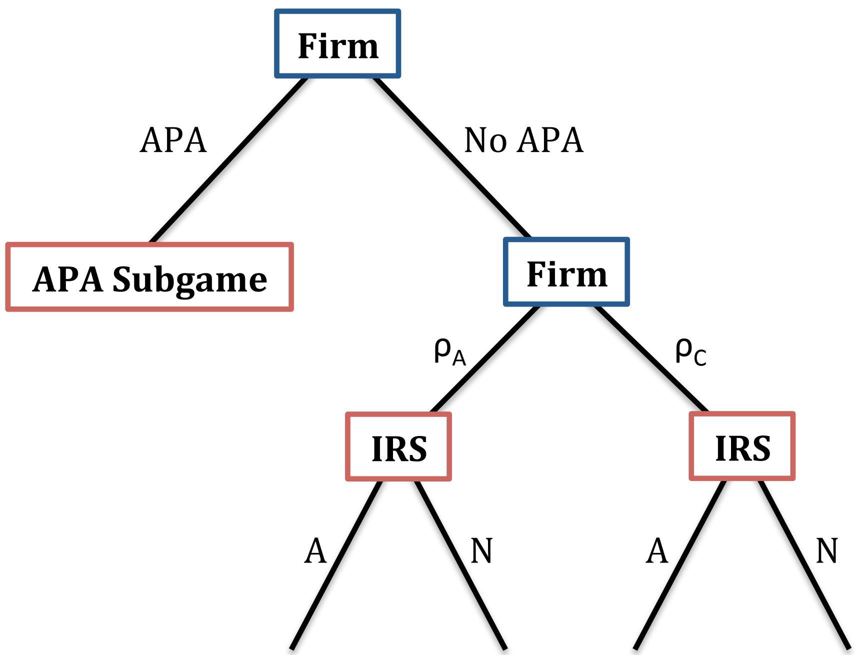

Figure 4.2: Timing of the Full Game

knowing if it is eligible to use the aggressive TPM. Let $\rho \in \{\rho_A, \rho_C\}$ denote the firm's choice of TPM where ρ_A is the "aggressive" TPM and ρ_C is the "conservative" TPM.

It is a commonly made assumption in the tax literature (see Graetz, Reinganum, and Wilde (1986) and Beck, Davis, and oh Jung (1996) for example) that taxpayers have binary types and binary actions that determine their tax payment to the IRS. These two assumptions can be mapped into the real world: for instance the firm must decide whether or not to take a specific action that would reduce its tax payment by a fixed amount[1] .

The IRS observes the firm's TPM choice ρ, and chooses to audit the firm or not (see Appendix C for the case where the IRS's audit decision is independent of ρ.). An IRS audit is costly to both the firm and IRS. The IRS must spend time examining the firm's records and interviewing firm employees to learn τ and the firm must incur documentation, legal, and time costs when its transfer pricing policy ρ is audited. The IRS's audit cost is $c > 0$ and the firm's audit cost is $k > 0$.

If the IRS audits the firm, its investigation reveals the firm's type $\tau \in \{\tau_E, \tau_N\}$ with certainty. This is a result of the IRS observing ex-post the nature of the firm's transaction and other information that may not be available ex-ante. If the firm chooses ρ_A and is found to be type τ_N, the IRS levies an exogenously determined fine $F > 0$ plus tax owed T_C[2] . A type τ_E firm can report either ρ_A or ρ_C without being fined

[1] However, it is just as plausible the type and action spaces are larger. Instead of choosing an aggressive or conservative transfer pricing methodology, in practice firms could use an intermediate methodology. Identically, the ambiguity of the tax system could ensure no firm is τ_E or τ_N with certainty.

[2] This model can handle without difficulty the case where the fine is proportional to the amount of

while a type τ_N firm can only report ρ_C without being fined.

IRS's Audit Subgame Strategy

Using backward induction, the IRS's auditing decision is analyzed first. In the audit subgame, after observing the firm's TPM ρ, the IRS chooses its audit probability $\beta(\rho) \in [0, 1]$. Since there are only two possible transfer pricing methodologies $\{\rho_A, \rho_C\}$ and auditing is costly, it is not beneficial for the IRS to audit if the firm uses TPM ρ_C. Thus, if ρ_C is observed, the IRS audits the firm with probability zero, $\beta(\rho_C) = 0$. Then denote $\beta(\rho_A) \equiv \beta \in [0, 1]$ as the IRS's probability of auditing the aggressive TPM ρ_A. The following assumption ensures the IRS will audit TPM ρ_A in the case it knew the firm was τ_N with certainty.

Assumption 4 *The IRS's benefit from auditing a firm reporting ρ_A when $\pi = 1$ is greater than the cost, i.e.*

$$T_C + F - T_A > c$$

If this assumption did not hold, the audit cost would be too large and the IRS would never audit any TPM.

Denote the IRS's payoff from auditing TPM ρ_A with probability $\beta \in [0, 1]$ in the audit subgame as $\hat{W}(\beta, \pi)$. The IRS chooses β to maximize

$$\max_{\beta \in [0,1]} \hat{W}(\beta, \pi) = \beta \left[\pi(T_C + F - c) + (1 - \pi)(T_A - c)\right] + (1 - \beta)T_A$$

---

underreported tax $F = p(T_C - T_A)$ where $p \in [0, 1]$ is the penalty rate.

The IRS's payoff from auditing ρ_A depends on the likelihood the firm is type τ_N, π, and the benefit from learning the firm's true type relative to not auditing. If the firm audits TPM ρ_A, with probability π the firm is type τ_N. In this case, the IRS collects penalty F and the correct tax T_C. With probability $1-\pi$, the firm using ρ_A is type τ_E and the IRS collects T_A. Regardless of the audit's outcome, the IRS incurs audit cost c. Denote $\bar{\pi}$ to be the probability in which the IRS is indifferent between auditing a firm that uses ρ_A and not auditing. If $\pi > \bar{\pi}$, the IRS's marginal payoff from auditing the firm is strictly positive so the IRS chooses to audit with certainty, $\beta = 1$; if $\pi < \bar{\pi}$, the IRS's marginal payoff from auditing is strictly negative so the IRS refrains from auditing the firm $\beta = 0$; and if $\pi = \bar{\pi}$, the IRS uses a mixed strategy with audit probability $\beta \in [0, 1]$. The IRS's choice of audit probability $\beta(\pi)$ is:

$$\beta(\pi) = \begin{cases} 1 & \text{if } \pi > \bar{\pi} \\ [0,1] & \text{if } \pi = \bar{\pi} \\ 0 & \text{if } \pi < \bar{\pi}. \end{cases} \tag{4.1}$$

The cutoff $\bar{\pi}$ depends on the IRS's cost of auditing c and the gain from detecting a non-eligible firm, $T_C + F - T_A$,

$$\bar{\pi} \equiv \frac{c}{T_C + F - T_A} \tag{4.2}$$

The IRS's threshold $\bar{\pi}$ is strictly positive and less than one by Assumption 4 and $c > 0$.

Remark 4.3.1 (Comparative Statics on $\bar{\pi}$) *The threshold* $\bar{\pi} \in (0, 1)$ *is:*

1. *increasing in audit cost c and the tax payment of type τ_E firms, T_A, and*

2. *decreasing in the tax payment of τ_N firms, T_C, and the fine F.*

The range of π in which the IRS audits with certainty $\pi \in (\bar{\pi}, 1)$, decreases as its cost of auditing increases and as the additional benefit from catching a non-eligible firm decreases ($T_C + F - T_A$ decreases).

Firm's Audit Subgame Strategy

Now that the IRS's audit strategy has been determined, consider the firm's choice of transfer pricing methodology in the audit subgame. In the audit subgame, the firm must choose an aggressive or conservative transfer pricing methodology. Denote $\alpha \in [0, 1]$ to be the probability the firm uses TPM ρ_A. Denote the firm's expected payoff from choosing α given π and β as $\hat{V}(\alpha; \beta(\pi), \pi)$. The firm chooses α to minimize the sum of its tax payment to the IRS and audit costs,

$$\min_{\alpha \in [0,1]} \hat{V}(\alpha; \beta(\pi), \pi) \tag{4.3}$$

where $\hat{V}(\alpha; \beta(\pi), \pi)$ is

$$\alpha\pi \left[\beta u(T_C + F + k) + (1 - \beta)u(T_A)\right] + \alpha(1 - \pi)\left[\beta u(T_A + k) + (1 - \beta)u(T_A)\right] + (1 - \alpha)u(T_C)$$

If the firm reports ρ_A but is in fact type τ_N and the IRS investigates the firm, the firm must pay fine F in addition to tax T_C and audit cost k. If the firm is not audited, it pays the IRS T_A and incurs no audit cost. By choosing ρ_C and paying T_C, the firm can shield itself from audit costs at the expense of a higher tax payment T_C.

Remark 4.3.2 (Comparative Statics) *The firm's expected payoff $\hat{V}(\alpha; \beta(\pi), \pi)$ is increasing in the likelihood the firm is type τ_N, π, the firm's cost of being audited k, the fine F, and tax payments T_A, T_C.*

Using the IRS audit strategy $\beta(\pi)$, the firm's reporting strategy $\alpha(\pi)$ that solves $\min_\alpha \hat{V}(\alpha; \beta(\pi), \pi)$ can be described as a function only of π,

$$\alpha(\pi) = \begin{cases} 1 & \text{if } \pi < \hat{\pi} \\ [0,1] & \text{if } \pi = \hat{\pi} \\ 0 & \text{if } \pi > \hat{\pi} \end{cases} \tag{4.4}$$

When $\pi < \hat{\pi}$, the firm is confident it is eligible for the aggressive TPM and hence chooses ρ_A with certainty. When the firm is risk neutral, the cutoff $\hat{\pi}$ is

$$\hat{\pi} \equiv \frac{T_C - T_A - k}{T_C - T_A + F} \tag{4.5}$$

where $\hat{\pi} \in (0, 1)$.

Assumption 5 *The tax payment differential is greater than the aggregate auditing cost,*

$$T_C - T_A > c + k$$

This assumption guarantees that the IRS's cutoff is less than the firm's cutoff, $\bar{\pi} < \hat{\pi}$. Now that the IRS's auditing strategy $\beta(\pi)$ and the firm's transfer pricing strategy $\alpha(\pi)$ have been determined, the audit subgame equilibrium can be seen in Figure 4.3.

Proposition 4.3.1 (Equilibrium of Audit Subgame) *Let Assumptions 4 and 5 hold. The equilibrium of the audit subgame is the IRS never audits ρ_C and:*

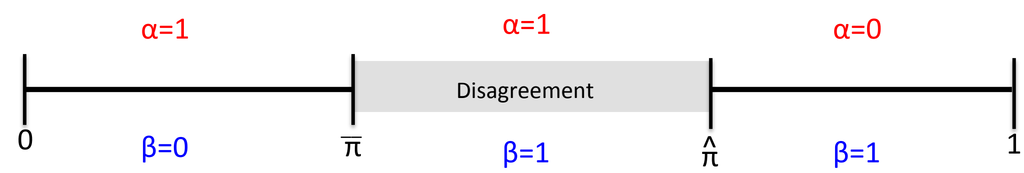

Figure 4.3: Audit Game equilibrium best-responses for firm $\alpha(\beta, \pi)$ and IRS $\beta(\pi)$.

1. *if* $\pi \in (0, \bar{\pi})$ *the firm chooses* ρ_A *and the IRS does not audit* $\beta = 0$*,*

2. *if* $\pi \in (\bar{\pi}, \hat{\pi})$ *the firm chooses* ρ_A *and the IRS chooses* $\beta = 1$*, and*

3. *if* $\pi \in (\hat{\pi}, 1)$ *the firm chooses* ρ_C *and the IRS audits* ρ_A *reports with certainty,* $\beta = 1$*.*

When π is sufficiently large or small, the IRS and firm choose actions that are non-confrontational. If $\pi < \bar{\pi}$, both players believe the firm is likely type τ_E, so the firm chooses ρ_A with certainty and the IRS does not audit. If $\pi > \hat{\pi}$, both parties believe the firm is likely type τ_N, and hence the firm will choose ρ_C and the IRS will refrain from auditing (the IRS will audit ρ_A reports with certainty). Only when $\pi \in (\bar{\pi}, \hat{\pi})$ will the firm and IRS choose actions that are confrontational: the IRS audits ρ_A with probability one and the firm will use TPM ρ_A with probability one. Thus in equilibrium, confrontation takes place between parties when there is the most uncertainty about the firm's true type.

The firm's equilibrium expected payoff in the audit subgame is

$$\hat{V}(\pi) = \begin{cases} u(T_A) & \text{for } \pi \in (0, \bar{\pi}) \\ \pi u(T_C + F + k) + (1 - \pi)u(T_A + k) & \text{for } \pi \in (\bar{\pi}, \hat{\pi}) \\ u(T_C) & \text{for } \pi \in (\hat{\pi}, 1) \end{cases} \tag{4.6}$$

The IRS's equilibrium expected payoff in the audit subgame is

$$\hat{W}(\pi) = \begin{cases} T_A & \text{for } \pi \in (0, \bar{\pi}) \\ \pi(T_C + F) + (1 - \pi)T_A - c & \text{for } \pi \in (\bar{\pi}, \hat{\pi}) \\ T_C & \text{for } \pi \in (\hat{\pi}, 1) \end{cases} \tag{4.7}$$

4.3.2 Advance Pricing Agreements Only

Now that the audit subgame has been analyzed, consider the APA subgame where the firm requests an advance pricing agreement. If requested, the IRS is obligated to conduct an investigation into the firm's true type. By investigating the firm, the IRS generates a publicly observed signal $x \in \{x_E, x_N\}$, where x_E is a signal the firm is eligible for an aggressive TPM and x_N is a signal indicating the firm is not eligible. The signal is correct with probability $\sigma \in (1/2, 1]$, i.e.

$$\sigma = Pr(x_E|\tau_E) = Pr(x_N|\tau_N)$$

and is incorrect with probability $1 - \sigma = Pr(x_N|\tau_E) = Pr(x_E|\tau_N)$.

Despite the ex-ante uncertainty surrounding the firm's transaction, the IRS's investigation is informative. In the United States, there are dedicated advance pricing teams that specialize in evaluating specific industries or types of transactions. These teams have experience evaluating firm transactions ex-ante and determining the appropriate TPM $\rho \in \{\rho_A, \rho_C\}$. The public signal $\sigma \in (1/2, 1]$ represents the specialized knowledge of the IRS's APA unit has when evaluating firms requesting an APA.

After observing $x \in \{x_E, x_N\}$, the firm and IRS update their belief about the firm's type. Denote $\mu(\pi, x) = Pr(\tau_N|\pi, x)$ as the posterior belief the firm is type τ_N given π and x. If $x = x_N$, the posterior belief is

$$\mu(\pi, x_N) = \frac{\pi\sigma}{\pi\sigma + (1-\pi)(1-\sigma)}$$

and if $x = x_E$, the posterior belief is

$$\mu(\pi, x_E) = \frac{\pi(1-\sigma)}{\pi(1-\sigma) + (1-\pi)\sigma}.$$

By requesting an APA, the firm incurs three types of costs. The first cost to the firm is the upfront fee charged by the IRS to process an APA, denoted $n > 0$. The second cost incurred by the firm is documentation costs. The firm may need to hire tax auditors to prepare the necessary documentation for the APA. Finally, there are additional costs involved in the APA process as the firm negotiates and responds to IRS inquiries. Denote the firm's documentation and APA cooperation costs as $k' > 0$.

Once the APA fee n is paid, the IRS investigates the firm at cost $\gamma > 0$. Proponents of APAs highlight that APAs are cheaper for both parties relative to ex-post auditing. Thus assume $c > \gamma > 0$ and $k > n > 0$. The savings are due to the cooperative nature of the APA and the avoidance of legal fees often required in ex-post audits.

As a benchmark case, suppose the IRS and firm can commit to the outcome of the

IRS's investigation x. Then conditional on posterior belief $\mu(\pi, x)$, the firm pays the IRS the tax payment of its expected type,

$$V_{APA}(\pi, x) = \mu(\pi, x)u(T_C + n + k') + (1 - \mu(\pi, x))u(T_A + n + k') \tag{4.8}$$

With probability $\mu(\pi, x)$, the firm is believed to be type τ_N and must pay the IRS $T_C + n$. With probability $1 - \mu(\pi, x)$, the firm is believed to be type τ_E and must pay the IRS $T_A + n$. The firm incurs additional compliance costs k' in the course of completing the APA. The firm's tax payment reflects the uncertainty remaining even after x is observed.

Ex-ante, before signal $x \in \{x_E, x_N\}$ is observed, the firm's expected payoff from requesting an APA is

$$V_{APA}(\pi) = \pi u(T_C + n + k') + (1 - \pi)u(T_A + n + k') \tag{4.9}$$

And the IRS's expected payoff before observing x is

$$W_{APA}(\pi) = \pi T_C + (1 - \pi)T_A + n - \gamma \tag{4.10}$$

4.3.3 Voluntary Advance Pricing Agreements

Now that the audit subgame and APA subgame have been analyzed, consider the voluntary APA game (shown in Figure 4.2). In this section, the firm can choose to request an APA or forgo the APA and choose $\rho \in \{\rho_A, \rho_C\}$ in the audit subgame. Denote the probability the firm requests an APA to be $z \in [0, 1]$ and the probability the firm forgoes the APA in favor of the audit game to be $1 - z$. In the voluntary APA game, the firm's strategy is the pair $(z, \alpha) \in [0,1] \times [0,1]$ and the IRS's strategy is the probability it audits aggressive TPM ρ_A, $\beta \in [0,1]$.

The firm chooses z to minimize its expected payment to the IRS,

$$V(z) = zV_{APA}(\pi) + (1-z)\hat{V}(\alpha; \pi, \beta)$$

where $\hat{V}(\alpha; \pi, \beta)$ is the firm's equilibrium audit subgame payoff given by (4.6) and V_{APA} is the firm's expected payoff from requesting an APA given by (4.9). The firm's optimal strategy $z(\pi)$ is

$$z(\pi) = \begin{cases} 0 & \text{for } \pi \in (0, \bar{\pi}) \\ 1 & \text{for } \pi \in (\bar{\pi}, \hat{\pi}) \\ 0 & \text{for } \pi \in (\hat{\pi}, 1) \end{cases} \tag{4.11}$$

Proposition 4.3.2 *The equilibrium of the voluntary APA game is:*

1. *if $\pi < \bar{\pi}$, the firm forgoes the APA and chooses ρ_A. The IRS chooses $\beta = 0$.*

2. *if $\pi \in (\bar{\pi}, \hat{\pi})$, the firm requests an APA.*

3. *if $\pi > \hat{\pi}$, the firm forgoes the APA and chooses ρ_C. The IRS chooses $\beta = 1$.*

This proposition shows the firm is weakly better off in the presence of APAs. In the audit game, when the IRS and firm disagree about the firm's type, $\pi \in (\bar{\pi}, \hat{\pi})$, the firm reports ρ_A with probability one and the IRS audits the firm with probability one. This is also the range of π in which the firm will optimally choose an APA. When APAs are available, the firm gains from avoiding the certain audit costs. When disagreement with the IRS is likely, the APA offers a convenient way to resolve the issue without the high audit cost.

Social Welfare

In this section, social welfare is characterized when APAs are voluntary, mandatory, and prohibited. By Proposition 4.3.2, the firm gains from the presence of APAs when uncertainty is greatest, $\pi \in (\bar{\pi}, \hat{\pi})$. But are APAs beneficial from the IRS's perspective?

Proposition 4.3.3 *The IRS's expected revenue when APAs are voluntary is weakly greater than expected revenue when APAs are prohibited.*

This follows directly from reduced IRS audit costs, $0 < \gamma < c$. Allowing advance pricing agreements benefits society in two ways: first, requiring the firm to pay fee $n > 0$ increases revenue and subsidizes the IRS's investigation costs. Second, the IRS's cost of conducting an APA is less than the cost of auditing the firm, $c > \gamma$. Since the APA is cooperative by nature, the IRS can avoid legal, time, and other costs by settling tax disputes through the APA program. This proposition is proved by comparing the IRS's expected payoffs in each regime. The IRS's expected payoff when APAs are prohibited is given by (4.7). The IRS's expected payoff when APAs are mandatory is

$$W_{APA} = \pi(T_C + F) + (1 - \pi)T_A + n - \gamma$$

And the IRS's expected payoff when APAs are voluntary is

$$W_{Voluntary}(\pi, \alpha, \beta) = \begin{cases} T_A & \text{for } \pi \in (0, \bar{\pi}) \\ \pi(T_C + F) + (1 - \pi)T_A + n - \gamma & \text{for } \pi \in (\bar{\pi}, \hat{\pi}) \\ T_C & \text{for } \pi \in (\hat{\pi}, 1) \end{cases}$$

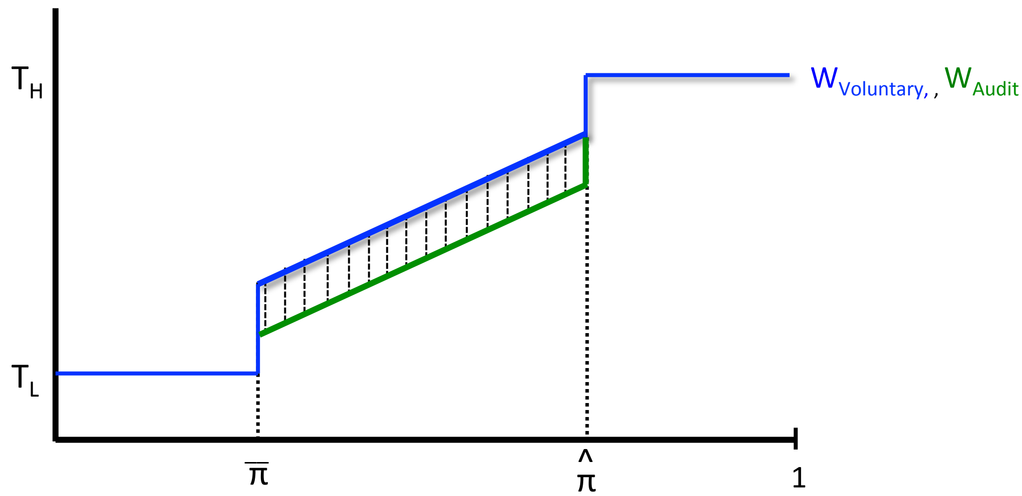

Figure 4.4: Social welfare from voluntary APAs is weakly higher than social welfare from ex-post auditing.

4.4 Alternative APA Settlements: Take-it-or-leave-it Offers

In this section, I consider a game in which the IRS and firm make take-it-or-leave-it tax payment offers as part of the APA process. The IRS's bargaining weight is $\lambda \in [0, 1]$ and is the probability the IRS is selected by nature to make the take-it-or-leave-it offer to the firm. With probability $1-\lambda$, the firm is selected to make a take-it-or-leave-it offer to the IRS. The IRS's (resp. firm's) take-it-or-leave-it offer is conditional on $x \in \{x_E, x_N\}$ and is denoted Q^x (resp. q^x). The bargaining weights are known ex-ante, but the identity of the proposer is not known until after the firm requests the APA.

First consider the case where the IRS proposes offer Q^x to the firm. If the firm accepts the settlement, the firm pays the IRS Q^x and the game ends. If the firm rejects the settlement, the IRS can commit to auditing the firm at a future date with probability one. Once the IRS spends resources in the IRS investigation, there is little additional work for the IRS to do at a later audit date, so the IRS's cost of auditing a firm that rejected the APA settlement is zero. To capture the idea that there is an additional penalty to the firm from exiting an APA, I assume the IRS will treat the firm as if the firm reported ρ_A. This exposes the firm to the risk of being penalized if it abandons its negotiations with the IRS.

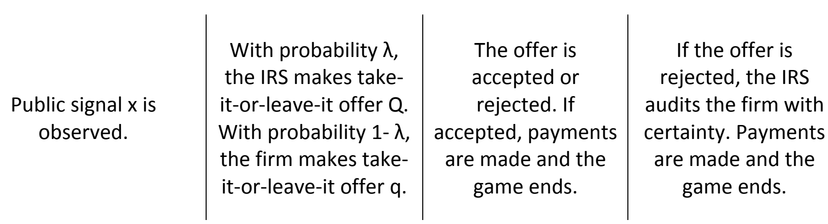

Figure 4.5: APA Subgame with Take-it-or-leave-it offers

The firm's expected payoff from rejecting Q^x when signal x is observed is

$$\mu(\pi, x)u(T_C + F + k' + n) + (1 - \mu(\pi, x))u(T_A + k' + n) \tag{4.12}$$

The IRS optimally selects Q^x so that the firm is indifferent between accepting Q^x and rejecting and receiving (4.12). Thus the IRS picks for $x \in \{x_E, x_N\}$

$$Q^x = \mu(\pi, x)u(T_C + F + k') + (1 - \mu(\pi, x))u(T_A + k') \tag{4.13}$$

The IRS's expected payoff before observing x (but after being given the right to make the take-it-or-leave-it offer) is

$$W_{APA}^{IRS} = Q^{x_E}\left[\pi(1-\sigma) + (1-\pi)\sigma\right] + Q^{x_N}\left[\pi\sigma + (1-\pi)(1-\sigma)\right]$$

where the superscript on W denotes the IRS makes the take-it-or-leave-it offer. With probability $\pi(1-\sigma) + (1-\pi)\sigma$ the investigation reveals $x = x_E$ so the IRS offers tax payment Q^{x_E} to the firm. And with probability $\pi\sigma+(1-\pi)(1-\sigma)$ the IRS's investigation yields $x = x_N$ so the IRS offers Q^{x_N}. Using (4.13), the IRS's expected payoff is

$$W_{APA}^{IRS} = \pi u(T_C + F + k') + (1-\pi)u(T_A + k') + n - \gamma \tag{4.14}$$

Though the IRS is risk neutral, the payment it receives from the firm (i.e. Q^x) depends on the firm's payoff function u and the firm's payoff if audited. The firm's expected payoff when the IRS makes the take-it-or-leave-it offer is

$$V_{APA}^{IRS} = \pi u(T_C + F + k') + (1-\pi)u(T_A + k') + u(n) \tag{4.15}$$

Firm's offer

With probability $1 - \lambda$, the firm makes take-it-or-leave-it offer q^x to the IRS. If the IRS rejects the firm's offer, conditional on x the IRS's expected payoff from auditing the firm at a later date is $\mu(\pi,x)(T_C+F-\gamma)+(1-\mu(\pi,x))(T_A-\gamma)+n$. Then the firm will offer the IRS

$$q^x = \mu(\pi,x)(T_C+F-\gamma)+(1-\mu(\pi,x))(T_A-\gamma)+n \tag{4.16}$$

The IRS's ex-ante expected payoff before observing x when the firm makes the take-it-or-leave-it offer W_{APA}^{firm}, is

$$W_{APA}^{firm} = \pi(T_C+F)+(1-\pi)T_A-\gamma+n$$

The firm's payoff from having offer q^x accepted is

$$V_{APA}^{firm} = \pi u(T_C+F-\gamma)+(1-\pi)u(T_A-\gamma)+u(n)$$

Clearly the firm prefers to propose the take-it-or-leave-it offer since $V_{APA}^{firm} < V_{APA}^{IRS}$. However, ex-ante neither the firm nor IRS know who will make the take-it-or-leave-it offer (but bargaining power $(\lambda, 1-\lambda)$ is known). Thus the IRS's ex-ante expected payoff from the APA without knowing x is

$$W_{APA}(\lambda) = \lambda W_{APA}^{IRS}+(1-\lambda)W_{APA}^{firm} \tag{4.17}$$

where W_{APA}^{IRS} is defined by (4.14). And the firm's expected payoff from requesting an APA is

$$V_{APA}(\lambda) = \lambda V_{APA}^{IRS}+(1-\lambda)V_{APA}^{firm} \tag{4.18}$$

where V_{APA}^{IRS} is defined by (4.15).

Full Game Equilibrium

Now that IRS and firm payoffs from the APA bargaining game have been determined, the full game can be analyzed. The firm must choose to request an APA or forgo the APA. To do so, it will compare its payoff from the audit subgame given by (4.6) with its expected APA payoff given by (4.18).

As in the main model, when $\pi \in (0, \bar{\pi})$, the firm and IRS are confident the firm's type is τ_E so the IRS does not audit and the firm uses TPM ρ_A. And when $\pi \in (\hat{\pi}, 1)$ the firm is likely type τ_N so the firm minimizes tax payments by forgoing the APA and using TPM ρ_C. If $\pi \in (\bar{\pi}, \hat{\pi})$, the firm will reduce its tax liability by requesting an APA only if $V_{APA}(\lambda) < \hat{V}(\pi)$, or

$$\lambda\left[\pi u(T_C + F + k') + (1-\pi)u(T_A + k')\right] + (1-\lambda)\left[\pi u(T_C + F - \gamma) + (1-\pi)u(T_A - \gamma)\right] + u(n)$$
$$< \pi u(T_C + F + k) + (1-\pi)u(T_A + k)$$

For exposition, consider the risk neutral case. This inequality reduces to

$$n + \lambda k' + (1-\lambda)(-\gamma) < k$$

When $\lambda = 0$, the firm will request an APA since $n < k$. If $\lambda = 1$, the firm minimizes tax payments by forgoing the APA if $n + k' > k$. Since the left hand side is monotone and continuous in λ, there exists $\lambda^* \in [0, 1]$ such that for all $\lambda < \lambda^*$, the firm requests an APA and for all $\lambda > \lambda^*$ the firm forgoes the APA. This holds when the firm is risk averse as well.

Proposition 4.4.1 *Suppose $\pi \in (\bar{\pi}, \hat{\pi})$. Then if $n + k' > k$, the firm requests an APA for all $\lambda < \lambda^*$ and forgoes the APA for $\lambda > \lambda^*$. If $n + k' < k$, the firm always requests an APA. The firm does not request an APA if $\pi < \bar{\pi}$ or $\pi > \hat{\pi}$.*

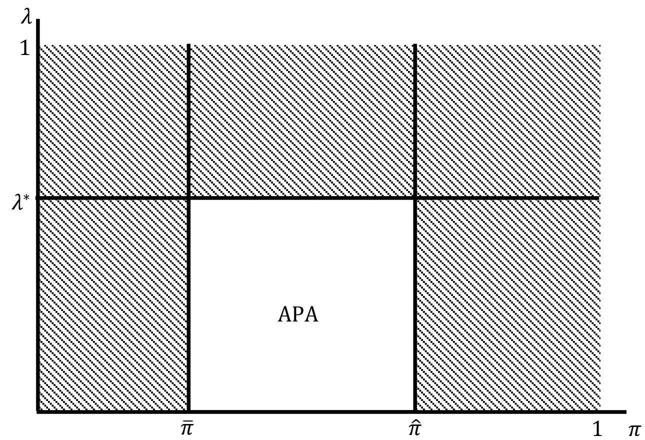

Figure 4.6: Firm's APA Choice: The shaded area represents the parameter ranges of π and λ the firm chooses not to request an APA.

The threshold λ^* is a function of the firm's and IRS's audit and APA costs n, k, k', γ. When the firm is risk neutral, the threshold λ^* is

$$\lambda^* = \frac{\gamma + k - n}{k' + \gamma}$$

By assumption $n < k$ so λ^* is strictly positive, but may not be less than one. If $\lambda^* > 1$, the firm always requests an APA if $\pi \in (\bar{\pi}, \hat{\pi})$. A sufficient condition for $\lambda^* > 1$ is $k > n + k'$, that the firm's administrative costs from requesting the APA are less than the cost of being audited ex-post. If the cost of being audited is not large relative to the cost of APAs, $k < n + k'$, the firm's decision to request an APA will depend on its bargaining power. The firm will request an APA only if it believes it has sufficiently large bargaining power, i.e. if λ is close to zero. If the IRS has sufficiently large bargaining power $\lambda > \lambda^*$, the firm will forgo the APA.

Proposition 4.4.2 *In the full game, the firm's choice to request an APA is given by Proposition 4.4.1. If the firm does not request an APA, the firm's reporting strategy is given by (4.4) and the IRS's audit strategy is given by (4.1). If the firm requests an APA, take-it-or-leave-it offers Q^x and q^x (given by (4.13) and (4.16) respectively) are proposed and accepted.*

The IRS always prefers using the APA when uncertainty is greatest, i.e. when $\pi \in (\bar{\pi}, \hat{\pi})$. If the firm is risk-neutral, the IRS prefers APAs to the audit game when $\pi \in (\bar{\pi}, \hat{\pi})$ only if $\lambda > \frac{\gamma - c - n}{k' + \gamma}$. By assumption, it is cheaper for the IRS to conduct an APA investigation relative to an ex-post audit, so $c > \gamma$. This implies the right hand side is negative, so for all $\lambda \in [0, 1]$, the IRS prefers the APA if $\pi \in (\bar{\pi}, \hat{\pi})$. Thus, the voluntary APA program is socially optimal and weakly dominates the case when APAs are prohibited.

Lemma 4.4.1 *When $\pi \in (\bar{\pi}, \hat{\pi})$, the IRS prefers disputes to be resolved through the APA program for all $\lambda \in [0, 1]$.*

This result suggests that if the APA program is voluntary, the IRS would be better off giving the firm more bargaining power. When the IRS makes the take-it-or-leave-it offer with certainty, the firm is better off by taking its chances in the audit game. The firm only gains from the APA when its bargaining power is a large. This result explains the widely-held belief that the IRS is lenient on firms that do request APAs.

4.5 Conclusion

In this chapter, I characterized firm and IRS incentives in the presence of an advance pricing agreement program. I showed that APAs are most useful to society and the firm when the tax system is most uncertain and there is ambiguity about the firm's tax liability. The chapter also provide a rationale for why the IRS is lenient when negotiating with the firm in the APA process. As the IRS's bargaining power decreases, the firm's payoff from the APA increases making the firm more likely to choose the socially beneficial APA. This result suggests the IRS cannot be too aggressive in the APA without discouraging firms from the program. The IRS's payoff from the APA, even if its negotiating power is reduced, is always greater than the payoff from ex-post auditing.

The model introduced in this chapter abstracts from many potentially important details. For instance, in the model the IRS could audit any firm as long as the cost of auditing c was less than the benefit from detecting a non-eligible firm $c < T_C - T_A + F$. However, a more realistic assumption would be that the IRS faces a budget constraint (see Andreoni, Erard, and Feinstein (1998) for the literature on budget constraints). A principal argument in favor of advance pricing agreements is that the cooperative nature of APAs lower the IRS's costs. An interesting question for future research is how the IRS should allocate its budget between ex-post auditing and APAs to maximize social welfare. A second interesting avenue for future research is allowing for repeated interaction between the IRS and firm. Over time, the IRS can learn from the firm's compliance and the nature of its business model. This would allow the IRS to choose its audit strategy based on more information and for the firm to build a reputation for compliance.

Chapter 5

Bibliography

References

ACHARYA, V., M. PAGANO, AND P. VOLPIN (2012): "Seeking Alpha: Excess Risk Taking and Competition for Managerial Talent," *Working Paper, NYU.*

ADDY, N. D., X. CHU, AND T. R. YODER (2009): "Recovering Bonuses After Restated Financials: Adopting Clawback Provisions," *Working Paper, SSRN.*

ANCTIL, R. M., AND S. DUTTA (1999): "Transfer Pricing, Decision Rights and Divisional versus Firm-wide Performance Evaluation," *The Accounting Review*, 74(1), 87–104.

ANDREONI, J., B. ERARD, AND J. FEINSTEIN (1998): "Tax Compliance," *Journal of Economic Literature*, 36(2), 818–860.

BABENKO, I., B. BENNETT, J. M. BIZJAK, AND J. L. COLES (2012): "Clawback Provisions," *Working Paper, Arizona State U.*

BECK, P., J. DAVIS, AND W. OH JUNG (1996): "Tax Advice and Reporting under Uncertainty: Theory and Experimental Evidence," *Contemporary Accounting Research*, 13(1), 49–80.

Bernard, A. B., J. B. Jensen, and P. K. Schott (2006): "Transfer Pricing by U.S.-Based Multinational Firms," *Working Paper, Yale SOM.*

Biais, B., and C. Casamatta (1999): "Optimal Leverage and Aggregate Investment," *Journal of Finance*, 54(4), 1291–1323.

Biais, B., T. Mariotti, J.-C. Rochet, and S. Villeneuve (2010): "Large Risks, Limited Liability, and Dynamic Moral Hazard," *Econometrica*, 78(1), 73–118.

Bond, P., and U. Axelson (2012): "Wall Street occupations: An equilibrium theory of overpaid jobs," *Working Paper, University of Minnesota.*

Bond, P., and V. Glode (2012): "Bankers and Regulators," *Working Paper, University of Minnesota.*

Border, K. C., and J. Sobel (1987): "Samurai Accountant: A Theory of Auditing and Plunder," *Review of Economic Studies*, 54(4), 525–540.

Bowe, C., and B. White (2007): "Record paybacks over options," *Financial Times.*

Chen, M. A., D. Greene, and J. Owers (2012): "The Costs and Benefits of CEO Clawback Provisions: Theory and Evidence," *Working Paper, GSU.*

Cho, I.-K., and D. M. Kreps (1987): "Signaling Games and Stable Equilibria," *Quarterly Journal of Economics*, 102(2), 179–221.

Choe, C., and C. Shekhar (2010): "Compulsory or Voluntary Pre-Merger Notification? Theory And Some Evidence," *International Journal of Industrial Organization*, 28(1), 10–20.

Clausing, K. (2003): "Tax-motivated transfer pricing and US intrafirm trade prices," *Journal of Public Economics*, 87, 2207–2223.

Daughety, A. F., and J. F. Reinganum (2000): "Appealing Judgments," *RAND Journal of Economics*, 31(3), 502–525.

Dehaan, E., F. Hodge, and T. Shevlin (2011): "Does Voluntary Adoption of a Clawback Provision Improve Financial Reporting Quality?," *Working Paper, University of Washington.*

DeMarzo, P. M., D. Livdan, and A. Tchistyi (2011): "Risking Other People's Money: Gambling, Limited Liability, and Optimal Incentives," *Working Paper, Stanford GSB.*

Dewatripont, M., and J. Tirole (1999): "Advocates," *Journal of Political Economy*, 107(1), 1–39.

Dutta, S., and F. Gigler (2002): "The Effect of Earnings Forecasts on Earnings Management," *Journal of Accounting Research*, 40(3), 631–655.

Equilar (2009): "Clawback Policy Report," *Redwood Shores.*

——— (2011): "Clawbacks are here to stay," *C-Suite Insight.*

Ferreira, D., and M. Rezende (2007): "Corporate strategy and information disclosure," *RAND Journal of Economics*, 38(1), 164–184.

Frankfurter, F. (1930): "Advisory Opinions," *Encyclopedia of the Social Sciences, Seligman and Johnson*.

Gennaioli, N., and A. Schleifer (2008): "Judicial Fact Discretion," *Journal of Legal Studies*, 37, 1–35.

Givati, Y. (2009): "Resolving Legal Uncertainty: The Unfulfilled Promise of Advance Tax Rulings," *Virginia Tax Review*, 29(2), 137–176.

Glaeser, E. L., and A. Shleifer (2002): "Legal Origins," *Quarterly Journal of Economics*, 117(4), 1193–1229.

Glode, V., and R. Lowery (2012): "Compensating Financial Experts," *Working Paper, Wharton Univ. of Pennsylvania*.

Gonzalez, A., and D. Benitez (2009): "Optimal Pre-Merger Notification Mechanisms: Incentives and Efficiency of Mandatory and Voluntary Schemes," *Working Paper, World Bank*.

Graetz, M. J., J. F. Reinganum, and L. L. Wilde (1986): "The Tax Compliance Game: Toward an Interactive Theory of Law Enforcement," *Journal of Law, Economics, and Organization*, 2(1), 1–32.

HAHN, V. (2000): “Antitrust Enforcement: Abuse Control or Notification?,” *European Journal of Law of Economics*, 10(1), 69–91.

HELLWIG, M. F. (2009): “A reconsideration of the Jensen-Meckling model of outside finance,” *Journal of Financial Intermediation*, 18(4), 495–525.

HOLMSTRON, B., AND J. TIROLE (1991): “Transfer Pricing and Organizational Form,” *Journal of Law, Economics and Organization*, 7(2), 201–228.

HYDE, C., AND C. CHOE (2005): “Keeping Two Sets of Books: The Relationship Between Tax and Incentive Transfer Prices,” *Journal of Economics & Management Strategy*, 14(1), 165–186.

INDERST, R., AND S. PFEIL (2011): “Securitization and Compensation in Financial Institutions,” *Working Paper, University of Frankfurt.*

IRS (2013): “IRS Accepts Settlement Offer in Largest Transfer Pricing Dispute,” *http://www.irs.gov.*

JAY, S. (1997): *Most humble servants : the advisory role of early judges.* Yale University Press.

JENSEN, M. C., AND W. H. MECKLING (1976): “Theory of the firm: Managerial behavior, agency costs and ownership structure,” *Journal of Financial Economics*, 3(4), 305–360.

Kaplow, L., and S. Shavell (1994): "Optimal Law Enforcement with Self-Reporting Behavior," *Journal of Political Economy*, 102(3), 583–606.

Kartik, N., M. Halac, and Q. Liu (2012): "Optimal Contracts for Experimentation," *Working Paper, Columbia GSB.*

Khalil, F. (1997): "Auditing without Commitment," *RAND Journal of Economics*, 28(4), 629–640.

Landes, W., and R. Posner (1976): "Legal precedent: A theoretical and empirical analysis," *Journal of Law and Economics*, 19, 249–307.

Lanz, R., and S. Miroudot (2011): "Intra-Firm Trade: Patterns, Determinants and Policy Implications," *OECD Trade Policy Papers No. 114.*

Levine, C. B., and M. J. Smith (2010): "The Relative Efficiency of Clawback Provisions in Compensation Contracts," *Working Paper, Boston University.*

Levitt, S. D., and C. M. Snyder (1997): "Is no news bad news? Information transmission and the role of "early warning" in the principal-agent model," *RAND Journal of Economics*, 28(4), 641–661.

Levy, G. (2005): "Careerist Judges and the Appeals Process," *RAND Journal of Economics*, 36(2), 275–297.

Li, W. (2007): "Changing One's Mind when the Facts Change: Incentives of Experts and the Design of Reporting Protocols," *Review of Economic Studies*, 75, 1175–1194.

Malcomson, J. M. (2011): "Do managers with limited liability take more risky decisions? An information acquistion model," *Journal of Economics and Management Strategy*, 20(1), 83–120.

Malik, A. S. (1993): "Self-Reporting and the Design of Policies for Regulating Stochastic Pollution," *Journal of Environmental Economics and Management*, 24, 241–257.

Maskin, E., and J. Tirole (2004): "The Politician and the Judge: Accountability in Government," *American Economic Review*, 94(4), 1034–1054.

Merryman, J. (1969): *The Civil Law Tradition.* Stanford University Press.

Miceli, T. J., and M. M. Cosgel (1994): "Reputation and Judicial Decision-Making," *Journal of Economic Behavior and Organization*, 23, 31–51.

Murphy, K. J. (1999): *Executive Compensation, Handbook of Labor Economics*vol. 3b, chap. 38, pp. 2485–2563. Orely Ashenfelter and David Card (eds). Elsevier Science North Holland.

Myerson, R. B. (2010): "Moral Hazard In High Office And The Dynamics Of Aristocracy," *Working Paper, University of Chicago.*

Note (1956): "Advisory Opinions on the Constitutionality of Statutes," *Harvard Law Review*, 69, 1302–1313.

OECD (2012): "OECD Launches Project on Improving the Resolution of Cross-Border Tax Disputes," http://www.oecd.org/tax/dispute/, Accessed May 1, 2012.

PALOMINO, F., AND A. PRAT (2003): "Risk Taking and Optimal Contracts for Money Managers," *RAND Journal of Economics*, 34(1), 113–137.

PAREDES, T. A. (2012): "Remarks at Society of Corporate Secretaries & Governance Professionals, 66th National Conference on "The Shape of Things to Come"," Online.

PERSKY, J. D. (2004-2005): ""Ghosts That Slay": A Contemporary Look at State Advisory Opinions," *Connecticut Law Review*, 37.

PRICEWATERHOUSECOOPERS (2012): "International Transfer Pricing 2012," http://www.pwc.com/gx/en/international-transfer-pricing/assets/pwc-international-transfer-pricing-2012.pdf.

ROGERS, J. R., AND G. VANBERG (2002): "Judicial Advisory Opinions and Legislative Outcomes in Comparative Perspective," *American Journal of Political Science*, 46(2), 379–397.

SHARP, S. (2012): "Whose Money Is It Anyway? Why Dodd-Frank Mandatory Compensation Clawbacks Are Bad Public Policy," *Georgetown Journal of Law & Public Policy*, 10(1), 321–341.

SHAVELL, S. (1995): "The Appeals Process as a Means of Error Correction," *The Journal of Legal Studies*, 24(2), 379–426.

SHEKHAR, C., AND P. L. WILLIAMS (2004): "Should the Pre-Notification of Mergers Be Compulsory in Australia?," *Australian Economic Review*, 37(4), 383–390.

Shin, H. S. (1998): "Adversarial and inquisitorial procedures in arbitration," *RAND Journal of Economics*, 29(2), 378–405.

Simone, L. D., R. Sansing, and J. Seidman (2011): "When are Enhanced Relationship Tax Compliance Programs Mutually Beneficial?," *Working Paper, University of Texas Austin.*

Spitzer, M., and E. Talley (2000): "Judicial Auditing," *The Journal of Legal Studies*, 29(2), 649–683.

Stevens, G. N. (1959): "Advisory Opinions–Present Status and an Evaluation," *Washington Law Review*, 34(1), 1–14.

Tomohara, A. (2004): "Inefficiencies of Bilateral Advance Pricing Agreements (BAPA) in Taxing Multinational Companies," *National Tax Journal*, 57(4).

Vanberg, G. (1998): "Abstract Judicial Review, Legislative Bargaining, and Policy Compromise," *Journal of Theoretical Politics*, 10(3), 299–326.

Waegenaere, A. D., R. Sansing, and J. L. Wielhouwer (2007): "Using Bilateral Advance Pricing Agreements to Resolve Tax Transfer Pricing Disputes," *National Tax Journal*, LX(2), 173–191.

Wrightman, L. S. (2006): *The Psychology of the Supreme Court*. Oxford University Press.

Appendix A

Proofs from Chapter 2

Derivation of Benchmark Contract

The incentive compatibility constraint binds, $U_H^b(z) - U_L^b(z) = \Delta c$ and by optimality $w_1(l) = 0$. Then $\pi_\theta^b = \rho_\theta w_1(h)$ so $\bar{w} = [z\rho_G + (1-z)\rho_B]w_1(h)$. Rewriting $U_H^b(z)$,

$$\begin{aligned}
U_H^b(z) &= \int_0^z \bar{w}dF_H(\mu) + \int_z^1 [\mu\rho_G w_1(h) + (1-\mu)\rho_B w_1(h)]dF_H(\mu) \\
&= w_1(h)\left[z[(\rho_G - \rho_B) + \rho_B]F_H(z) + (\rho_G - \rho_B)\int_z^1 \mu dF_H(\mu) + \rho_B[1 - F_H(z)]\right] \\
&= w_1(h)\left[(\rho_G - \rho_B)zF_H(z) + (\rho_G - \rho_B)\int_z^1 \mu dF_H(\mu) + \rho_B\right] \\
&= w_1(h)\left[(\rho_G - \rho_B)\left(zF_H(z) + \int_z^1 \mu dF_H(\mu)\right) + \rho_B\right] \\
&= w_1(h)\left[(\rho_G - \rho_B)E_H[\max(\mu, z)] + \rho_B\right]
\end{aligned}$$

Similarly, $U_L^b(z) = w_1(h)\left[(\rho_G - \rho_B)E_L[\max(\mu, z)] + \rho_B\right]$. The binding incentive compatibility constraint reduces to

$$w_1(h) = \frac{\Delta c}{(\rho_G - \rho_B)[E_H[\max(\mu, z)] - E_L[\max(\mu, z)]]}$$

and $\bar{w}$ can be found using $\bar{w} = z\pi_G^b + (1-z)\pi_B^b$.

Proof of Proposition 2.6.1

Suppose the contracting regime is the benchmark contract. The principal chooses z_b to maximize her expected payoff given the contract described in (2.10) and (2.11). By Assumption 1, f_H is concave implying the principal's problem is concave. When $z = 1$, the marginal cost is $+\infty$ and the marginal benefit is $-f_H(1)\eta_G < 0$, so the marginal cost is greater. By Assumption 2, at $z = 0$, the marginal benefit is greater than the marginal cost. By the intermediate value theorem, there exists a $z_i^* \in (0, 1)$ that satisfies the first order condition

$$-\left[z_b^*\eta_G + (1 - z_b^*)\eta_B\right] f_H(z_b^*) = K'_H(z_b^*) \tag{A.1}$$

Proof of Theorem 2.6.1

It is possible to order the expected compensation cost $K_H^i(z)$ for each contracting regime and arbitrary z. The first result is that $K_H^c(z) \leq \min\{K_H^b(z), K_H^d(z)\}$. Since K is convex, it is also the case that $K_H^{'c}(z) \leq \min\{K_H^{'b}(z), K_H^{'d}(z)\}$. Hence the marginal cost is lowest when the clawback contract is used for any z. Then $z_c^* > \max\{z_d^*, z_b^*\}$ follows from the principal's first order condition. The marginal cost of deferred compensation is lower than it is for the benchmark contract when $\delta > \hat{\delta}$, so it follows that $z_b^* < z_d^*$ for $\delta > \hat{\delta}$ and $z_b^* > z_d^*$ for $\delta < \hat{\delta}$.

Proof of Theorem 2.6.3

The principal's payoff from eliciting low effort is independent of a. Define $B_H^i(z_i^*) = \int_{z_i}^1 [\mu\eta_G + (1-\mu)\eta_B] a\mu^{a-1} d\mu$ as the principal's expected benefit from eliciting high effort at the optimal z_i^*, and $B_L(z^*) = \int_{z^*}^1 [\mu\eta_G + (1-\mu)\eta_B] d\mu$ as the benefit from eliciting low effort and using the first-best cutoff z^*. Take the limit as a decreases to one, $a \to 1$, then the principal's total payoff from low effort is $B_L(z^*) - c_L$. The principal's payoff from high effort is $B_H(z_i^*) - K_H^i(z_i^*)$. In the limit, $\lim_{a\to 1} K_H^i(z_i^*) = \infty$, the principal is unable to distinguish between high and low effort, making incentives for high effort infinitely costly, implying the principal prefers low effort, $B_H(z_i^*) - K_H^i(z_i^*) < B_L(z^*) - c_L$. And as $a \to \infty$, $B_H(z_i^*) - K_H^i(z_i^*) > B_L(z^*) - c_L$ in the limit. Since $B_H(z_i^*) - K_H^i(z_i^*)$ is continuous in a and $B_L(z^*) - c_L$ is independent of a, by the intermediate value theorem, there exists $a^* \in (1, \infty)$ such that the result holds.

Proof of Theorem 2.6.2

When $\kappa = 0$, $K_H^c > K_H^d$ for $\delta > \hat{\delta}$. When $\kappa = 1$, for all $\delta \in (0,1)$, $K_H^c < K_H^d$ and at $\delta = 1$, $K_H^c = K_H^d$. Since K_H^c is monotone and continuous, there exists $\kappa^*(\delta)$ such that $K_H^c(\kappa^*(\delta)) = K_H^d(\delta)$. When $\kappa > \kappa^*(\delta)$, $K_H^c < K_H^d$, so $z_c^* > z_d^*$.

Proof of Securitization

First consider the case high effort is elicited. The first order condition of the principal's payoff is

$$-f_H(z)[\eta(z) - \zeta(\tilde{R}(z) - p)] - K_H'(z) = 0$$

Denote the left hand side as $g(z,\zeta)$. Using the implicit function theorem, $\frac{\partial \zeta}{\partial z} = \frac{-\frac{\partial g(z,\zeta)}{\partial z}}{\frac{\partial g(z,\zeta)}{\partial \zeta}}$,

$$\frac{\partial \zeta}{\partial z} = \left(\frac{az^{a-2}\left((a-1)[\eta(z) - \zeta(\tilde{R}(z) - p)] + z[(\gamma_G - \gamma_B)R_h(1-\zeta)]\right) + K''(z)}{az^{a-1}[\tilde{R}(z) - p]} \right)$$

where $\frac{\partial g(z,\zeta)}{\partial z} = -a(a-1)z^{a-2}[\eta(z) - \zeta(\tilde{R}(z) - p)] - az^{a-1}[(\gamma_G - \gamma_B)R_h(1-\zeta)] - K''(z)$ and $\frac{\partial g(z,\zeta)}{\partial \zeta} = az^{a-1}[\tilde{R}(z) - p]$. The numerator is positive since $a > 1$, $K''(z) > 0$ and $\zeta \in [0,1]$. By the claim below, the denominator is negative since $\tilde{R}(z) < p$ and the result follows. If low effort is elicited, the first order condition is

$$-[\eta(z) - \zeta(\tilde{R}(z) - p)] = 0$$

Using the method above, $\frac{\partial \zeta}{\partial z} < 0$.

Claim A.0.1

$$\tilde{R}(z) < p$$

Suppose $(1+\Gamma)\tilde{R}(z) \leq (1+\Gamma)E[\tilde{R}(\mu)|\mu > z]$. Then

$$\begin{aligned}
\tilde{R}(z) &\leq E[\tilde{R}(\mu)|\mu > z] \\
\tilde{R}(z) &\leq \left(\frac{1}{1-F(z)}\right)\int_z^1 \tilde{R}(\mu)dF \\
\tilde{R}(z) - F(z)\tilde{R}(z) &\leq (\gamma_G - \gamma_B)R_h \int_z^1 \mu dF + \gamma_B R_h \int_z^1 \mu dF \\
\tilde{R}(z) &\leq (\gamma_G - \gamma_B)R_h \left[zF(z) + \int_z^1 \mu dF\right] + \gamma_B R_h \\
z(\gamma_G - \gamma_B)R_h + \gamma_B R_h &\leq (\gamma_G - \gamma_B)R_h E[\max(\mu, z)] + \gamma_B R_h \\
z(\gamma_G - \gamma_B) &\leq (\gamma_G - \gamma_B)E[\max(\mu, z)] \\
z &\leq E[\max(\mu, z)]
\end{aligned}$$

Since $\Gamma > 0$, $\tilde{R}(z) < (1+\Gamma)\tilde{R}(z) \leq p$ so $\tilde{R}(z) < p$.

Appendix B

Proofs from Chapter 3

Derivation of Probabilities

Here, I derive society's beliefs about the court's ability given the court's reporting strategy. The probability the court is high ability is

$$\hat{\mu}(H|r_1 = r_2) = Pr(H|s_1 = s_2)Pr(s_1 = s_2|r_1 = r_2) \tag{B.1}$$

The first term on the right hand side is

$$Pr(H|s_1 = s_2) = \frac{\mu}{\mu + (1-\mu)p_L}$$

The second conditional probability

$$Pr(s_1 = s_2|r_1 = r_2) = \frac{Pr(s_1 = s_2)Pr(r_1 = r_2|s_1 = s_2)}{Pr(r_1 = r_2)} \tag{B.2}$$

Then

$$Pr(r_1 = r_2) = Pr(s_1 = s_2)Pr(r_1 = r_2|s_1 = s_2) + Pr(s_1 \neq s_2)Pr(r_1 = r_2|s_1 \neq s_2) \tag{B.3}$$

Then (B.3) becomes

$$Pr(r_1 = r_2) = [\mu + (1-\mu)p_L]F(\bar{Q}) + [(1-\mu)(1-p_L)](1 - F(\underline{Q}))$$

And so (B.2) becomes

$$Pr(s_1 = s_2 | r_1 = r_2) = \frac{[\mu + (1-\mu)p_L]F(\bar{Q})}{[\mu + (1-\mu)p_L]F(\bar{Q}) + (1-\mu)(1-p_L)(1 - F(\underline{Q}))} \tag{B.4}$$

The market's belief about the court's ability using (B.1) and (B.4) provide the beliefs stated in the text.

Main Result Proofs

To understand how society's belief about court ability changes as the court's preferences change, it is necessary to prove the following results. The following assumption dictates how quickly f converges to its limit and is used for welfare comparisons.

Assumption 6 $\lim_{x\to\infty} f(x+k)/f(x) = 0$ *where* $k > 0$ *is a constant.*

Lemma B.0.1

$$\lim_{\lambda\to\infty} \hat{\mu}(H|r_1, r_2) = \mu$$

Suppose not, then as $\lambda \to \infty$ the market updates its belief about court ability, $\hat{\mu}(H|r_1 = r_2) \neq \hat{\mu}(H|r_1 \neq r_2)$. But as $\lambda \to \infty$ both consistent and inconsistent courts follow the same strategy: report $r_1 = r_2$. This means the market observes consistent reports from both types, so it cannot update its belief about court ability.

Lemma B.0.2

$$\lim_{\lambda\to\infty} \lambda[\hat{\mu}(H|r_1 = r_2) - \hat{\mu}(H|r_1 \neq r_2)] = \infty$$

Suppose not for contradiction, instead suppose $\lambda[\hat{\mu}(H|r_1 = r_2) - \hat{\mu}(H|r_1 \neq r_2)] = 0$. Then $\lim_{\lambda\to\infty} \bar{Q} = \alpha W$. Since $\lambda \to \infty$, the value to consistent and inconsistent courts from reporting $r_1 = r_2$ increases. But as the value of repeating reports increases, since $\lim_{\lambda\to\infty} \bar{Q} = \alpha W$, the $\lim_{\lambda\to\infty} Pr(r_1 \neq r_2|s_1 = s_2) = \lim_{\lambda\to\infty} 1 - F(\bar{Q})$ is increasing. This means as a consistent court cares more about his reputation, he reports $r_1 \neq r_2$ more. Contradiction.

Lemma B.0.3

$$\lim_{\lambda\to 0} \lambda[\hat{\mu}(H|r_1 = r_2) - \hat{\mu}(H|r_1 \neq r_2)] = 0$$

Since the probabilities $\hat{\mu}(H|r_1 = r_2)$ and $\hat{\mu}(H|r_1 \neq r_2)$ are bounded by zero and one, as $\lambda \to 0$, the product $\lambda[\hat{\mu}(H|r_1 = r_2) - \hat{\mu}(H|r_1 \neq r_2)] = 0$ does as well.

Proof of Proposition 3.3.1

By Lemma B.0.3 it can be shown that $\lim_{\lambda\to 0} E[V_{12}] = E[V_2]$. To see that $E[V_{12}]$ converges to $E[V_2]$ from above, the first order condition is

$$\frac{\partial E[V_{12}]}{\partial \lambda} : f(\alpha W_2 + \lambda\Delta^*)Pr(s_1 = s_2) - f(\alpha W_2 - \lambda\Delta^*)Pr(s_1 \neq s_2)$$

where $\Delta^* = \hat{\mu}(H|r_1 = r_2) - \hat{\mu}(H|r_1 \neq r_2)$. Then

$$\lim_{\lambda\to 0} \frac{\partial E[V_{12}]}{\partial \lambda} = f(\alpha W_2)Pr(s_1 = s_2) - f(\alpha W_2)Pr(s_1 \neq s_2) > 0$$

It is the case that $Pr(s_1 = s_2) > Pr(s_1 \neq s_2)$ (when the court knows its ability the probabilities are $\mu\gamma_1^H + (1-\mu)\gamma_1^L > \mu(1-\gamma_1^H) + (1-\mu)(1-\gamma_1^L)$ and when it does not they are $\mu + (1-\mu)p_L > (1-\mu)(1-p_L)$). Then the derivative is positive implying $E[V_{12}]$ has a positive derivative when λ is sufficiently small. Thus, at $\lambda = 0$, $E[V_{12}] = E[V_2]$; and as λ begins to increase, $E[V_{12}]$ increases but $E[V_2]$ is constant. Hence, $E[V_{12}] > E[V_2]$ when λ is a small and positive.

Proof of Proposition 3.4.2

Let $E[V_{12}] = W_2Pr(r_2 = \theta|FG) + K_2(1 - Pr(r_2 \neq \theta|FG))$ where the probabilities are defined by (3.14) and $E[V_1] = W_1Pr(r_1 = \theta|AO) + K_1(1 - Pr(r_1 \neq \theta|AO))$ where $Pr(r_1 = \theta|AO) = \mu + (1-\mu)p_L$. Then $E[V_1] - E[V_{12}] > 0$ if

$$W_1Pr(r_1 = \theta|AO) - K_1(1 - Pr(r_1 = \theta|AO)) > W_2Pr(r_2 = \theta|FG) + K_2(1 - Pr(r_2 = \theta|FG))$$

Solving for W_1 yields the result.

Proof of Corollary 3.3.2

Let $W_1 = W_2 = W$. By Lemma B.0.2, as $\lim_{\lambda\to\infty} \lambda\Delta^* = \infty$ which implies $F(\bar{Q}) \to 1$ and $F(\underline{Q}) \to 0$. In the limit, $E[V_{12}] = E[V_1]$. To see that $E[V_{12}]$ decreases to $E[W_1]$ the first order condition is $\frac{\partial E[V_{12}]}{\partial \lambda} : f(\alpha W + \lambda\Delta^*)Pr(s_1 = s_2) - f(\alpha W - \lambda\Delta^*)Pr(s_1 \neq s_2)$. Then

$$\lim_{\lambda\to\infty} \frac{\partial E[V_{12}]}{\partial \lambda} = \lim_{\lambda\to\infty} f(\alpha W_2 + \lambda\Delta^*)Pr(s_1 = s_2) - \lim_{\lambda\to\infty} f(\alpha W_2 - \lambda\Delta^*)Pr(s_1 \neq s_2)$$

By Assumption 6, $f(\alpha W_2 + \lambda\Delta^*)$ converges to zero faster than $f(\alpha W_2 - \lambda\Delta^*)$. Since $Pr(s_1 = s_2) > Pr(s_1 \neq s_2)$, the derivative is negative for $\lambda \to \infty$ implying that $E[V_{12}]$ is decreasing to $E[V_1]$ in the limit.

Proof of Proposition 3.4.1

To determine if the court will report truthfully, incentive compatibility constraints can be checked for the full game. As I show below, the court reports its first period signal truthfully, so I will only present the court's second period reporting strategy given a truthful first period report. The analysis below is for high-ability courts, but the same logic applies to low-ability courts.

Case one: Suppose $s_1 = s_2$ and the bias "favors" reporting the signal s_2. Then the court will report $r_2 = s_2$ if

$$\gamma_2^H[\alpha W_2 + Q + \lambda\hat{\mu}(H|r_1 = r_2)] + (1-\gamma_2^H)[\alpha K_2 + Q + \lambda\hat{\mu}(H|r_1 = r_2)]$$
$$\geq \gamma_2^H[\alpha K_2 + \lambda\hat{\mu}(H|r_1 \neq r_2)] + (1-\gamma_2^H)[\alpha W_2 + \lambda\hat{\mu}(H|r_1 \neq r_2)]$$

Simplified this expression becomes

$$\alpha(W_2 - K_2)(2\gamma_2^H - 1) + \lambda\Delta^* + Q \geq 0 \tag{B.5}$$

which is always true since $W_2 > K_2$, $\gamma_2^H > 1/2$, $\Delta^* \geq 0$, and by definition $Q \geq 0$. When bias, reputation, and social welfare all work in favor of the court reporting $r_2 = s_2$, the court has no reason to lie in the second period. The more interesting case is when the court's bias works against its second period signal s_2.

Step two: Given $s_1 = s_2$ and the bias is not in favor of s_2, the court will report truthfully if

$$Q_c^H \equiv \alpha(W_2 - K_2)(2\gamma_2^H - 1) + \lambda\Delta^* \geq Q \tag{B.6}$$

In this case the court's reputational concerns and bias work in opposite directions. Reputation dominates the court's decision making only if the bias is sufficiently small.

Otherwise, if (B.6) is violated, the court reports $r_2 \neq s_2$.

Step three: Suppose the court receives inconsistent signals, $s_1 \neq s_2$ and the bias favors reporting s_2. The court reports $r_2 = s_2$ if

$$\alpha(W_2 - K_2)(2\gamma_2^H - 1) - \lambda\Delta^* + Q \geq 0 \tag{B.7}$$

Step four: If $s_1 \neq s_2$ and the bias is against reporting $r_2 = s_2$, the court reports $r_2 = s_2$ when

$$Q_i^H \equiv \alpha(W_2 - K_2)(2\gamma_2^H - 1) - \lambda\Delta^* \geq Q \tag{B.8}$$

To simplify the analysis, note the symmetry between (B.5) and (B.6) (similarly for (B.7) and (B.8)). Define $Q_c^H = \alpha(W_2 - K_2)(2\gamma_2^H - 1) + \lambda\Delta^*$ and rewrite (B.5) as $Q_c^H \geq -Q$ and (B.6) as $Q_c^H \geq Q$. Since F is symmetric, it is true that $F(Q_c^H) = 1 - F(-Q_c^H)$. The probability (B.5) is satisfied is $Pr(Q > -Q_c^H) = 1 - F(-Q_c^H) = F(Q_c^H)$ and the probability (B.6) is satisfied is $Pr(Q < Q_c^H) = F(Q_c^H)$. Hence both are satisfied with the same probability. The same argument can be made with constraints (B.7) and (B.8). This means we only need to consider two constraints for the high-ability court, (B.6) and (B.8) (if (B.6) is satisfied, so is (B.5); if (B.8) is satisfied, so is (B.7)). The same analysis as above can be done with low-ability courts as well.

Proof of Theorem 3.3.3

The proof of Proposition 3.3.4 and Theorem 3.3.3 follow from Ferreira and Rezende (2007). First consider the equilibrium in which consistent and inconsistent courts report differently. Then society believes $\hat{\mu}(H|r_1 = r_2) = \frac{\mu}{\mu+(1-\mu)p_L}$ and $Pr(L|r_1 \neq r_2) = 1$. By Proposition 3.3.4, consistent courts will report $r_1 = r_2$. An inconsistent court's payoff from $r_1 \neq r_2$ is

$$\alpha W + \lambda \hat{\mu}(H|r_1 \neq r_2)$$

and its payoff from $r_1 = r_2$ is

$$\alpha \cdot 0 + \lambda \hat{\mu}(H|r_1 = r_2)$$

Then a necessary condition for separation to occur is

$$\alpha W + \lambda \hat{\mu}(H|r_1 \neq r_2) \geq +\lambda \hat{\mu}(H|r_1 = r_2)$$

or

$$\alpha W \geq \lambda \left[\hat{\mu}(H|r_1 = r_2) - \hat{\mu}(H|r_1 \neq r_2)\right]$$

Because society believes $\hat{\mu}(H|r_1 \neq r_2) = 0$, this condition simplifies to

$$\alpha W \geq \lambda \hat{\mu}(H|r_1 = r_2) = \lambda \frac{\mu}{\mu + (1-\mu)p_L}$$

Now consider a consistent court's incentives, by Proposition 3.3.4 it will always report $r_1 = r_2$. Thus, a consistent court always has less incentive to deviate and report $r_1 \neq r_2$.

Since consistent courts always report $r_1 = r_2$, for a pooling equilibrium to exist, it must be that inconsistent courts also report $r_1 = r_2$. Because both types report identically, society cannot update its belief given $r_1 = r_2$, i.e., $\hat{\mu}(H|r_1 = r_2) = \mu$. Off-equilibrium, if $r_1 \neq r_2$ is observed, Cho and Kreps (1987)'s D1 Criterion implies society will believe the inconsistent court deviated. This means that $\hat{\mu}(L|r_1 \neq r_2) = 1$ if the event is ever observed.

To check incentives, the inconsistent court's payoff from $r_1 = r_2$ is

$$\alpha \cdot 0 + \lambda\hat{\mu}(H|r_1 = r_2)$$

and from $r_1 \neq r_2$

$$\alpha W + \lambda\hat{\mu}(H|r_1 \neq r_2)$$

Hence inconsistent courts will pool and report $r_1 = r_2$ only if

$$\lambda\hat{\mu}(H|r_1 = r_2) \geq \alpha W + \lambda\hat{\mu}(H|r_1 \neq r_2)$$

which simplifies due to the above arguments to

$$\lambda\mu \geq \alpha W$$

Again, if an inconsistent court has no incentive to deviate from this equilibrium, a consistent court has less incentive to do so. And by appealing to the D1 Criterion, the off-equilibrium belief is unique, meaning the pooling equilibrium is unique as well.

For $\alpha W \in \left(\lambda\mu, \frac{\lambda\mu}{\mu+(1-\mu)p_L}\right)$ there is a hybrid equilibrium where inconsistent courts report $r_1 = r_2$ with probability π and $r_2 \neq r_2$ with probability $1 - \pi$. Consistent courts report $r_1 = r_2$.

Clearly only the inconsistent court is willing to randomize his action. Suppose with probability π it plays $r_1 = r_2$. To mix, an inconsistent court must be indifferent between reporting $r_1 = r_2$ and $r_1 \neq r_2$,

$$\alpha W = \lambda \hat{\mu}(H|r_1 = r_2)$$

If $r_1 \neq r_2$ is observed, society knows a low ability court was responsible, $Pr(L|r_1 \neq r_2) = 1$. Because inconsistent courts report $r_1 \neq r_2$ with positive probability, society can update its belief given $r_1 \neq r_2$, i.e., $\hat{\mu}(H|r_1 = r_2) > \mu$. This implies that

$$\alpha W > \lambda \mu$$

The probability society assigns to consistent reports is $\hat{\mu}(H|r_1 = r_2) \cdot Pr(H|s_1 = s_2)$, or

$$\frac{p}{p + (1-p)\pi} Pr(H|s_1 = s_2)$$

From $\alpha W = \lambda \hat{\mu}(H|r_1 = r_2)$, it must be that

$$\alpha W = \frac{p}{p + (1-p)\pi} \lambda Pr(H|s_1 = s_2)$$

which implies

$$\pi = \left(\frac{p}{1-p}\right)\left(\frac{\lambda \hat{\mu}(H|r_1 = r_2)}{\alpha W} - 1\right)$$

For $\pi > 0$, it must be that $\alpha W \leq \lambda \hat{\mu}(H|r_1 = r_2) = \frac{\lambda \mu}{\mu + (1-\mu)p_L}$

Second, suppose $\alpha W \in \left(\lambda \mu, \frac{\lambda \mu}{\mu + (1-\mu)p_L}\right)$. I need to show there exists $\pi \in (0, 1)$. Clearly, $\pi > 0$ if the conditions stated in the lemma hold. But $\pi < 1$ only if

$$p[\lambda \hat{\mu}(H|r_1 = r_2) - \alpha W] < (1-p)\alpha W$$

or

$$\lambda\hat{\mu}(H|r_1 = r_2) < \alpha W$$

which is again true under the conditions of the lemma.

Appendix C

Additional Material from Chapter 4

Random Audits

In practice, it may be difficult for the IRS to infer the firm's transfer pricing methodology $\rho \in \{\rho_A, \rho_C\}$ from its tax filing without first investing time and effort. In this section, I consider the case in which the IRS cannot observe the firm's TPM $\rho \in \{\rho_A, \rho_C\}$ before choosing its audit strategy $\beta \in [0, 1]$.

The IRS's payoff when audits are random is denoted $W_{Random}(\beta; \alpha, \pi)$, which is equal to:

$$\beta \left(\pi \left[\alpha(T_C + F - c) + (1-\alpha)(T_C - c)\right] + (1-\pi)\left[\alpha(T_A - c) + (1-\alpha)(T_A - c)\right]\right)$$
$$+(1-\beta)\left(\pi \left[\alpha T_A + (1-\alpha) T_C\right] + (1-\pi)\left[\alpha T_A + (1-\alpha) T_C\right]\right)$$

With random auditing, the firm can no longer insulate itself from the risk of IRS audits by using TPM ρ_C. Since the IRS cannot condition β on ρ, even a firm that chooses the conservative TPM ρ_C will be audited with positive probability. The IRS's marginal benefit from auditing is positive only if $\pi > \bar{\pi}_{Random}$, where

$$\bar{\pi}_{Random} \equiv \frac{(T_C - T_A)(1-\alpha) + c}{T_C - T_A + \alpha F}$$

Where $\pi > \bar{\pi}_{Random}$ only if the firm chooses ρ_A with probability $\alpha > \bar{\alpha}_{Random}$, where

$$\bar{\alpha}_{Random} \equiv \frac{(T_C - T_A)(1-\pi) + c}{\pi F + T_C - T_A}$$

Then the IRS's audit strategy when it cannot observe the firm's TPM ρ is given by:

$$\beta_{Random}(\alpha, \pi) = \begin{cases} 1 & \text{if } \alpha > \bar{\alpha}_{Random} \\ [0,1] & \text{if } \alpha = \bar{\alpha}_{Random} \\ 0 & \text{if } \alpha < \bar{\alpha}_{Random}. \end{cases} \tag{C.1}$$

With random audits, the firm's audit subgame payoff $V_{Random}(\alpha; \pi, \beta)$ is

$$\begin{aligned} V_{Random} = {} & \alpha \left[\beta \left(\pi u(T_C + F + k) + (1-\pi)u(T_A + k)\right) + (1-\beta)u(T_A)\right] \\ & + (1-\alpha)\left[\beta \left(\pi u(T_C + k) + (1-\pi)u(T_A + k)\right) + (1-\beta)u(T_C)\right] \end{aligned}$$

The firm is indifferent between reporting ρ_A and ρ_C if the IRS chooses audit probability $\bar{\beta}_{Random}$,

$$\bar{\beta}_{Random} = \frac{u(T_C) - u(T_A)}{\pi \left[u(T_C + F + k) - u(T_C + k)\right] + u(T_C) - u(T_A)}$$

Naturally, the threshold $\bar{\beta}_{Random}$ is decreasing in π and F: as the IRS becomes more confident the firm is type τ_N and as the expected benefit from detecting a non-eligible firm increases, the firm chooses TPM ρ_C for a greater range of audit probabilities. The firm's strategy $\alpha(\beta_{Random}, \pi)$ is

$$\alpha(\beta_{Random}, \pi) = \begin{cases} 1 & \text{if } \beta > \bar{\beta}_{Random} \\ (0,1) & \text{if } \beta = \bar{\beta}_{Random} \\ 0 & \text{if } \beta < \bar{\beta}_{Random}. \end{cases} \tag{C.2}$$

Lemma C.0.4 (Random Audit Equilibrium) *In the audit subgame with random audits, the firm reports* ρ_A *with probability* $\bar{\alpha}_{Random}$ *and* ρ_C *with probability* $1-\bar{\alpha}_{Random}$. *The IRS audits with probability* $\bar{\beta}_{Random}$ *and does not audit with probability* $1-\beta_{Random}$.

Interestingly, unlike the main model where the audit subgame equilibrium involved only pure strategies, when audits are random the IRS and firm use mixed strategies. Since random audits do not affect the firm's payoff from the APA subgame, the firm's payoff from requesting an APA is given by (4.9) and the IRS's payoff is given by (4.10).

Proposition C.0.1 *Suppose IRS auditing is independent of* ρ. *Then,*

1. *if the firm's APA costs* $n + k'$ *are sufficiently less than its expected audit costs* $\beta\,(k + \alpha\pi F)$, *the firm requests an APA with probability* $z = 1$,

2. *if the firm's APA costs* $n + k'$ *are sufficiently greater than its expected audit costs* $\beta\,(k + \alpha\pi F)$, *the firm will forgo the APA (i.e.* $z = 0$*) and the equilibrium audit subgame strategies are given by Lemma C.0.4.*

3. *if the firm's APA costs* $n + k'$ *and expected audit costs* $\beta\,(k + \alpha\pi F)$ *are approximately the same, part one or two may hold depending on* $\pi \in [0, 1]$.

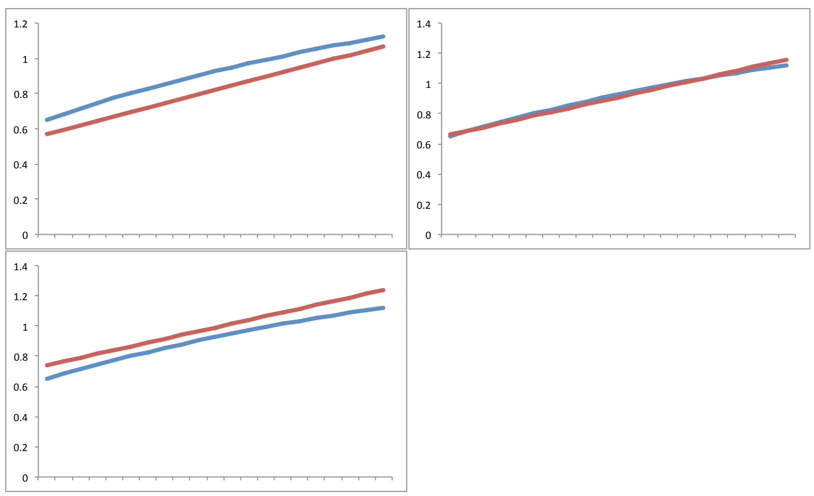

Figure C.1: The top left panel shows part one of Proposition C.0.1, when APA costs $n + k'$ are low relative to audit costs, the APA is cheaper for the firm for all π. The bottom left panel shows when the APA costs are large relative to audit costs, the firm forgoes the APA for all π. The top right panel shows when APA and audit costs are approximately the same.

This proposition shows how the behavior of the firm and IRS are sensitive to the costs of advance pricing agreements and ex-post auditing. When APA costs are low relative to the firm's expected cost of ex-post auditing, the firm benefits by taking advantage of the APA. The interesting case is when the costs of the APA and ex-post auditing are approximately equal. In numerical simulations, a tripartite strategy emerges for the firm. This can be seen by the upper right panel of Figure C.1. When π is sufficiently small (resp. large), the IRS and firm will coordinate their actions to avoid confrontation: the firm will choose ρ_A (resp. ρ_C) and the IRS will not audit (resp. $\beta = 1$). For moderate values of π, when uncertainty is greatest, the firm and IRS disagree about the firm's true type and hence the firm–wishing to avoid confrontation with the IRS–will opt for the APA. This is reminiscent of the equilibrium behavior in the full game when the IRS's audit strategy can be conditioned on ρ.