Applications of Information Design to Finance, Politics, and Education

A DISSERTATION
SUBMITTED TO THE FACULTY OF THE
UNIVERSITY OF MINNESOTA
BY

Isaac Douglas Swift

IN PARTIAL FULFILLMENT OF THE
REQUIREMENTS FOR THE DEGREE OF
Doctor of Philosophy

David Rahman, Advisor

July, 2020

Acknowledgements

I would first and foremost like to acknowledge my wife, Kim. None of this would have been possible without her. She is the reason I got through the hardest times. I will be forever grateful to her and all she has done for me.

I had the best advisor a graduate student could hope for in David Rahman. He was always supportive and encouraging. He was infinitely knowledgable and gave the best advice whenever I discussed my work with him. My only regret is not going to him for help more often.

I would like to thank the countless people who listened as I tried to explain my nonsensical thoughts and gave valuable feedback. They are the reason this work transformed from incoherent babble to semi-coherent research. In addition to seminar participants, I would like to thank Jan Werner, Aldo Rustichini, Martin Szydlowski, Beth Allen, Anmol Bhandari, Anton Kolotilin, Jaroslav Borovicka, Bruno Strulovici, Francoise Forges, Debraj Ray, Kim-Sau Chung, Vijay Krishna, Brennan Platt, Brian Albrecht, Amy Neff, and Jordan Pandolfo.

Finally, to the giants that went before me.

Dedicated to

Joslyn, Dorian, and Monroe

Abstract

Information design is the study of how an informed player optimally communicates to influence the actions of another player. In the first chapter, Ill give a brief review of a few main results from the literature that I will be using in this dissertation. In the remainder of this dissertation, I will present novel applications and extensions of the information design framework to solve problems in finance, politics, and education. In chapter 2, I study a trader that has private information relating to the value of an asset. The traders trades are publically observable, so their actions reveal information to the rest of the market. I find the optimal trading strategy and show that it rationalizes the commonly assumed price dynamics assumed in the finance literature. In chapter 3, I study the problem of gerrymandering. A state gets to elected several members of congress. The designer chooses how to divide the state into districts to each elect one congress member. I first show how the designer can maintain a congressional majority for their preferred political party, even when that party becomes the minority among the population. I then document in the data that state seat-vote curves (the fraction of seats a party wins as a function of the-fraction of votes the party wins in the state) are highly responsive to changes in the vote share, and that this responsiveness is negatively correlated with the size of the state. This is explained in my equilibrium model. In chapter 4, I study university grading schemes. A worker has no way to credibly signal their productivity to potential employers on their own, so they hired a university to evaluate them a publisha rating. The university needs to balance their desire to help as many students as possible get hired (so they can charge high tuition) with their need to build and maintain a good reputation. The model explains observed patterns of grade inflation, especially concentrated among

good schools. In chapter 5, I give a few concluding remarks and thoughts about future avenues for research.

Contents

List of Tables

List of Figures

Chapter 1

Information Design

Information design is the study of how an informed player optimally communicates to influence the actions of another player. Think of a lobbyist explaining their position to congress members, lawyers presenting evidence to a judge or jury, an applicant making their resume for future employers, or countless other everyday interactions. One player has private information. A second player needs to make some decision that will affect both players. The first player much choose what information they want to reveal and what they want to conceal. The two players may not have completely aligned incentives. Both players are rational Bayesians, so the second player knows the first player may have adverse incentives and reponds to messages appropriately.

You may wonder if any useful information can be transmitted when the informed player (sender) and the uninformed player (receiver) have different incentives. A large and recently rapidly growing field has arisen to show how useful information is optimally transmitted in such situations. In this chapter, I'll give a brief review of a few main results from the literature that I will be using in this dissertation. First I'll start with the situation where the sender can costlessly send any message they want to the receiver. These "cheap talk" games were studied initially in Crawford and Sobel (1982). I will then look at how the optimal strategy changes when the sender can commit themselves ex ante to a reporting strategy. These "Bayesian Persuasion" games have been very popular recently, after their introduction by Kamenica and Gentzkow (2011). After this, I will show how these games can be extended into a dynamic setting. These sections are all based on existing papers in the literature. In the final section of the chapter I will give an outline of the remainder of this dissertation, which is comprised of my original research.

1.1 Salesman Example

You walk into the store to do some shopping for the weekend. What you would like to buy depends on the weekend weather. If you buy and umbrella, you will get a payoff of one if it rains and a payoff of minus two if it doesn't rain (you carry it around all that time for nothing). If you buy a new grill the payoffs are the opposite, plus one if it's dry and minus two if it rains. You also have to option to not buy anything and get a payoff of zero. Unfortunately, you don't know whether it is going to rain this weekend or not. You optimal decision depends on your perceived probability of sun this weekend, μ. If $\mu < \frac{1}{3}$, you will buy the umbrella. If $\mu = \frac{1}{3}$, you are indifferent between buying the umbrella or not and willing to do any mixed strategy of the two. If $\frac{1}{3} < \mu < \frac{2}{3}$, you will buy nothing. If $\mu = \frac{2}{3}$, you are indifferent between buying the grill or not and willing to do any mixed strategy of the two. Finally if $\mu > \frac{2}{3}$, you will buy the grill.

Luckily, the salesman has seen the latest forcast, knows if it is going to rain or not, and can share any information he likes with you before you make your purchase. However, the salesman doesn't simply wish for your wellbeing. The salesman gets a commission if you make a purchase. Since the grill is more expensive, the salesman gets a bigger commission for selling the grill than the umbrella. We'll say the salesman's payoff is one if you buy the umbrella, two if you buy the grill, and zero if you don't buy anything. What information will the salesman communicate in equilibrium?

We can graph the salesman's payoff as a function of your beliefs at the time of your purchase.

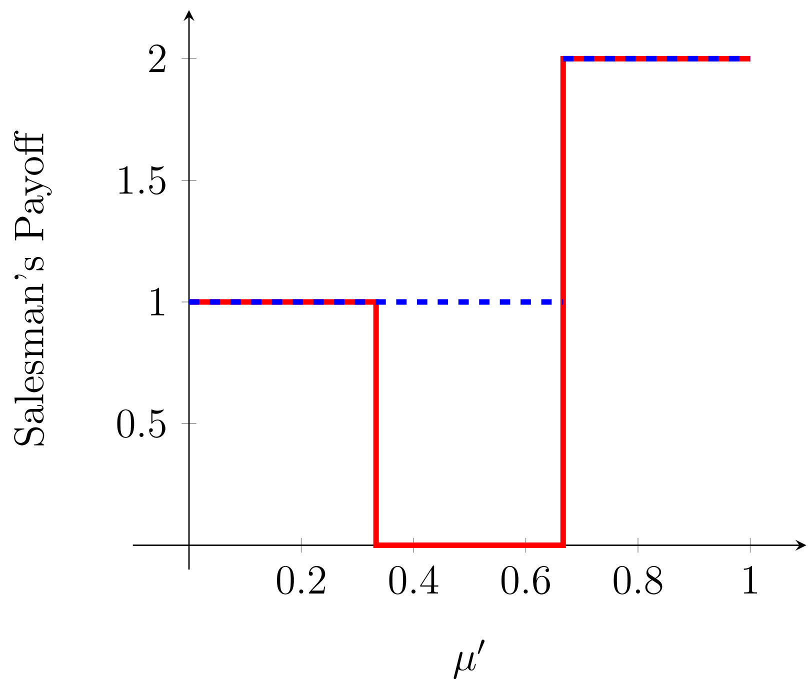

Figure 1.1: The solid line (blue) is the payoff the salesman receives as a function of the customer's beliefs at the time of purchase. The dashed line (red) is the maximum expected payoff the salesman can receive in equilibrium as a function of the customer's prior belief.

If your prior belief when you walk into the store is already above $\frac{2}{3}$, the salesman has no need to communicate anything to you about the weather. You are already going to buy the item with the highest commission without any further information. It will be more obvious why in a minute, but if your beliefs are below $\frac{1}{3}$ the salesman will again communicate nothing meaningful in any equilibrium and you will buy the umbrella. The interesting case is when your prior beliefs are strictly between $\frac{1}{3}$ and $\frac{2}{3}$. There are many different equilibria of this game. The one that gives the lowest payoff to the salesman is one where no information is conveyed. You think that anything the salesman says will not convey meaningful information, therefore regardless of what the salesman says, your beliefs are unchanged and you buy nothing. Since the salesman's words don't have any effect the salesman optimally chooses not to say anything useful. The salesman gets a payoff of zero in this equilibrium.

The Perfect Bayesian Equilibrium with the highest payoff for the salesman looks quite different. Take the prior belief to be $\frac{1}{2}$. Suppose that the salesman is playing a mixed strategy. If it is going to be sunny, the salesman will say it is going to be sunny. On the other hand, if it is going to be rainy the salesman will say it is going to be rainy $\frac{1}{2}$ of the time and say it is going to be sunny the other $\frac{1}{2}$ of the time.

How will you optimally respond to the salesman's information? When the salesman tells you it is going to be rainy, you know they are telling the truth and you buy the umbrella. When the salesman tells you if is going to be sunny, you use Bayes rule to get an updated probability of sun equal to $\frac{2}{3}$. At this belief you are just indifferent between buying the grill, buying nothing, or any mixed strategy of the two. Suppose you decide to mix fifty-fifty between buying the grill or not (which is in your best response).

Let us look at the salesman's payoffs. If the salesman says it is going to be rainy, the customer buys an umbrella and the salesman gets a payoff of one. If the salesman says it is going to be sunny, the customer buys the grill half the time and buys nothing the other half the time. This gives an expected payoff of one. This means that as long as any other messages are interpretted as uninformative, the salesman's best response is to play any mixed strategy of saying sunny or rainy. This includes the strategy written above. This means that the above strategies and beliefs constitute an equilibrium.

It also must be the case that this is the highest payoff that can be obtained in equilibrium. To get a higher payoff, there must be at least one message to which the customer responds by buying the grill more than fifty percent of the time. Since belief must be a martingale, there is necessarily a message that makes beliefs go down from $\frac{1}{2}$ and earns the salesman a payoff less than or equal to one. Clearly, when it is time to give the message to make beliefs fall the salesman could profitably deviate by sendig the message to make the customer buy the grill. Thus, this is not an equilibrium.

In Lipnowski and Ravid (ming) the authors show an easy way to find the payoff maximizing equilibrium of games like this one. They say that the sender has "transparent motives" if the sender payoff does not depend on the state of the world. This is like how the salesman's commission doesn't depend on if it rains or not. This will be the case in each of the examples in this dissertation. The theorem in Lipnowski and Ravid (ming) shows that the highest payoff the sender can obtain in an equilibrium of a game with transparent motives (as a function of the prior) is equal to the quasi-concave envelope of the sender's payoff as a function of the posteriors. If you graph the sender's payoff against the receiver's posterior, the sender's maximum equilibrium payoff is equal to the lowest quasi-concave function that is above that line. In 1.1 this is the dashed line. This concept will be used again in chapter 2.

1.2 Commitment

Suppose now that the salesman can commit ex ante to a message strategy. For example, they could commit to always telling the truth or any mixed strategy depending on the state. What strategy should they commit to?

Perhaps, commitment doesn't seem like a reasonable assumption in this situation. You can just interpret it differently. Maybe they can't commit to always telling the truth, but they could have the weather forecast playing on the store tv for the customer to see. Imagine that there are many available weather forecasts available to play on the tv with varying levels of accuracy at predicting rain or sun. Another way to interpret the problem is to ask, which is the optimal weather forecast station to have playing in the store.

The salesman could commit to always truthfully revealing the state. This means that beliefs will always move to zero or one after the message and the customer will always make a purchase. The expected payoff to the salesman is then $1 + \mu$. If $\mu \in [0, \frac{2}{3})$ this is better than the best payoff the salesman got in any equilibrium without commitment. If $\mu \in [\frac{2}{3}, 1]$ this is worse than the best equilibrium payoff without commitment. This is not yet the optimal strategy the salesman can commit to.

The optimal strategy is as follows. As before if $\mu \in [\frac{2}{3}, 1]$, the customer is already going to purchase the more expensive item. So, providing any information can not lead to an improvement. If μ is anything lower than $\frac{2}{3}$ the optimal strategy is a mixed strategy that is dependent on the state. If it is going to be sunny, the salesman says it will be sunny. If it is going to be rainy, the salesman says it will be sunny $\frac{1}{2}\frac{\mu}{1-\mu}$ fraction of the time and says it will be rainythe rest of the time. When the salesman says it will be rainy, the customer believes him and buys the umbrella (giving the salesman a payoff of one). When the salesman says it will be sunny, the customer uses Bayes rule to arrive at a posterior beliefs equal to $\frac{2}{3}$. This is just high enough to convince them to buy the grill (giving the salesman a payoff of two). This means that the salesman's expected payoff when the prior is below $\frac{2}{3}$ is equal to $1 + \frac{3}{2}\mu$. This is greater than the payoff from always telling the truth or any payoff obtained in equilibrium without commitment.

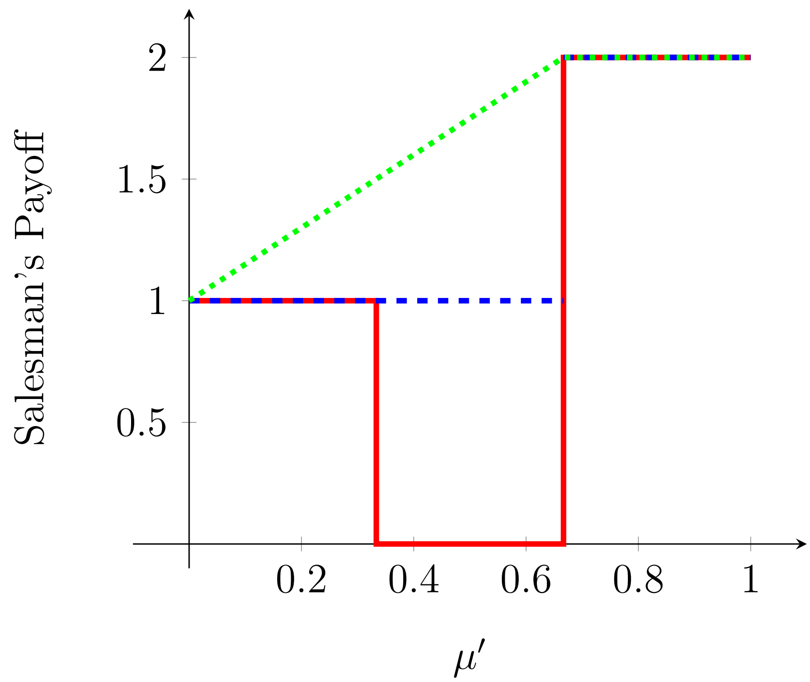

Figure 1.2: The dotted line (green) shows the highes payoff the salesman can obtain with commitment.

In Kamenica and Gentzkow (2011) the authors show the simple way to solve games such as this one. They show first that you can focus simply on the distribution of posterior beliefs that is induced rather than on the messages used. This is possible because they show that every distribution of posteriors that averages out to the prior can be induced by some messaging strategy. This means that if you look at the graph of the sender's payoff as a function of the receiver's posterior, you can choose any points on the graph (as long as some are one opposite sides of the prior) and the sender can get a payoff equal to the weighted average of those points. The weights are the weights needed to make the posteriors average to the prior. The highest of these payoffs can be seen in an intuitive graphical way. The sender's highest expected payoff as a function of the prior is simply the "concavification" of the sender's payoff as a function of the receiver's posterior. This means that it is the lowest concave function that lies weakly above the payoff function.

This "concavification" technique will be used in chapters 3 and 4.

1.3 Dynamics

In Ely (2017) the author shows how Bayesian Persuasion games can be extended into a dynamic setting. Call $u(\mu)$ the payoff the sender gets as a function of the receiver's posterior belief as before. Now suppose that there are many periods. Every period, the sender sends a message, the receiver takes an action, and the sender gets payoff $u(\mu_t)$ with a discount factor δ. Suppose the state of the world is not changing each period. Ely (2017) first proves the "Obfuscation Principle" which shows that any stochastic process (μ_t) such that each period the posterior averages to the prior ($\mathbb{E}[\mu_{t+1} \mid \mu_t] = \mu_t$) can be induced by some dynamic messaging strategy. Using this they show that a "concavification" can be used with the standard dynamic programming approach to solving dynamic problems.

$$V = cav\left[u + \delta V\right] \tag{1.1}$$

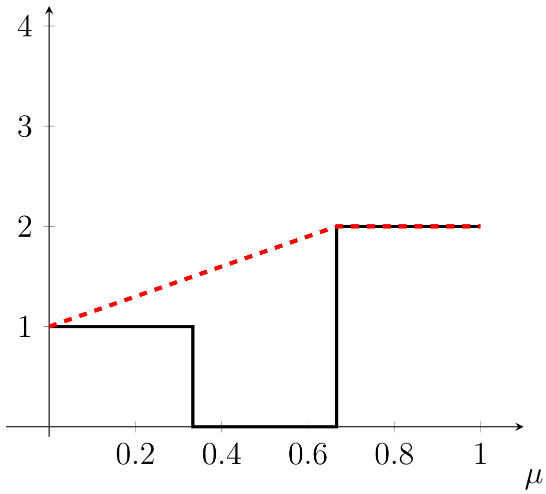

(a) Take an initial guess $V_0 = 0$. The solid line is $u(\mu')$. The dashed line is the value function after the first iteration V_1.

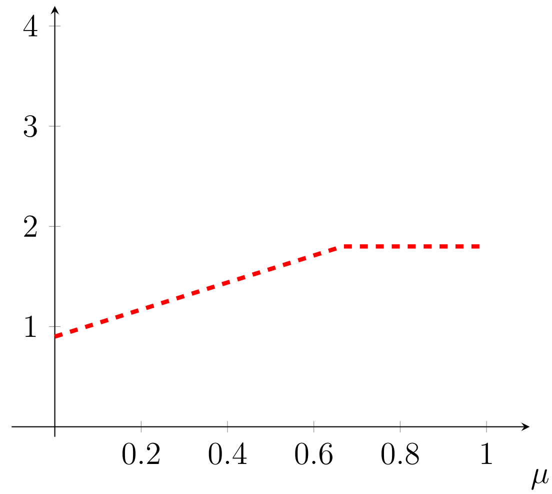

(b) The discounted value function after one iteration, δV_1

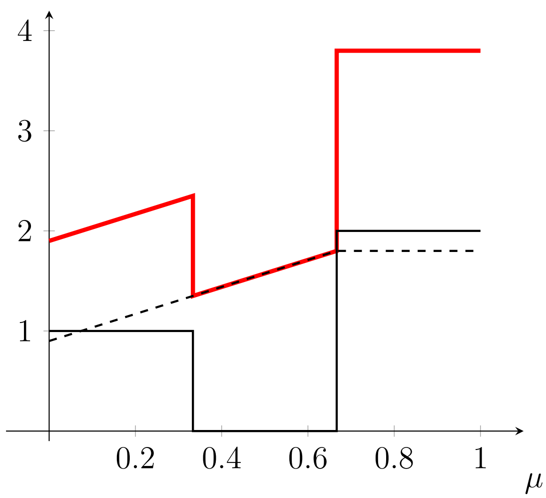

(c) The thin lines are u and δV_1. The thick line is $u + \delta V_1$.

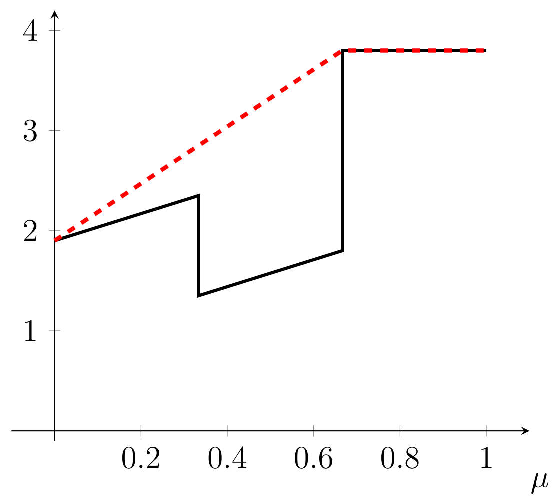

(d) The new value function after the second iteration V_2.

These principles will be used in chapter 2.

1.4 Outline of Dissertation

While there is an extensive literature within micro theory extending these methods and working on information problems, there have been few papers with serious applications of information design to other fields and situations without communication between players. In the remainder of this dissertation, I will present novel applications and extensions of the information design framework to solve problems in finance, politics, and education. In chapter 2, I study a trader that has private information relating to the value of an asset. The trader's trades are publically observable, so their actions reveal information to the rest of the market. I find the optimal trading strategy and show that it rationalizes the commonly assumed price dynamics assumed in the finance literature. In chapter 3, I study the problem of gerrymandering. A state gets to elected several members of congress. The designer chooses how to divide the state into districts to each elect one congress member. I first show how the designer can maintain a congressional majority for their preferred political party, even when that party becomes the minority among the population. I then document in the data that state seat-vote curves (the fraction of seats a party wins as a function of the fraction of votes the party wins in the state) are highly responsive to changes in the vote share, and that this responsiveness is negatively correlated with the size of the state. This is explained in my equilibrium model. In chapter 4, I study university grading schemes. A worker has no way to credibly signal their productivity to potential employers on their own, so they hired a university to evaluate them a publish a rating. The university needs to balance their desire to help as many students as possible get hired (so they can charge high tuition) with their need to build and maintain a good reputation. The model explains observed patterns of grade inflation, especially concentrated among good schools. In chapter 5, I give a few concluding remarks and thoughts about future avenues for research.

Chapter 2

Dynamics of Price Discovery

Abstract

Changes in the price of a financial asset represent learning as the market updates its expectation about fundamentals. In this paper I characterize what price dynamics are possible when the information is being released strategically by a profit maximizing trader and market participants are Bayesian. I study how information is incorporated into prices over time in model with general trading strategies that allow for the spread of false information and price manipulation. Every period an informed trader reveals their information by buying or selling an asset. After observing the trade, beliefs and prices are updated. The informed trader's prefered equilibrium is characterized with and without commitment leading to starkly different results. Regardless of how beliefs impact prices, the optimal strategy for the informed trader is to release their information gradually mixed with a nearly equal amount of misinformation. This strategy leads to volatile price paths that bounce back and forth each period. In the continuous time limit, the price process converges to a Brownian motion. Moving prices back and forth in this way hinges critically on the informed traders ability to commit ex ante to a strategy. Without such commitment power, the optimal strategy is to release nearly all information suddenly at randomized times. The optimum resembles a pump-and-dump price manipulation scheme and can lead to sudden crashes or spikes in the price of the asset. In the limit, the price converges to a Poisson process. This paper gives a micro-foundation to price processes commonly assumed in the literature.

JEL Classification: D82, D83, G14

2.1 Introduction

In an efficient well functioning financial market an asset's price reflects the market's beliefs about fundamentals. Dynamics of prices are then driven by learning as the market updates its expectation of these fundamentals. Does supposing the market to be rational and use Bayesian updating put any restriction on the dynamics we would expect to see in asset prices? Going even further, suppose that the new information that leads the the market learning doesn't just arrive exogenously. If the new informataion is being released by a strategic trader maximizing their profit, can we infer anything about what price dynamics should be possible? In fact, the informed agent behaving optimally along with the market being Bayesian implies a strict characterization of possible price dynamics. The subject of this paper is deriving these dynamics.

In this paper I use a simple model to study how information gets incorporated into prices. Begin with a trader holding some information not commonly known to the rest of the market. As the informed trader acts over time they reveal something about their information. The market price will continuously adjust as the market updates its beliefs to reflect the information revealed. I study the dynamics of prices through this process.

To understand how prices move from revealed information, we need to understand what information is being revealed by trades. We must first know how an informed trader optimally incorporates their information into their trading strategy. The informed trader buys or sells an asset primarily based on what they think the market price will do in the future. How the market price responds to their trade is of fundamental importance in creating the optimal trading stratey. Thus, the price process and optimal trading strategy are equilibrium objects that must be determined jointly.

I will use a simpler model of market structure to allow for more general trading strategies and a few incentives that are typically ignored in the literature. If the informed agent's trades can move prices, there will always be an incentive to manipulate those prices. The informed trader may potentially choose to buy the asset today, try to convince people the state is good to drive up the price, then just sell it tomorrow. This paper shows how rumors and misinformation can interact with incentives to reveal true information and how much it can hinder price discovery.

In an efficient market, prices change only as new information is incorporated. Observed stock prices jump around continually throughout the trading day, seemingly too much the be justified by new information alone (see Shiller (1981) among others). In this paper, I find the maximum amount of variation that can be caused solely by information.

I address these questions using a dynamic information disclosure game in which information is the only factor moving prices. Each period an informed trader buys or sells an asset and their actions reveal their information to the market maker over time. The informed trader profits from having a long position when the price goes up and a short position when the price goes down. I then take the limit to go to continuous time and characterize the stochastic process the asset price follows.

I find that regardless on how the information affects the asset's price, the optimal way for the informed trader to release the information is Brownian. The informed trader gradually releases the information by giving small pieces of true or false news with nearly equal frequency each period. This leads to prices following an Itó process. The asset has continuous but spiky price paths. This leads to the maximal amount of variation over time. Only asymptotically is all information revealed.

However, implimenting this policy requires a strong amount of commitment on the side of the informed trader. The informed trader must commit to a distribution of actions before observing the state. This is the same type of commitment used in the Bayesian persuasion literature. Without this commitment ability, the optimal policy looks very different. The informed agent is no longer able to break up the information and release it gradually in a profitable way, but they retain a lot of control over the timing of information release. The optimal policy becomes a Poisson process with drift that looks a lot like a pump-and-dump scheme. This leads to a price path with sudden discontinuous crashes or jumps.

The Itó and Poisson processes arise completely endogenously as the optimal strategy of an informed trader. Nothing in the setup of the model is Normal or Poisson or game to lead to these distributions. Throughout finance and particularly option pricing, an Itó process possibly with a Poisson process added on form the dynamics of a stock price. This is typically assumed for tractability, but my model gives a micro-foundation for these to arise as the natural stochastic process for prices. In papers of learning and information aquisition the arrival of information is typically modeled to be either Brownian or Poisson. My model shows that such information flows may be motivated by the information being obtained from a strategic player.

Example

Consider a very simple example that will allow me to fix ideas and explain the general concepts of the model. While the model can be much more general than this example, the intuition is similar.

There is an asset with a payoff, ω, equal to either zero or one with equal probability. This asset is just a one dollar bet. This payoff is received once at the end of the game. Suppose further that this asset is highly liquid, all agents are risk neutral, and there is no discounting. What I mean by these is that at any time, you can walk down to the market and buy or sell a share of this asset at the posted price, which is equal to its expected value. Since either state is equally likely, the initial price is $\frac{1}{2}$.

If a trader privately learns what the payoff is going to be, how can she best profit from this information? For simplicity, say that there are two trading periods before the asset payoffs and that the trader can take a long or short position, but faces a capacity constraint of one share (holdings $x \in [-1, 1]$). The obvious candidate strategy is that if she learns the asset's payoff will be high ($\omega = 1$) she should buy the asset and if she learns that the asset's payoff will be low ($\omega = 0$) she should sell the asset. This strategy would earn her a total payoff of $\frac{1}{2}$. If she realizes the asset is good, she will buy it in the first period for a price of $\frac{1}{2}$. She then doesn't need to do anything in the second period because she is already holding her max amount. Then the payoffs from the asset realize and she gets 1. If the asset is actually bad, her payoffs are just the mirror. She sells the asset in the first period for a gain of $\frac{1}{2}$, then at the end of the game doesn't need to pay back anything ($\omega = 0$). Her payoff is the same in either state.

Now I'm going to propose a candidate strategy that can do better than simply buying if the asset is good and selling if it is bad. In this strategy the trader will utilize the fact that there are multiple periods of trading by misleading the market in the first period and manipulating prices to get a higher return in the second period. Consider the following randomized strategy. In the first period if the trader observes the asset is good, she will buy with probability $\frac{3}{4}$ and sell with probability $\frac{1}{4}$. If she observes the asset is bad she will do the opposite, buy with probability $\frac{1}{4}$ and sell with probability $\frac{3}{4}$. Then in the second period she will do the simple strategy of holding a long position if the asset is good and a short position if the asset is bad, just like the previously proposed strategy above.

This strategy will give a higher payoff. To see this, we first need to know what the prices will be in the second period. Since the price is always equal to the expected payoff, this is computed using Bayes rule. If she buys the asset in the first period, the price in the second period will be $\frac{\frac{1}{2}\frac{3}{4}}{\frac{1}{2}\frac{3}{4}+\frac{1}{2}\frac{1}{4}} = \frac{3}{4}$. If she sells in the first period, the price move the opposite way, $\frac{\frac{1}{2}\frac{1}{4}}{\frac{1}{2}\frac{1}{4}+\frac{1}{2}\frac{3}{4}} = \frac{1}{4}$.

Let's now compute the trader's expected payoff. First take the good state ($\omega = 1$). Three quarters of the time she will buy in the first period at a price of one half, in the secon period maintains that long position with no cost, then receives a payment of one at the end of the game. The other one quarter of the time she will sell at a price of one half in the first period, buy at a price of one quarter in the second period, and receive a payoff of one at the end of the game. This gives an expected payoff of

$$\frac{3}{4}\left(-\frac{1}{2}+1\right)+\frac{1}{4}\left(\frac{1}{2}-2\frac{1}{4}+1\right)=\frac{5}{8}$$

which is larger than $\frac{1}{2}$. If the state is good we can get the payoff with a similar calculation.

$$\frac{3}{4}\left(\frac{1}{2}-0\right)+\frac{1}{4}\left(-\frac{1}{2}+2\frac{3}{4}-0\right)=\frac{5}{8}$$

Most of the time the trader does the usual strategy of buying when the asset is good and selling when it is bad to get a payoff of $\frac{1}{2}$. Occasionally, she trades opposite of her information in the first period. This allows her to transact at a more favorable price in the second period. This manipulation earns a higher payoff whether the asset is good or bad.

If there were three periods of trading, the trader would be able to manipulate prices for two periods before taking the obvious trade in the last period. If there are many trading periods, the optimal strategy involves the trader potentially moving prices back and forth between higher and lower levels many times before the end. In fact, in the limit as you took an infinite number of trading periods the price would approach a Brownian motion. Every period she would choose to buy or sell with nearly equal probability. This will cause the price to continue to bounce up or down by infinitesimal amounts.

In turns out that $\frac{5}{8}$ is the highest payout the trader can garuntee herself in this two period game. Notice, however, that acheiving this requires a strong amount of commitment on the part of the trader. If not committed ex ante, the trader has an incentive to deviate from the outlined strategy some of the time. When the asset is good, the trader mixes between buying and selling in the first state. When she buys in the first stage the trader gets a payoff of $\frac{1}{2}$, but when she sells in the first stage the trader gets a payoff of 1. Recall that the market doesn't obvserve the state (ω), only the trade. Hold fixed the market prices($p_1 = \frac{1}{2}$, $p_2 = \frac{3}{4}$ if buy is observed in the first period, and $p_2 = \frac{1}{4}$ if sell is observed in the first period). When the asset is good, the trader will always prefer to sell and manipulate the price because it gives a higher payoff. However, the manipulation strategy was able to effectively move prices simply because it is done infrequently.

Thus, without commitment power this equilibrium would fall apart completely. The best the trader can do in an equilibrium without commitment is exactly the simple strategy of buying if the asset is good and selling if the asset is bad. This is done in one period (the last one). In the more general form of the model with infinite periods, it is still the case that the best the trader can do without commitment is to reveal nearly all their information at once. The trader still will maintain a lot of power over the timing of this information dump. The profit is maximized by randomizing of the timing of the information release. In the continuous time limit the price will converge to a Poisson process.

2.2 Model

In this section, I present the model in a much more general form than the example, but I still strive for simplicity. I've stripped away all aspects that distract from the main result. To see how it generalizes or how it relates to more standard models, see sections five and six respectively. What is essential for the results is simply a strategic informed trader that profits from price changes and a rule for how beliefs are converted to prices. I'll first solve the model here without commitment. After examining the solution briefly, I take the continuous time limit and arrive at a Poisson process. Then in section four I will solve the model with commitment. Again the continuous time limit will be studied and a Brownian motion will arise. In section five, I will give extensions of this model.

There is a permanent state of nature that takes one of two possible values, $\omega \in \{0, 1\}$. I'll sometimes refer to the state as the asset being bad or good. The probability that $\omega = 1$ is denoted by μ. There is a long lived asset whose value is affected by the state.

There are discrete time periods $t = 1, 2, \ldots, T$ and usually T will be taken to be infinite. There are two players. Each period the market maker chooses a price at which they are willing to buy or sell. Then the informed trader that knows the value of ω chooses to buy or sell the asset.

Informed Trader

The informed trader is risk neutral chooses how much of the asset to hold each period, x_t, to maximize expected discounted profit. Because the objective is linear, a bang-bang solution obtains and the insider would like to hold a infinite number of shares of the asset. I restrict the number of shares that can be held to $x_t \in [-1, 1]$. When T is finite, I will force $x_T = 0$ so that the payoffs to the insider don't include any final value of holding the asset at the end of the game.

The payoff to the trader is the following.

$$V(\mu_0) = \mathbb{E}\left[\int_0^T e^{-rt} x_t dp_t\right] \tag{2.1}$$

Though the main results are continuous processes, the model is in discrete time and the main theorems are about the limits as the time periods get smaller. This means that through most of the paper, I will be using a discrete approximation to this payoff function. Call $\delta = e^{-r\Delta t}$. The discrete payoff used is

$$V(\mu_0) = \mathbb{E}\left[\sum_{t=0}^{T} \delta^t (p_{t+1} - p_t) x_t\right]. \tag{2.2}$$

If the price increases by three dollars today, the trader receives three times the number of shares they are holding.

The expectation won't be important at this point of the model. The only natural uncertainty in the model is ω, which is known to the informed trader. There will be further uncertainty due to the fact that the informed trader will randomize, but again that is known to the trader. The only time the expectation would be meaningful would be if the market maker is randomizing their price, but that will be ruled out in equilibrium.

The asset pays no dividends (Japanese yen, gold, Amazon stock, etc.). This means that the only profit to the informed trader comes from changes in the price. You can also think of the next dividend of the asset being far enough in the future to be beyond the horizon of the game.

Market Maker

The market maker is meant as a stand in for whatever market process determines the prices and creates liquidity. What is important is simply that for any beliefs the market will arrive at some price and there will be traders willing to buy or sell.

For concreteness sake, I will say that the market maker chooses a price each period to maximize some flow payoff and stands willing to buy or sell up to one share at that price. Call $P(\mu_t)$ the optimal price given beliefs μ_t. I will assume this function to always have a unique value.

$$P(\mu_t) = \operatorname*{argmax}_{p_t \in \mathbb{R}} U(p_t, \mu_t) \tag{2.3}$$

for some period payoff function $U(p_t, \mu_t)$.

You may think of $U(p_t, \mu_t)$ as capturing profit from unspecified liquidity traders in the market.

The simplest example is thinking of a robot market maker that simply sets the price equal to some expected value like in Kyle (1985).

$$P(\mu_t) = \mathbb{E}[z(\omega)] \tag{2.4}$$

where z represents the value of the asset to the market maker (resale value) in each state of the world.

Another example is to consider a large market where $P(\mu_t)$ is a standard price equation for stochastic discount factor m and asset payoffs z by having

$$U(p_t, \mu_t) = \mathbb{E}_{\mu_t}\left[(p_t - m(\omega)v(\omega))^2\right]. \tag{2.5}$$

All that matters for the results is that there is some function $P(\mu)$ that gives the price for any beliefs and is single valued.

Equilibrium

I will solve for the Perfect Bayesian Equilibrium that gives the highest payoff to the informed trader. Given the market maker's pricing strategy, the informed trader will choose asset holdings each period to maximize their expected discounted profit. Given the informed trader's trading strategy and beliefs about the state, the market maker will choose the price to maximize period utility $U(p_t, \mu_t)$. Beliefs at all nodes will be obtained using Bayes rule from the informed trader's strategy when possible.

I will focus throughout the paper on Markov equilibria only. Other equilibria of the game will be discussed in the appendix.

2.3 Profit Maximizing Equilibrium

The strategy to solving for the informed trader's maximal equilibrium profit is as follows. Write the recursive formulation of the problem. Translate the problem into one of choosing posteriors rather than trades. Assume the function $P(\mu)$ is continuous and monotone. Do value function iteration by hand with an initial guess of $V(\mu) = 0$. We will see that our iteration will converge to the optimum in only two steps. After studying the solution, I take the limit to continuous time and further analyze the dynamics of prices.

Similar to a mechanism design problem, I will set this up as the the trader choosing all variables (prices, beliefs, and trades) subjects to contraints for incentive compatability and Bayesian updating to insure that it is an equilibrium of the game. The incentive constraint for the market maker to be maximizing is simply $p_t = P(\mu_t)$ here. We can plug this straight into the problem. This will give the following recursive formulation.

$$V(\mu) = \max_{x \in [-1,1]} \mathbb{E}\left[\left(P(\mu') - P(\mu)\right)x + \delta V(\mu')\right] \tag{2.6}$$

where μ' is obtained using Bayes rule.

Take any two posteriors on opposite sides of the prior, $\bar{\mu} \geq \mu \geq \underline{\mu}$. By rearranging Bayes rule, we can see that there always exists a mixed strategy for x such that these posteriors would be induced after buying and selling.

$$\pi(x=1|\omega=1)=\frac{\bar{\mu}(\mu-\underline{\mu})}{\mu(\bar{\mu}-\underline{\mu})};\qquad \pi(x=1|\omega=0)=\frac{(1-\bar{\mu})(\mu-\underline{\mu})}{(1-\mu)(\bar{\mu}-\underline{\mu})} \tag{2.7}$$

So, the problem can be written as one of choosing two posteriors on opposite sides of the prior subject to incentive compatability. It will be made clear in the next section that incentive compatability here reduces to the two posteriors giving the same payoff.

Iteration

Assume $P(\mu)$ is continuous and monotone. The usual arguements apply for the Contraction Mapping Theorem, so we can iterate to find the value function. Take an initial guess of $V(\mu)=0$ and consider the first iteration.

$$V_1(\mu)=\max_{x\in[-1,1]}\mathbb{E}[(P(\mu')-P(\mu))x] \tag{2.8}$$

Before telling you what is an equilibrium, it will be illustrative to tell you what isn't an equilibrium. It would seem natural to think that the informed trader should buy if the state is good and sell if the state is bad. This won't be supported in any equilibrium. In this candidate equilibrium, after observing a buy beliefs would go to one and the price would rise to $P(1)$. After observing a sell price would fall to $P(0)$. Hold the updated prices after observing buy or sell fixed and consider the best response of the informed trader. The traders gets a profit of $P(1)-P(\mu)$ if they buy and $P(\mu)-P(0)$ if they sell. Since price is continuous and monotone, there is one knife edge case where these will be equal. Call that belief μ^*.

$$\mu^*=P^{-1}\left(\frac{P(1)-P(0)}{2}\right) \tag{2.9}$$

If the price is currently lower than that, buying is more profitable than selling. Holding fixed the market maker's strategy, always buying is a profitable deviation.

In essance, it isn't credible for the informed trader to completely reveal the state because revealing the good state is more valuable that revealing the bad state. They have an incentive to lie when the state is bad.

However, always buying (or always selling) won't be profitable for the informed trader. The action will be uninformative. This means that beliefs, and therefore prices, won't change. This gives zero profit to the informed trader.

Any equilibrium with positive profits must then have the informed trader playing a mixed strategy. For them to be willing to play a mixed strategy, it must be that they are indifferent between buying and selling. Consider two posteriors on opposite sides of the prior that give the same payoff.

$$P(\bar{\mu}) - P(\mu) = P(\mu) - P(\underline{\mu}) \tag{2.10}$$

If the market maker is playing $P(\bar{\mu})$ after observing a buy and $P(\underline{\mu})$ after observing sell, then the informed trader is indifferent between the two actions. In fact, the informed trader's best response contains any mixed strategy of the two actions. As we saw above, there exists a mixed strategy such that $\bar{\mu}$ and $\underline{\mu}$ are the correct Bayesian updates. Thus, both players are playing a best response. These strategies constitutes an equilibirium. If we draw a graph of the the profit to the informed trader against the posterior of the market maker, any flat line on the graph connecting two points on opposite sides of the prior is an equilibrium profit.

Consider the highest profit the informed trader can obtain in equilibrium. This would be the highest flat line on such a graph. If $P(\mu)$ is monotone, then the highest payoff on either side of the posterior is at the boundary. Consider the boundary with the lower payoff of the two. Clearly the informed trader cannot get a higher payoff than that. If $P(\mu)$ is continuous, the Intermediate Value Theorem says there must

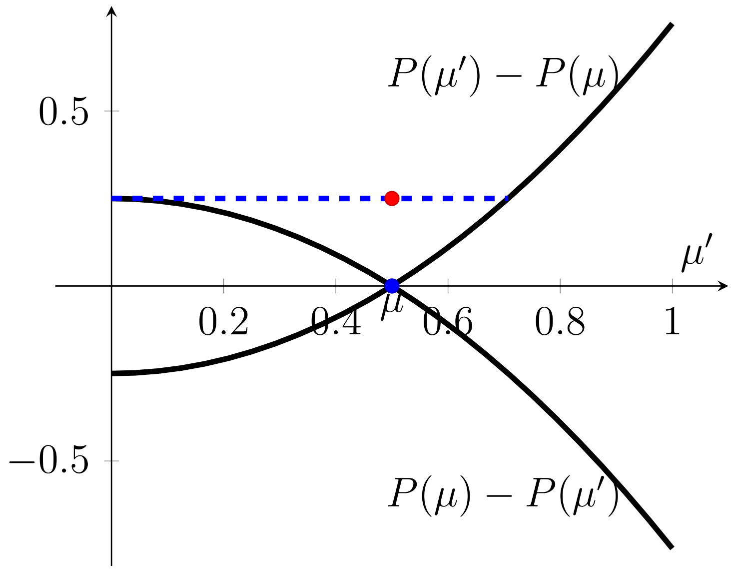

Figure 2.1: The upward sloping solid line is the profit from buying the asset graphed against the posterior induced. The downward sloping line is the profit from selling the asset. An equilibrium is a flat line on this graph that crosses the prior. The highest equilibrium profit is shown by the dashed line.

be a posterior on the other side of the prior that gives the same payoff.

This gives the new value function after one iteration.

$$V_1(\mu) = \min\{P(\mu) - P(0), P(1) - P(\mu)\} \tag{2.11}$$

Solution

To do the next iteration, all we need to do is find the highest flat line on

$$W(\mu') = |P(\mu') - P(\mu)| + \delta V_1(\mu') \tag{2.12}$$

graphed against μ'.

Take $\mu < \mu^*$. For any $\delta < 1$, $W(\mu')$ is decreasing on $[0, \mu]$. Take $\hat{\mu} \in [0, \mu]$.

$$W(\hat{\mu}) = P(\mu) - P(\hat{\mu}) + \delta(P(\hat{\mu}) - P(0)) \tag{2.13}$$

$$< P(\mu) - P(0) \tag{2.14}$$

So, the highest payoff to the left is still at the boundary. Similarly, the highest payoff to the right is at the other boundary. Again, by continuity the minimum of these can be obtained by an equilibrium. This gives our value function. This is the same value we had in the previous iteration. Thus, we have found a fixed point.

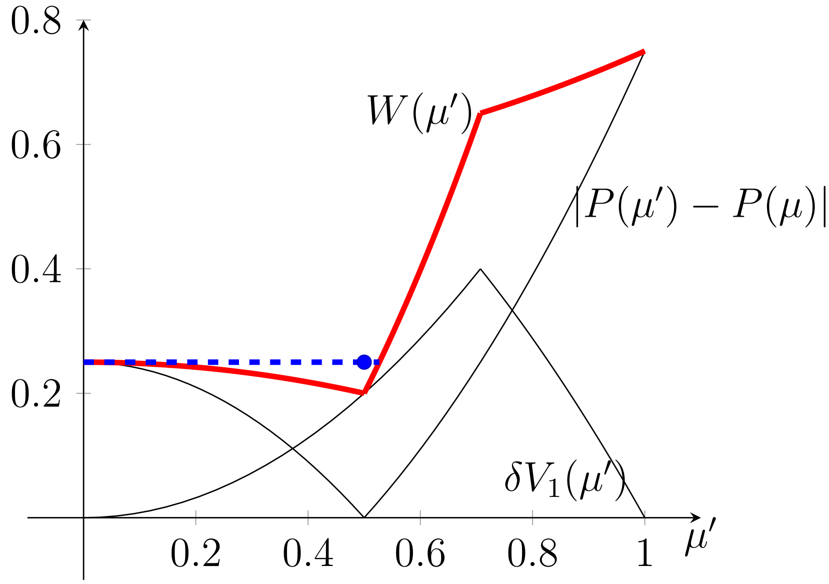

Proposition 1. *Let $P(\mu)$ be continuous and monotone. For any discount factor, $\delta \in [0, 1)$, the value is*

$$V(\mu) = \min\{P(\mu) - P(0),\ P(1) - P(\mu)\} \tag{2.15}$$

Price Dynamics

The endpoints of the flat line giving the value are the posteriors that are induced by the optimal strategy. We saw that one endpoint is always at the boundary, but the other is generally interior. Take $\mu < \mu^*$. When the informed trader sells, this perfectly reveals the state to be bad and beliefs fall to zero. When the insider buys, beliefs increase to an interior point just high enough to make the informed trader indifferent between buying and selling.

The payoff from inducing a posterior μ' can be obtained by putting the the value function from the previous proposition.

$$W(\mu') = \begin{cases} P(\mu) - \delta P(0) - (1-\delta)P(\mu') & \text{if } \mu' \leq \mu \\ (1+\delta)P(\mu') - P(\mu) - \delta P(0) & \text{if } \mu < \mu' \leq \mu^* \\ (1-\delta)P(\mu') - P(\mu) + \delta P(1) & \text{if } \mu' > \mu^* \end{cases} \tag{2.16}$$

To the left of the prior we have a strictly decreasing function of μ', so the left endpoint will be zero. The right endpoint could be on the second or third segment of the curve $W(\mu')$.

$$\bar{\mu} = \begin{cases} P^{-1}\left(\frac{2P(\mu)-(1-\delta)P(0)}{1+\delta}\right) & \text{if } \mu \leq P^{-1}\left(\frac{P(1)(1+\delta)+P(0)(3-\delta)}{4}\right) \\ P^{-1}\left(\frac{2P(\mu)-P(0)-\delta P(1)}{1-\delta}\right) & \text{otherwise} \end{cases} \tag{2.17}$$

This equation doesn't have a nice intuition to it, but it is important to note that unlike the left endpoint, the right endpoint does depend on δ.

Continuous Time Limit

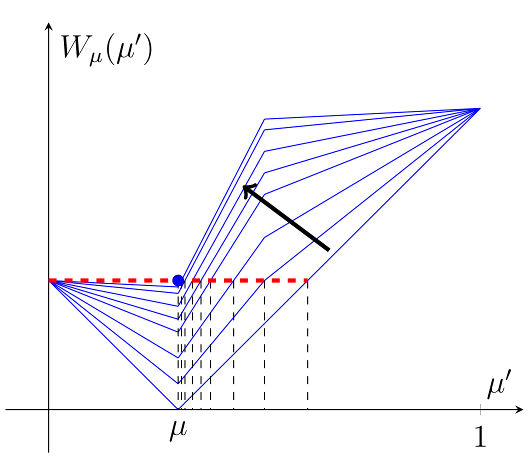

Figure 2.2: The blue lines are $W(\mu')$ for

Consider the limit as δ goes to 1. After selling beliefs still drop to 0. After buying beliefs move up to the right endpoint described above. This right endpoint is moving closer and closer to the prior as

δ goes to 1. Since beliefs are a martingale, this means that the frequency of the jumps down to $P(0)$ needs to be going to 0. The likelihood that beliefs just drift up the small amount in a given period goes to 1. Though not obvious from looking at it, in the appendix I show that the probability of a jump is shrinking to 0 at a linear rate in $1-\delta$. This gives use the first main result of the paper.

Theorem 1. *For any differentiable strictly monotone price function, $P(\mu)$, as δ goes to one the price proccess converges to a Poisson process.*

- *If $\mu_t < \mu^*$,*

$$dP(\mu_t) = \frac{1}{2}(P(\mu_t) - P(0))dt - (P(\mu_t) - P(0))dN_t \tag{2.18}$$

  where N_t is a standard Poisson process with arrival rate $\lambda = \frac{P(\mu_t)-P(0)}{2\mu_t P'(\mu_t)}$.

- *If $\mu_t > \mu^*$,*

$$dP(\mu_t) = -\frac{1}{2}(P(1) - P(\mu_t))dt + (P(1) - P(\mu_t))dN_t \tag{2.19}$$

  where N_t is a standard Poisson process with arrival rate $\lambda = \frac{P(1)-P(\mu_t)}{2(1-\mu_t)P'(\mu_t)}$.

- *If $\mu_t = \mu^*$, all information is revealed immediately and the price jumps to either $P(1)$ or $P(0)$.*

Consider a low initial price ($\mu < \mu^*$). With a Poisson arrival rate the informed trader will sell the asset and completely reveal that the state is bad. This makes

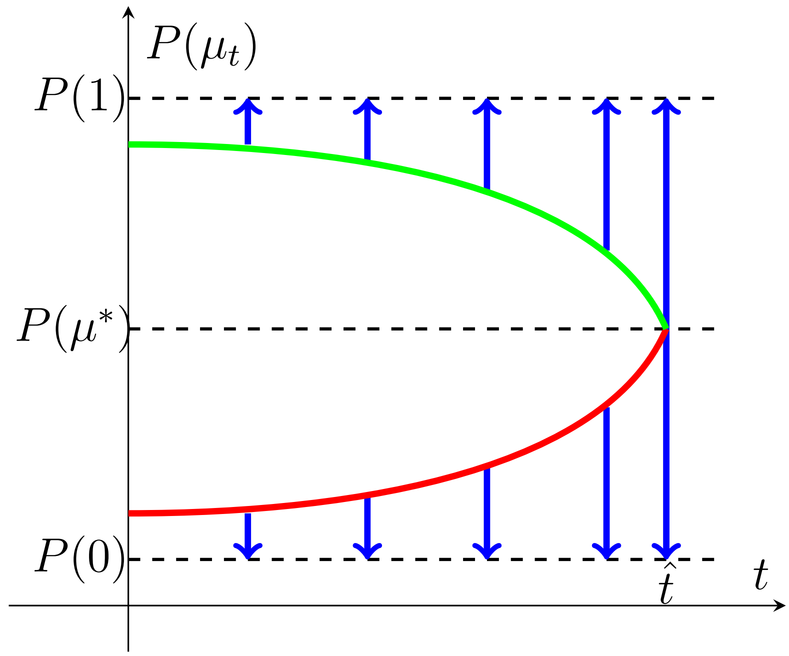

Figure 2.3: Sample price paths when the price starts high (green) and when the price starts low (red).

the price crash to $P(0)$. While this Poisson information hasn't arrived, the informed trader will hold a long position in the asset and the price will slowly drift upward. If the price starts high, the dynamics are simply the mirror image. With a Poisson arrival rate the informed trader will buy and the price will spike to $P(1)$. At all other times the informed trader holds a short position and the price drifts down. If the beliefs ever reach μ^*, the informed trader perfectly reveals the state, good or bad.

The arrival rate can be obtained intuitively by considering the informed trader's incentive compatability constraint. The informed trader must be indifferent between revealing the state to be bad today or letting the price drift up a little and revealing the state to be bad tomorrow. Call λ the arrival rate of the Poisson process. Once it arrives, beliefs jump down to zero. In order for beliefs to be a martingale, the drift must equal $\lambda_t \mu_t$. The drift in price is then equal to $\lambda_t \mu_t P'(\mu_t)$.

$$\underbrace{P(\mu_t) - P(0)}_{\text{reveal today}} \;\approx\; \underbrace{P'(\mu_t)\mu_t\lambda_t}_{\text{drift today}} \;+\; \delta(\underbrace{P'(\mu_t)\mu_t\lambda_t + P(\mu_t) - P(0)}_{\text{reveal tomorrow}}) \tag{2.20}$$

As δ gets large, this gives a linear relationship between λ and $1 - \delta$.

$$\Rightarrow \; \lambda_t \approx \frac{P(\mu_t) - P(0)}{2\mu_t P'(\mu_t)}(1 - \delta) \tag{2.21}$$

Intuition

If the price starts low, the dynamics look like a pump-and-dump scheme. The informed trader buys every period and spreads good information to pump up the value of the asset. Then, at a random arrival date they dump all their shares and reveal the asset to be bad. This causes the price to crash suddenly. If the asset actually is good that crash date never comes, hence the positive drift.

If the price starts high we have the mirroring dynamics. The informed trader sells each period causing the price to drift down. If the asset is actually good, then at a random arrival rate they buy back all the shares and reveal the state causing a sudden spike in price. This is a short-and-distort scheme. This matches the empirical fact that pump-and-dumps are usually done on cheap stocks and short-and-distorts on more expensive stocks.

The form of the optimal strategy comes from the intuition in the one period model for why buy when good sell when bad isn't an equilibrium. If the price is initially low, the informed trader can't credibly reveal the state because revealing the good state is better than revealing the bad state. They'd like to lie and always say it's the good state.

The dynamics give them that credibility. When the informed trader says the state is bad, the market maker believes them and moves the price all the way to $P(0)$. When the informed trader says the state is good, the market maker mostly doesn't believe it because there is a much bigger potential gain from the state being good. Price increases only a small amount. After many periods of repeatedly saying the state is good, beliefs eventually drift up to μ^* which is the cutoff point for when the good state can be credibly revealed. They small price jump can be done immediately, but the informed trader needs to spend time to build credibility before they can get the big price jump.

The arrival rate of the Poisson process depends on the function $P(\mu)$. If $P(\mu)$ is linear then is simplifies to $\lambda = \frac{1}{2}$. If $P(\mu)$ is concave, then the arrival rate will be bigger than $\frac{1}{2}$ and it will be increasing over time. If $P(\mu)$ is convex, then the arrival rate will be smaller than $\frac{1}{2}$ and it will be decreasing over time. You can see this by writing the Taylor series for $P(0)$.

$$P(0) = P(\mu_t) - P'(\mu_t)\mu_t + \frac{1}{2}P''(\mu_t)\mu_t^2 + \dots \tag{2.22}$$

This gives us an approximate equation for the arrival rate.

$$\lambda(\mu_t) = \frac{P(\mu_t) - P(0)}{2\mu_t P'(\mu_t)} \approx \frac{1}{2} - \frac{1}{4}\mu_t \frac{P''(\mu_t)}{P'(\mu_t)} \tag{2.23}$$

We can see there how the arrival rate depends on the concavity of $P(\mu)$.

However, regardless of the curvature of $P(\mu)$, the price will always reach $P(\mu^*)$ and then jump to the boundaries in finite time. We can see this because the magnitude of the drift is increasing in time.

Proposition 2. *The maximum time to full information revelation is*

$$t^{max} = 2\log\left(\frac{P(\mu^*) - P(0)}{P(\mu_0) - P(0)}\right) \tag{2.24}$$

if $\mu_0 \leq \mu^*$ *and*

$$t^{max} = 2 \log \left(\frac{P(1) - P(\mu^*)}{P(1) - P(\mu_0)} \right) \tag{2.25}$$

if $\mu_0 > \mu^*$.

The proof of this result is in the appendix.

2.4 Commitment

In the previous section buy if the state is good sell if the state is bad was not an equilibrium because the informed trader could not credibly commit to that strategy. In this section I'll derive the optimal policy when they do have that commitment power. The informed trader is allowed to commit ex ante to their mixed strategy.

Rather than choose an action (buy or sell) each period the informed trader can choose a distribution of actions contingent on the state. For example they can buy when good and sell when bad, or when the state is good randomize fifty-fifty between buying and selling. This is the same type of commitment power that is assumed in the Bayesian persuasion literature. We could equivalently think of commitment as saying the informed trader has verifiable information to reveal. This essentially allows us to ignore the incentive compatability constraint in the previous problem. This unconstrained problem is much more difficult to solve analytically, but I can still characterize the continuous time limit as in the previous section.

Even though buy when the state is good sell when the state is bad is clearly going to be optimal in the one period model in the dynamic model the solution is nearly the exact opposite.

Setup

As before the solution will require the informed trader to mix, but they can now mix between any two points on the payoff graph, not just ones that are equal.

I can write the problem as choosing the posteriors that will be induced after buying and selling subject to the constraint that beliefs are a martingale. If I write the posteriors as

$$\mu_{buy} = \mu + \bar{\epsilon}, \quad \text{and} \quad \mu_{sell} = \mu - \underline{\epsilon} \tag{2.26}$$

then the martingale requirement stipulates the probabilities must be

$$p(buy) = \frac{\underline{\epsilon}}{\bar{\epsilon} + \underline{\epsilon}} \quad \text{and} \quad p(sell) = \frac{\bar{\epsilon}}{\bar{\epsilon} + \underline{\epsilon}}. \tag{2.27}$$

The constraints $\bar{\epsilon} \in [0, 1-\mu]$ and $\underline{\epsilon} \in [0, \mu]$ ensure that beliefs stay between zero and one.

The Bellman equation can then be written simply.

$$V(\mu) = \max_{\bar{\epsilon},\underline{\epsilon}} \ \left(|P(\mu+\bar{\epsilon}) - P(\mu)| + \delta V(\mu+\bar{\epsilon})\right) \frac{\underline{\epsilon}}{\bar{\epsilon} + \underline{\epsilon}} + \left(|P(\mu) - P(\mu-\underline{\epsilon})| + \delta V(\mu-\underline{\epsilon})\right) \frac{\bar{\epsilon}}{\bar{\epsilon} + \underline{\epsilon}} \tag{2.28}$$

$$= \max_{\bar{\epsilon},\underline{\epsilon}} \ \mathbb{E}_{\mu'} \left[|P(\mu') - P(\mu)| + \delta V(\mu')\right] \tag{2.29}$$

We'll characterize the the approximate solution when it's assumed $\bar{\epsilon}$ and $\underline{\epsilon}$ are small. Then when we take the limit to continuous time, the approximation will give us the exact solution. First notice that $|P(\mu+\bar{\epsilon}) - P(\mu)| \approx |P'(\mu)|\bar{\epsilon}$. Now take a second order approximation to $V(\mu+\bar{\epsilon})$.

$$V(\mu+\bar{\epsilon}) \approx V(\mu) + V'(\mu)\bar{\epsilon} + \frac{1}{2}V''(\mu)\bar{\epsilon}^2 \tag{2.30}$$

Put these into equation (2.28).

$$(1-\delta)V(\mu) = \max_{\bar{\epsilon},\underline{\epsilon}} \ 2|P'(\mu)|\frac{\bar{\epsilon}\underline{\epsilon}}{\bar{\epsilon} + \underline{\epsilon}} + \frac{1}{2}\delta V''(\mu)\bar{\epsilon}\underline{\epsilon} \tag{2.31}$$

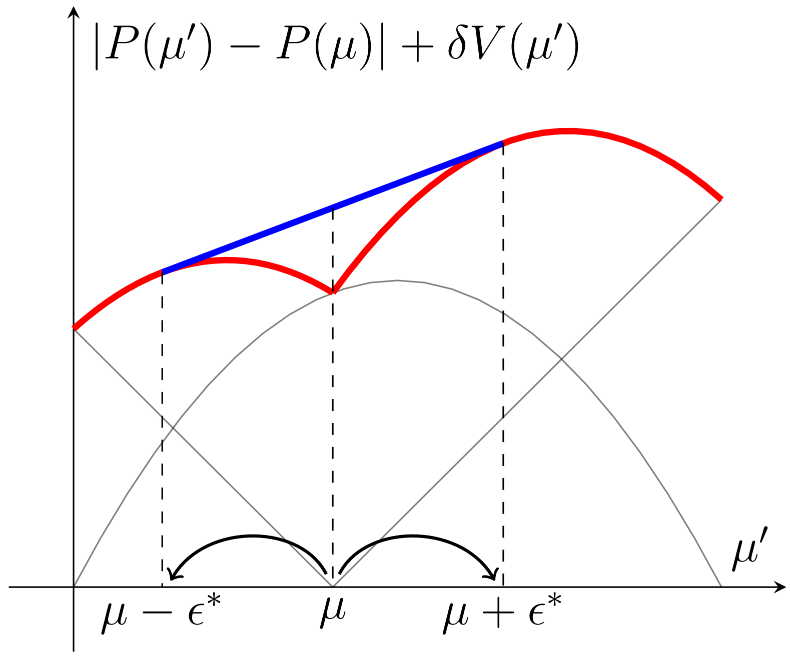

Figure 2.4: The faded black lines are the one period value, $|P(\mu') - P(\mu)|$, and the continuation value, $\delta V(\mu')$. The thick red line is the sum of those two. The blue segment connecting it gives the optimal policy and value. Beliefs either jump up or down by step size ϵ^* each period.

The main tradeoff can be seen in the two right hand side terms. Giving away more information gives a higher payoff this period because it leads to larger price changes. On the other hand since $V(\mu)$ is concave, the more information I give away the worse my expected continuation value.

Price Dynamics

Holding the product $\bar{\epsilon}\underline{\epsilon}$ fixed, we'd like to minimize the sum $\bar{\epsilon} + \underline{\epsilon}$. This is always accomplished when $\bar{\epsilon} = \underline{\epsilon}$. So, this is only a problem of one variable when we have an interior solution.

$$(1 - \delta)V(\mu) = \max_{\epsilon} \ |P'(\mu)|\epsilon + \frac{1}{2}\delta V''(\mu)\epsilon^2 \tag{2.32}$$

Taking the derivative and setting it equal to zero yields the optimal ϵ.

$$\epsilon^* = -\frac{|P'(\mu)|}{\delta V''(\mu)} \tag{2.33}$$

Beliefs now follow a random walk. Each period beliefs either jump up or down by a small step of size ϵ^*. As δ goes to one, $V(\mu)$ and $V''(\mu)$ both grow in magnitude toward infinity. This means that the step size, ϵ^*, is going to zero. The key is that it is going to zero slowly (at a rate of $\sqrt{1-\delta}$). As we take the step size shrinking to zero, this converges to a Brownian Motion for beliefs. Itó's Lemma gives us the process for prices which are a smooth function of beliefs. Thus, we have the second main result. Details are given in the appendix.

Theorem 2. *Let $P(\mu)$ be any $\mathcal{C}^2$ function. As δ goes to one, the price converges to an Itó Process over time.*

$$dP(\mu_t) = \frac{1}{2}P''(\mu_t)\sigma^2(\mu_t)dt + P'(\mu_t)\sigma(\mu_t)dB_t \tag{2.34}$$

The B_t here is a standard Brownian Motion.

Intuition

In this section I ellaborate on the form of the solution and give intuition.

Even though optimal information release is Brownian, there is still a non-constant drift and variance term. The function

$$\sigma(\mu_t) = \frac{|P'(\mu_t)|}{\hat{V}''(\mu_t)} \tag{2.35}$$

is the standard deviation multiplying the Brownian increments in beliefs. $\hat{V}(\mu_t)$ is the value function after rescaling for the discount factor. Itó's Lemma tells us that the standard deviation multiplying the Brownian increment on prices is then $P'(\mu_t)\sigma(\mu_t)$. There is higher variance when the price is more sensitive to information.

The drift in prices is also pinned down. Since beliefs need to be a martingale, they must have zero drift. This doesn't mean that prices won't have a drift. In fact, Jensen's inquality tells us that the drift needs to be positive (negative) when price is a convex (concave) function of beliefs. Itó's Lemma confirms that the drift is $\frac{1}{2}P''(\mu_t)\sigma^2(\mu_t)dt$.

This is not a complete solution because $\sigma(\mu_t)$ was defined in terms of the value function for which a complete analytical solution cannot always be given. In the special case of a linear price function, $P(\mu_t) = \mu_t$, the analytic solution can be written. We then have that

$$\sigma(\mu_t) = n(N^{-1}(\mu_t)) \tag{2.36}$$

where $n(\cdot)$ is the normal distribution pdf and $N(\cdot)$ is the normal distribution cdf.

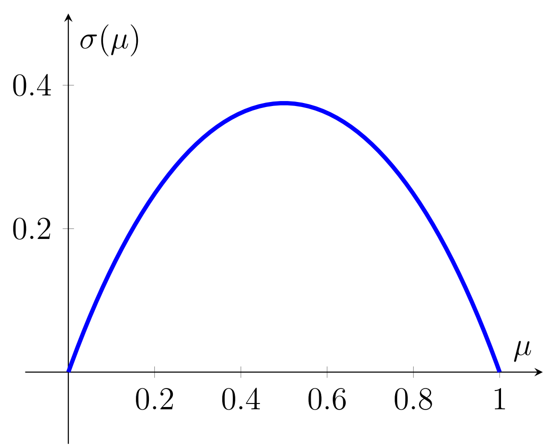

Figure 2.5: The standard deviation of prices as a function of beliefs.

The function $\sigma(\mu_t)$ is the normal distribution evaluated at the μ_t quantile. It is similar to a geometric Brownian motion in that the standard deviation goes to zero linearly in the price. This ensures that prices can never drop below zero. Prices are the most volatile when there is the most uncertainty.

From the perspective of the informed trader, the price still follows and Itó process. The variance is the same, but

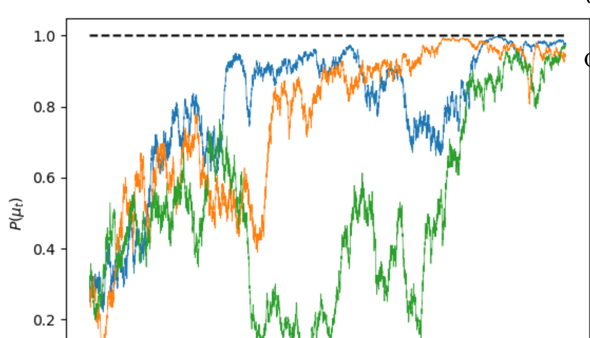

the drift is different. With the linear price function, the drift conditional on the state being good is $\frac{\sigma^2(\mu_t)}{\mu_t}$. When the state is bad, the drift is $-\frac{\sigma^2(\mu_t)}{1-\mu_t}$. Beliefs always drift toward the true state. When beliefs are far from the true state the drift is large, and when they are close to the truth the drift is small. Beliefs converge to the truth over time in the weak$-^*$ topology.

Comparison

In this section I discuss the relation between the solutions with and without commitment and their importance for finance and economics.

The price paths induced with and without commitment are intuitively the opposite of each other. With commitment, the information release happens gradually and all information is revealed only in the limit. The price paths are continuous and highly volatile. This represents inside information being slowly leaked to the public and being encorporated into the price. Without commitment, all information is released suddenly in finite time. The price paths are smooth until a discontinuous jump. This is a pump-and-dump scheme to manipulate the price of the asset.

The dramatically different price dynamics can inform us on whether observed information leaks are likely to be strategic. Consider the commitment assumption to be about whether the informed trader can generate verifiable evidence or not. A strategic insider releases information with verifiable evidence gradually, but unverifiable information is released suddenly.

Together these form a micro-foundation of the price processes commonly assumed throughout finance.

2.5 Extensions

Persistence

Say that the state is not permanent. Assume for this section that the state follows a Markov process. Call π_1 and π_0 the probability that $\omega_{t+1} = 1$ conditional on $\omega_t = 1$ or 0 respectively.

The timing of the game requires a bit more care in this section. At the beginning of period t beliefs are μ_t. Then, the informed trader can choose to buy or sell the asset at price $P(\mu_t)$. The market maker immediately observes the trade and updates beliefs according to Bayes rule to μ'_t. The price is updated right away and the informed trader closes their position at price $P(\mu'_t)$. Then, after trading is done ω_{t+1} is drawn from a Markov process. Beliefs at this point are updated for the next period, $\mu_{t+1} = \pi_0 + (\pi_1 - \pi_0)\mu'_t$. In the model with a permanent state, $\mu_{t+1} = \mu'_t$. It wasn't important at that point to say that the informed trader closes their position at the end of each period, because the price at the end of each period was the same as the price at the start of the next period. That is no longer the case. Between periods, the state could change. Thus, beliefs and prices will also change between each period.

There are two different ways we can think about private information of a persistent state. The first is that the informed trader is able to see ω_t every period. The second is that the informed trader is able to see the state only in the first period, ω_0. In the first, the fact that the state is persistent rather than permanent is a good thing for the informed trader. It means that there is more information flowing to them each period. The informed trader then has more opportunity for profit. In the second, the fact that the state is persistent rather than permanent is a bad thing for the informed trader. It means that their information has less predictive power of the state as time passes. The informed trader's information is becoming less valuable each period.

Interestingly, both cases incentivize the informed trader to reveal information at a faster rate. I will show in this section that even though the value and the price process will look very different in the two cases, the optimal strategy is identical. I am not aware of any other paper that shows this kind of relationship between the two types of private information of a persistent state.

One Time Information

The informed trader observes ω_0 but not ω_t for $t > 0$. In this section, i will assume a linear price function for simplicity. $P(\mu) = \mu$. It is still the case that the informed trader can choose to buy when the price is about to go up and sell when the price is about to go down. The objective is

$$\mathbb{E}\left[\sum_{t=0}^{\infty} \delta^t |\mu_t' - \mu_t|\right] \tag{2.37}$$

where $\mu_{t+1} = \pi_0 + (\pi_1 - \pi_0)\mu_t'$.

The main difference here is that the informed trader no longer has complete control over beliefs in every period. μ_t is the beliefs about the state in period t, but the informed trader only knows the state in period 0. Call $\tilde{\mu}_t$ the beliefs in date t about

ω_0. These are the beliefs that the informed trader can control.

$$\mu_0 = \tilde{\mu}_0$$
$$\mu_1 = \pi_0 + (\pi_1 - \pi_0)\tilde{\mu}_1$$
$$\vdots$$
$$\mu_t = \sum_{\tau=0}^{t-1} \pi_0(\pi_1 - \pi_0)^\tau + (\pi_1 - \pi_0)^t \tilde{\mu}_t$$

In equation 2.37 the informed trader is constrained by Bayes plausibility, incentive compatability, and shrinking bounds on where beliefs can be sent due to the informativeness of their signal deteriorating. The level of persistence puts an upper and lower bound on beliefs each period.

$$\sum_{\tau=0}^{t-1} \pi_0(\pi_1-\pi_0)^\tau \leq \mu_t \leq \sum_{\tau=0}^{t-1} \pi_0(\pi_1-\pi_0)^\tau + (\pi_1-\pi_0)^t \tag{2.38}$$

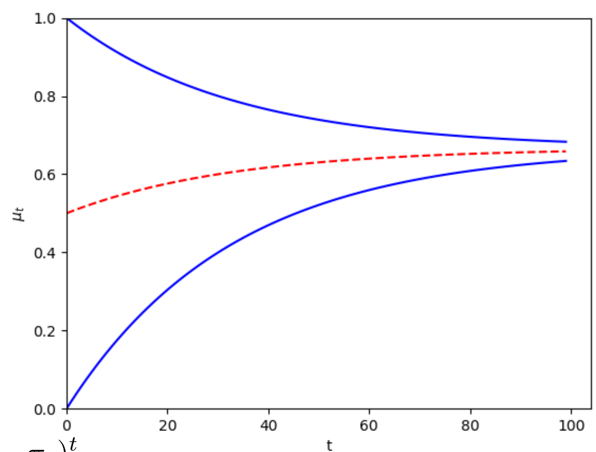

Figure 2.7: The solid lines are the upper and lower bounds on beliefs. The dashed line shows the drift. $\pi_1 = .99$, $\pi_0 = .02$

As time goes on, beliefs must ultimately converge to $\frac{\pi_0}{1-\pi_1+\pi_0}$ regardless of the informed trader's actions. The martingale condition only holds within each period. Between periods the beliefs have a drift determined solely by the persistence of the states. Intuitively, the size of the game is just shrinking over time. We can see this precisely by rewriting the problem in terms of $\tilde{\mu}_t$.

$$\mathbb{E}\left[\sum_{t=0}^{\infty} \delta^t |\mu_t' - \mu_t|\right] = \mathbb{E}\left[\sum_{t=0}^{\infty} (\delta(\pi_1 - \pi_0))^t |\tilde{\mu}_t' - \tilde{\mu}_t|\right] \tag{2.39}$$

The only effect of persistence on the strategy is to reduce the discount factor. To take limits set $\delta = 1 - r\Delta t$, $\pi_1 = 1 - \lambda_1 \Delta t$, and $\pi_0 = \lambda_0 \Delta t$. For small time periods, $\tilde{\delta} = \delta(\pi_1 - \pi_0) \approx 1 - (r + \lambda_1 + \lambda_0)\Delta t$. The arrival rate of state switches simply adds on the the discount rate. The rest of the calculations from the proofs of theorem 1 and 2 go through the same as before.

Without commitment beliefs follow,

$$d\tilde{\mu}_t = \frac{r + \lambda_1 + \lambda_0}{2} \tilde{\mu}_t dt - \tilde{\mu}_t dN_t \tag{2.40}$$

if $\tilde{\mu}_t < \frac{1}{2}$, the symmetric equation if $\tilde{\mu}_t > \frac{1}{2}$, or jump immediately to 0 or 1 if $\tilde{\mu}_t = \frac{1}{2}$. The arrival rate of the Poisson process is $\lambda(\tilde{\mu}_t) = \frac{r+\lambda_1+\lambda_0}{2}$.

With commitment beliefs follow

$$d\tilde{\mu}_t = \frac{r + \lambda_1 + \lambda_0}{2} \phi(\tilde{\mu}_t) dB_t. \tag{2.41}$$

Prices are not based on beliefs about what the state was in date zero, $\tilde{\mu}_t$, but on beliefs about the current state, μ_t. In the limit, μ_t can still be written as a function of $\tilde{\mu}_t$ and time.

$$\mu_t = \frac{\lambda_0}{\lambda_1 + \lambda_0} + (1 - \lambda_1 - \lambda_0)^t \left(\tilde{\mu}_t - \frac{\lambda_0}{\lambda_1 + \lambda_0} \right) \tag{2.42}$$

Prices without commitment must then follow

$$dP(\mu_t) = \left(\frac{r + \lambda_1 + \lambda_0}{2} + \log(1 - \lambda_1 - \lambda_0) \left(\tilde{\mu}_t - \frac{\lambda_0}{\lambda_1 + \lambda_0} \right) \right) (1-\lambda_1-\lambda_0)^t dt - \tilde{\mu}_t (1-\lambda_1-\lambda_0)^t dN_t \tag{2.43}$$

whenever $\mu_t < \mu_t^*$, the symmetric equation when $\mu_t > \mu_t^*$, and jump to the shrinking boundaries immediately when $\mu_t = \mu_t^*$. Once the price hits a boundary, it remains on the boundary for the rest of the game and continues to drift toward $P\left(\frac{\lambda_0}{\lambda_1+\lambda_0}\right)$. The midpoint, μ_t^* is also changing over time now.

$$\mu_t^* = \frac{1}{2}(1-\lambda_1-\lambda_0)^t + \left(1-(1-\lambda_1-\lambda_0)^t\right)\frac{\lambda_0}{\lambda_1+\lambda_0} \tag{2.44}$$

This is the dashed line in the previous figure.

With commitment prices follow

$$dP(\mu_t) = \log(1-\lambda_1-\lambda_0)\left(\tilde{\mu}_t - \frac{\lambda_0}{\lambda_1+\lambda_0}\right)(1-\lambda_1-\lambda_0)^t dt + (1-\lambda_1-\lambda_0)^t \frac{r+\lambda_1+\lambda_0}{2}\phi(\mu_t)dB_t. \tag{2.45}$$

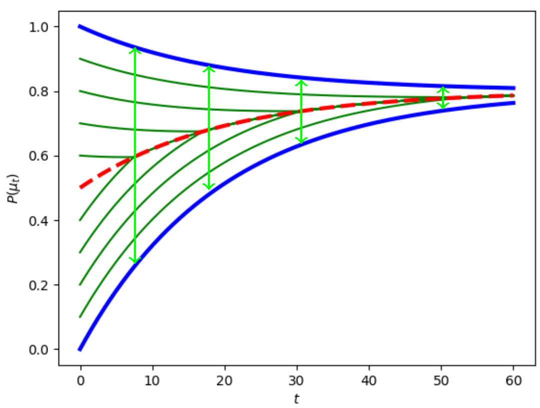

(a) Without commitment

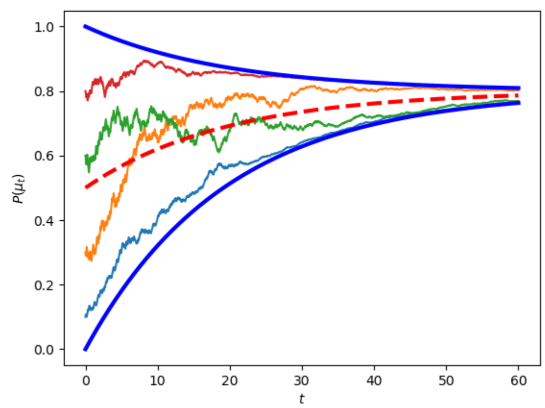

(b) With commitment

Figure 2.8: Sample paths for prices.

Information Flows Every Period

Basically, the persistence flattens out the continuation value function, but the current period payoff slope remains the same. This means you still always hit a boundary. The difference is that when you hit the boundary, the game doesn't end. Posteriors drift back interior because the state may switch and you can trade again. It is easy to verify that the value function is now

$$V(\mu) = \min\left\{P(\mu) - P(0) + \frac{P(\pi_0) - P(0)}{1-\delta}, P(1) - P(\mu) + \frac{P(1) - P(\pi_1)}{1-\delta}\right\}. \quad (2.46)$$

Assume the price function is linear to solve for the policy function and take the limit. Take $\mu < \mu^*$ (μ^* is not equal to $\frac{1}{2}$ anymore). The left endpoint of the strategy jumps to zero still. Some algebra reveals that the right endpoint is

$$\mu' = \frac{2\mu + (1-\delta)\pi_0}{1 + \delta(\pi_1 - \pi 0)}. \quad (2.47)$$

This approaches the prior, μ, as δ goes to one. This is approximately equal to

$$\mu' \approx \frac{2\mu}{2 - (r - \lambda_1 - \lambda_0)\Delta t}. \quad (2.48)$$

Write out the probability of jumping down to zero and simplify.

$$\frac{\mu' - \mu}{\mu' \Delta t} = \frac{\mu(r + \lambda_1 + \lambda_0)}{2\mu + \Delta t} \quad (2.49)$$

Then we have our arrival rate.

$$\lambda(\mu_t) = \lim_{\Delta t \to 0} \frac{prob_{sell}}{\Delta t} = \frac{r + \lambda_1 + \lambda_0}{2} \quad (2.50)$$

Notice that this is exactly the same strategy we had for $\tilde{\mu}_t$ when the informed trader only knew the state in period zero. The prices and value function are different, however. The value is scaled up by a factor depending on the persistence of the state. A less persistent state leds to a higher payoff, because the informed trader is getting more information. The price process is derived using Itò's Lemma. The drift comes from the potentially switching state and from the information conveyed by the Poisson event not arriving. There are Poisson jumps that send beliefs all the way to the boundary.

If $\mu_t < \mu^*$,

$$dP(\mu_t) = P'(\mu_t)(\lambda_0 + \mu_t(1 + \lambda(\mu_t) - \lambda_1 - \lambda_0))dt - (P(\mu_t) - P(0))dN_t. \tag{2.51}$$

The process is the reverse if $\mu_t \mu^*$.

N states

Say that there are N possible states of nature, $\omega \in \{0, 1, \ldots, N-1\}$. If the informed trader is choosing $x_t \in [-1, 1]$, it is no longer the case that they will always choose only $x = -1$, or $x = 1$ in the best equilibrium. In fact, the best equilibrium is not going to exist. Since there are N states, the informed trader could generally benefit from sending N messages. That means choosing orders of N different sizes. The informed trader would like to place orders as large as possible accompanying each message that makes price go up, but cannot have the order sizes be equal. The message that makes price increase the most will be accompanied with buying $x = 1$ shares. The informed trader would like the message that makes price increase the second most to be as high as possible without being equal to 1. Such a number does not exist without some constraint forcing discrete increments on orders. In this section I will focus on a the asset holdings of the insider being chosen from the discrete set, $x_t \in \{-1, 1\}$.

Lack of Commitment

Suppose there are N possible states of the world, $\omega \in \{\omega_1, \omega_2, \ldots, \omega_N\}$. You can relabel the states so they are in order of lowest price to highest price. Call μ_i the probability of state i. Let the price be linear function of the probabilities.

$$P(\mu) = \sum_{i=1}^{N} p_i \mu_i \tag{2.52}$$

By linearity, the support of the distribution of posteriors induced will still always contain a boundary point. The boundary point will either be one where the probability $\omega = \omega_1$ is zero or where the probability $\omega = \omega_N$ is zero.

In the limit the value is obtained with a series of Poisson processes all with arrival rate $\lambda = \frac{1}{2}$. If the price is low today, at a Poisson rate the price will jump down to the boundary where $\mu_N = 0$ (the trader sells) and the rest of the time the price will slowly drift up (the trader buys). Once it is revealed that the state is not $\omega = \omega_N$, the process starts over as before but with only $N - 1$ remaining possible states. The informed trader reveals the state of the world piece by piece with a series of Poisson processes. Each time a Poisson arrival happens, the trader is revealing one of the states that is not the true state.

Commitment

The problem is still equivalent to one where the informed trader chooses two posteriors rather than choosing trades directly. With two states the informed trader only had to choose how much beliefs increase after a buy and how much they decrease after a sell. With many states, the informed trader is still choosing the size of the increase after buying and decrease after selling, but they also need to choose a direction in $\mathbb{R}^{N-1}$. Let μ be an $N-1$ length vector representing the belief probabilities for all but the last state. The informed trader chooses $\tilde{\epsilon}$ from among unit vectors in $\mathbb{R}^{N-1}$. The informed trader also chooses $\underline{c}$, $\overline{c}$ from $\mathbb{R}_+$. Beliefs will move to $\mu + \overline{c}\tilde{\epsilon}$ after observing a buy and to $\mu - \underline{c}\tilde{\epsilon}$ afterobserving a sell. The martingale condition on beliefs will ensure that the two posteriors induced need to be in exactly opposite directions of each other. This means that the insider does not get to choose to directional vectors, only two magnitudes on one direction.

The setup of the problem still looks similar.

$$V(\mu) = \max_{\overline{c},\underline{c},\tilde{\epsilon}} \left(P(\mu+\overline{c}\tilde{\epsilon}) - P(\mu) + \delta V(\mu+\overline{c}\tilde{\epsilon})\right)\frac{\underline{c}}{\overline{c}+\underline{c}} + \left(P(\mu) - P(\mu-\underline{c}\tilde{\epsilon}) + \delta V(\mu-\underline{c}\tilde{\epsilon})\right)\frac{\overline{c}}{\overline{c}+\underline{c}} \tag{2.53}$$

There rest looks very similar to the proof of theorem 2. Taking the same approximations as in the two state case, we can simplify the problem.

$$(1-\delta)V(\mu) = \max_{\overline{c},\underline{c},\tilde{\epsilon}} \ 2\nabla P(\mu)^T\tilde{\epsilon}\frac{\overline{c}\underline{c}}{\overline{c}+\underline{c}} + \delta\tilde{\epsilon}^T\mathbf{H}_V(\mu)\tilde{\epsilon}\ \overline{c}\underline{c} \tag{2.54}$$

We can again see that in the optimum $\overline{c} = \underline{c}$ since $\mathbf{H}_V$ is negative definite. Define $\epsilon = c\tilde{\epsilon}$. The problem is again simplified.

$$(1-\delta)V(\mu) = \max_{\epsilon} \ \nabla P(\mu)^T\epsilon + \delta\epsilon^T\mathbf{H}_V(\mu)\epsilon \tag{2.55}$$

subject to the constraint that the vectors $\mu+\epsilon$ and $\mu-\epsilon$ each have elements that are positive and sum to less than one.

The maximum is obtain by differentiating.

$$\epsilon^* = \frac{1}{\delta}\mathbf{H}_V(\mu)^{-1}\ \nabla P(\mu) \tag{2.56}$$

We can now put ϵ^* into the objective function and get

$$(1-\delta)V(\mu) = -\frac{1}{\delta}\nabla P(\mu)^T\mathbf{H}_V(\mu)^{-1}\nabla P(\mu). \tag{2.57}$$

From this differential equation, we see that $\hat{V}(\mu) = V(\mu)\sqrt{\delta(1-\delta)}$ is constant in δ. If we substitute $\mathbf{H}_V(\mu) = \frac{\mathbf{H}_{\hat{V}}(\mu)}{\sqrt{\delta(1-\delta)}}$ into the equation for ϵ^* we get an easy expression.

$$\epsilon^* = \mathbf{H}_{\hat{V}}(\mu)^{-1}\ \nabla P(\mu)\sqrt{\frac{1-\delta}{\delta}} \tag{2.58}$$

With $\Delta t = \frac{1-\delta}{\delta}$ we see that beliefs converge to a Brownian motion. More precisely, we see that there is one Brownian motion, and the probabilities of each of the N states are driven by different weights of these states. Price is a smooth function of beliefs, so price also follows a Brownian motion. Explain more here later.

Multiple Informed Traders

In this section, I will study the game when there are multiple informed traders with the same information.

Pump-and-Dump

If there are two traders with the same information, the pump-and-dump strategy will still be an equilibrium.

Consider the same indifference condition when the beleifs are low. The trader needs to be the same between revealing the state to be bad today (by selling the asset), and claiming the state to be good today (by buying) to let the price drift up and revealing the state to be bad tomorrow. Call π_1 and π_2 the probability with which each player sells when the state is bad. Suppose that the market takes either trader selling to be fully revealing that the asset is bad.

$$P(\mu)-P(0) = (1-\pi_2)\left(P'(\mu)\mu(\pi_1+\pi_2) + \delta\left(P'(\mu)\mu(\pi_1+\pi_2) + P(\mu) - P(0)\right)\right) - \pi_2\left(P(\mu) - P(0)\right) \tag{2.59}$$

On the left hand side, the value of revealing the state to be bad is the same, $P(\mu) - P(0)$. The value of waiting one more period now had two parts. There is a chance the other player will reveal the state today and you will lose money because you are long, $-\pi_2\left(P(\mu)-P(0)\right)$. If the other trader doesn't reveal the state first $(1-\pi_2)$, then the payoff is the same. You get the positive drift today which depends on the probabilities of revealing a bad state $(P'(\mu)\mu(\pi_1+\pi_2))$, plus you get the payoff from revealing tomorrow $\delta\left(P(\mu)-P(0)\right)$.

$$\left(P(\mu)-P(0)\right)\left(1-\delta+\pi_2(1-\delta)\right)=(1-\pi_2)(1+\delta)P'(\mu)\mu(\pi_1+\pi_2) \tag{2.60}$$

Rearranging a little, we can get a more convinient expression.

$$\pi_1+\pi_2=\frac{P(\mu)-P(0)}{P'(\mu)\mu}\frac{1-\delta+\pi_2(1-\delta)}{(1-\pi_2)(1+\delta)} \tag{2.61}$$

Now, suppose that player 2 is using a Poisson pump-and-dump strategy as before. We will see that it will also be optimal for player 1 to use a Poisson process. Let $\pi_2=\lambda_2(1-\delta)$. Now we can see that for large δ, π_1 will also be linear in $1-\delta$.

$$\pi_1+\pi_2\approx\frac{P(\mu)-P(0)}{P'(\mu)\mu}(1+2\lambda_2)(1-\delta) \tag{2.62}$$

Call $\pi_1=\lambda_1(1-\delta)$. Now we can solve for the equilibrium arrival rate. We can find the symmetric equilibrium by letting $\lambda_1=\lambda_2$.

$$2\lambda\left(1-\frac{P(\mu)-P(0)}{2P'(\mu)\mu}\right)=\frac{P(\mu)-P(0)}{2P'(\mu)\mu} \tag{2.63}$$

Solving for λ gives the arrival rate of the Poisson procces.

$$\lambda=\frac{1}{2}\frac{\frac{P(\mu)-P(0)}{2P'(\mu)\mu}}{1-\frac{P(\mu)-P(0)}{2P'(\mu)\mu}} \tag{2.64}$$

More Revealing Equilibria

Even though the pump-and-dump scheme is an equilibrium just as in the one trader model, it is no longer the case that it will be the best equilibrium. This equilibrium gives a payoff to the informed traders of $\min\{P(\mu) - P(0), P(1) - P(\mu)\}$ as before.

Consider another equilibrium. If the state is good, both informed players will buy the asset in the first period. If the state is bad, both informed players will sell the asset in the first period. When both traders buy the market maker's beliefs move to one. When both traders sell the market maker's beliefs move to zero. If the two informed traders do different actions (off path), the market maker's beliefs go to zero if $P(\mu) - P(0) \leq P(1) - P(\mu)$ and one otherwise.

Rather than the informed traders getting the minimum of the value of sending beliefs to zero or one, they get the average in expectation. This wasn't an equilibrium when there was only one informed player. With only one informed player if the price was low today, when the state is bad the trader would like to deviate to buying and pretending the state is good. The price is already low, so revealing the state to be bad only moves price a little but revealing the state to be good moves the price a lot.

This deviation is no longer profitable when there are multiple informed traders. When the state is bad, you know the other informed trader is going to reveal the state to be bad. So if you try to deviate by buying, the price will still fall to $P(0)$ because the other informed trader revealed it to be bad and you will lose money because you are long. Essentially, since the other informed trader is revealing information that will move the price, you always want to match what they are doing. You need to buy when they buy (and make the price go up) and sell when they sell (and make the price go down). This gives to strictly positive profits where any deviation is going to give you negative profits.

In fact, this arguement reveals that there are many move equilibria that can be supported. Many of these give an even higher payoff to the informed traders. If beliefs move up to $\mu_{up}(\mu)$ after both informed traders buy, down to $\mu_{down}(\mu)$ after both traders sell, and they move to whichever of those beliefs is farther from μ^* when the informed traders take different actions (off path), the an informed trader always has an incentive to match whatever action the other informed trader is doing. There will be no profitable way to deviate. If the traders are playing a mixed strategy, this requires them both to be able to see one randomization device. The informed trader needs to know what the other informed trader is doing (or supposed to do in equilibrium). The most profitable equilibrium would of course then be the same strategy and payoff the single informed trader used whith commitment.

They informed traders aren't colluding. There are no trigger strategies to prevent deviations. They are simply coordinating, and it's never profitable to go against the grain of what the other informed traders are doing. We can see that if there are any number of informed traders greater than one, and they can coordinate (see a common randomization device), the best payoff is going to be the same as with commitment in the single player case and lead to prices following a Brownian motion. Coordiantion becomes a substitute for commitment.

Other Information

A natural question would be how this model could generate both the Brownian motion and Poisson jumps at the same time. After all, this is what we seem to see in the data and option pricing models with jumps typically have a Brownian motion with jumps not just the jumps by themselves. This is achieved by having multiple informed traders with independent pieces of information. Say the asset is Amazon stock. There may be one trader that can obtain verifiable (commitment) information about the aquisition of Whole Foods, and another trader with unverifiable (no commitment) information about the cloud computing services. Both of these pieces of information may be relevent to the value of Amazon stock, but they don't necessarily need to be correlated.

I will again explain this using a linear price function. Suppose that there are N independent pieces of information, $\omega_i \in \{0,1\}$, relevent to the value of the asset.

$$P(\mu_1, \mu_2, \ldots, \mu_N) = \sum_{i=1}^{N} \mu_i z_i \tag{2.65}$$

Let's solve the problem of some trader that knows the value of ω_j, when there may potentially be other traders that know the other pieces of information.

Trader j holds fixed the stochastic processes for μ_i for all $i \neq j$, and chooses the optimal process for μ_j.

$$V_j(\mu_1, \mu_2, \ldots, \mu_N) = \max_{\Delta(\mu_j') \in \Delta([0,1])} \mathbb{E}\left[\left(P(\mu_1', \mu_2', \ldots, \mu_N') - P(\mu_1, \mu_2, \ldots, \mu_N)\right) x_j + \delta V_j(\mu_1', \mu_2', \ldots, \mu_N' \tag{2.66}$$

subject to Bayes plausibility and incentive compatability if solving the non-commitment problem. Trader j is also not allowed to correlate their strategy for μ_j with any other μ_i as they don't know those pieces of information and they are independent.

With the linear price function and independence, the problem simplifies.

$$V_j(\mu_1, \mu_2, \ldots, \mu_N) = \max_{\Delta(\mu_j') \in \Delta([0,1])} \mathbb{E}[(\mu_j' - \mu_j) z_j x_j] + \sum_{i \neq j} (\mathbb{E}[\mu_i'] - \mu_i) z_i x_j + \delta \mathbb{E}\left[V_j(\mu_1', \mu_2', \ldots, \mu_N')\right] \tag{2.67}$$

The beliefs about each piece of information must be a martingale. Regardless of what strategy the other players use, each term in the sum is zero from trader j's perspective.

$$V_j(\mu_1, \mu_2, \ldots, \mu_n) = \max_{\Delta(\mu_j) \in \Delta([0,1])} \mathbb{E}[|\mu'_j - \mu_j | z_j] + \delta \mathbb{E}[V(\mu_1, \mu_2, \ldots, \mu_N)] \tag{2.68}$$

Hence while prices and profits are quite different and will move randomly outside the control of player j, the strategy and expected profits of player j remain the same as in the version with only one strategic informed trader. This holds rather the other $N-1$ pieces of information are being release by other strategic players of if they are just randomly arriving by some exogenous process.

2.6 Conclusion

Literature Review

The topic of how prices move in an efficient capital market goes back to early work by Paul Samuelson and greatly expanded with Eugene Fama's papers, see Samuelson (1965) and Fama (1970). I study this by merging the literature of strategic informed trading with that of information design.

Questions of strategic informed trading often studied using a variant of Kyle (1985). Many of these study whether the trader's private information gets full incorporated into the price. Some examples of this are Holden and Subrahmanyam (1992), Foster and Viswannathan (1996), Ostrovsky (2012), and many others. Other models of prices in financial markets with informed strategic trading that don't use Kyle (1985) usually derive from Glosten and Milgrom (1985), Hellwig (1980), or Hanson

(2003, 2007). Van Bommel (2003) extends a Kyle model to allow for false information and rumors.

My model solves for the maximum amount of volatility in prices cause solely by information. Shiller (1981) shows empirically how prices are seemingly too volatile to be explained by information alone.

My model without commitment becomes a model of cheap talk popularized in Crawford and Sobel (1982). I use techniques for solving this like those in Lipnowski and Ravid (ming). They are able to simplify cheap talk games considerably when the sender's preference does not depend on the state.

My model with commitment is akin to a Bayesian persuasion model like Kamenica and Gentzkow (2011) and going back to Aumann and Maschler (1995). Dynamic information design models (like mine) are popular recently, Ely (2017), Renault et al. (2017), Orlov et al. (2019), and Hörner and Skrzypacz (2016). My model with commitment generalizes the results of Ely et al. (2015) where they study how much an informed agent can surprise and uninformed player.

I use tools developed by Mertens and Zamir (1977) who study the maximal variation of a bounded martingale. This gives results to the model with commitment similar to De Meyer and Saley (2003); De Meyer (2010) who seek a strategic foundation for Brownian motion in finance.

My results from the model without commitment look a lot like the solution of Zhong (2020), who studies dynamic information aquisition. They show how a Poisson process gives the most uncertainty across time, and that eponentially discounting agents are risk loving over lotteries across time. This same logic applies to the strategy of my trader without commitment power.

Chapter 3

Equilibrium Gerrymandering

Abstract

Through gerrymandering, a state drawing congressional districts can have a large effect on who gets elected. This in turn affects the policy chosen by elected representatives. This paper studies the optimal gerrymandering in an equilibrium of the fifty states electing members of the United States House of Representatives. First I find the optimal districting strategy when a party seeks to maximize the expected number of seats they win. This strategy always employs "cracking" (splitting up the opponent's base to spread them out in many districts) and it sometimes employs "packing" (cramming one district full of exlusively the opponent's base.) The optimal strategy can be found using techniques from information design. When the district drawer seeks to maximize the welfare of the state's citizens the care not just about the average seats won by each party, but the entire seat-vote curve. A seat-vote curve is a graph of the fraction of seats in congress that go to a political party against the fraction of votes obtained by that party. The national social optimal is for each

state to have a seat-vote curve that is less responsive (flatter) than proportional (45 degree line). However, each state has an incentive individually to choose a highly responsive seat-vote curve to disproportionately swing policy in their favor. In equilibrium each state chooses an extreme seat-vote curve close to a winner-take-all election. This is a prisoner's dilemma situation where every state is worse off in equilibrium, but it is the dominant strategy of each state to choose a highly responsive seat-vote curve. I then empirically estimate the seat-vote curve for each state and observe a few motivating facts. First, seat-vote curves are highly responsive. Every state's seat-vote curve has a slope much steeper than one (the "proportional" seat-vote curve). Second, the size of the state is predictive of the responsiveness. Smaller states have steeper curves.

3.1 Introduction

Through gerrymandering, the group drawing politcal districts can affect who gets elected to congress. Getting different people elected will change the policy chosen and have a national impact.

The House of Representatives of the United States Congress has 435 members. Each state is allocated some number of these representatives to elect based on its population. The state of Minnesota gets to elect eight representatives. The state is divided up into eight districts and each of the districts elects one of the representatives by popular vote. How the state is divided into districts will have a large impact on who will win the elections.

For example, suppose you could perfectly predict how everyone would vote and the state population is 51 percent Democrat and 49 percent Republican. You could put all the Republicans in district 1 through 4 and put all the Democrats in district 5 through 8. Then the representatives from the state would be half Republicans and half Democrats. Another possible districting would be to make each of the eight districts perfectly representative of the whole state. If each district was 51 percent Democrat and 49 percent Republican, the all of the elected representatives would be Democrat.

I take a simple model of voters (largely the same as Coate and Knight (2007)) and find the optimal districting strategies. There are three types of citizens in the state with policy preference distributed on the interval $[0, 1]$: Democrats (0), Republicans (1), and Independents $\in (0, 1)$. The Independents' preferences follow some distribution, but the mean of the distribution in a given election is unknown. Thus the districter doesn't know how many Independents will vote for each party. The objective of the districter plays a large role in designing the optimal districts.

In many states the districts are chosen by the state congress. Here the majority party can attempt to maximize the number of seats their party will win in the next election. This is the controversial gerrymandering that is frequently seen in court cases on the news. In section three, I will show the optimal way to do this partisan gerrymandering in this model.

The optimal districting plan has three types of districts. First, there are districts the party is garunteed to win. The districter fill these districts fifty percent (plus ϵ) of the way with voters from their own party to garuntee the win. Any more voters from their own party would be wasteful. The rest of this district is filled up with the opposition party's voters (if there are enough. If not, use Independents after.) The second type of district are the ones the party is garunteed to lose. The districter only need to have any of these if the opposing party has a larger base than their own party. Since these districts are going to be lost, they might as well be filled 100 percent with the opposition party. The third type of district are the competitive ones. These districts will be filled exclusively with Independent voters and thus will be equally likely to vote for either party. I will show in section 3 that this districting scheme maximizes the expected number of seats the districter's party will win.

However, in many states, the districts are chosen by a non-partisan committee or judicial branch. For these states, I model the districter as trying to maximize social welfare in the state. The policy that maximizes welfare (and thus the represetatives you want elected) depend on the particular realization of the Independent distribution in each election. If more of the Independents are Democrats this year, the districter would like to elect more Democrats and have a policy that is more Democrat. For this reason, the districter doesn't care only about the expected number of seats won by each party but rather about the entire seat-vote curve.

A seat-vote curve is the expected fraction of elected representatives that come from the Democratic party graphed against the fraction of the state-wide vote that was for the Democratic party. It is the proportion of seats the party wins graphed against their vote share. An intuitive example of a seat-vote curve is the 45 degree line. That is, a line with an intercept of zero and a slope of one. With this curve, if the Democratic party wins x percent of the votes in the states they will win that same x percent of the seats in congress from the state. Generally, the curve may be non-linear, steeper, flatter, or even biased toward one party or the other.

In section 4, I solve for the equilibrium seat-vote curves. While the preferred policy reacts modestly to changes in the vote share, the equilibrium effects lead to highly responsive seat-vote curves. Essentially, when a state wants policy to move by one percent to the left they need one percent more of congress to be on the left. But, the state only chooses one-fiftieth of congress on average. Thus, they need to elect fifty percent more of their state's representatives from the left.

This leads to every state having a highly responsive seat-vote curve. In section 5, I show that this is not what will be socially optimal for the country as a whole. Political districting is a prisoners dilemma between the states. Total welfare is optimized when all states have modest seat-votes curves. However, each individual state has an incentive to deviate to a highly responsive curve. When each state chooses highly responsive curves, everyone is made worse off.

In section 6, I estimate the seat-vote curves in each of the 50 states and present a few stylized facts. The first fact is that seat-vote curves highly responsive. Rather than having a slope of one, on the interval 45 percent to 55 percent of the vote, the average slope is around four. If one percent more of the state votes Democrat, about four percent more of the elected representatives will be Democrats. The second fact is that the responsiveness has a strong negative correlation with the size of the state. Smaller states have especially steep seat-vote curves. I show that the inverse of the number of representatives from a state is a good predicter of seat-vote curve responsiveness, The responsiveness is not predicted by how much control a political party has in the state. The responsiveness is partially explained by who was in charge of drawing the districts (state congress, bipartisan committee, courts, etc.) I take these facts as validation of the model.

3.2 Model

Here a present a model of voters help understand optimal gerrymandering and explain the stylized facts in the data presented in section 6. The model is very similar to Coate and Knight (2007), but extended to have many states.

Model Setup

The payoff relevant object of interest is the policy chosen. A policy is a number that lies in the interval $[0, 1]$. Think of 0 as the preferred policy of Democratic party and 1 as the preferred policy of the Republican party. After the representatives from all states are elected, the policy will be the average of the representatives' preferences. In section 7, I explore the model where the policy chosen is equal to that of the median representative and the main results are qualitatively similar.

The strategic players in the game are the states' district designers. There are fifty states indexed by $i \in \{1, 2, \ldots, M\}$. Each state is characterized by the distribution of voters in the state. The districter will choose how to split the distribution of voters into different districts. Through this, they affect who gets elected and what policy is chosen. Two different objectives for the state district designer will e treated seperately. In section three, I will study a districter that seeks to maximize the expected number of seats for a given party. In section four, I will study a districter that seeks to maximize the aggregate welfare of the citizens in the state.

Every voter has a private preference over the policy chosen. In each state, there are three different groups of voters. There is a mass π_{Di} for state i of Democrats with a preferred policy $\theta = 0$. There is a mass π_{Ri} for state i of Republicans with a preferred policy $\theta = 1$. There is a mass π_{Ii} for state i of Independents with preferred policy distributed over $\theta \in [0,1]$. Call m_i the mean preference among the independent voters, and call $2\tau_i$ the width. The independent voters have a preferred policy uniformly distributed on the interval $[m_i - \tau_i, m_i + \tau_i]$. This mean, m_i, is unknown for every to the districter. The districter is also unable to distinguish between independent voters that lie in different parts of the interval. At the time of the election the mean is drawn from a uniform distribution $m_i \sim U\left(\left[\frac{1}{2} - \tau_i, \frac{1}{2} + \tau_i\right]\right)$. This means that the fraction of independents that will vote Democrat in a given election is uniformly distributed over $[0,1]$.

Each voter wants the policy to be as close as possible to their preferred policy. They face a quadratic loss funtion. If a voter's preferred policy is $\hat{\theta}$ and the policy chosen is θ, the voter recieves a payoff of $-\left(\theta - \hat{\theta}\right)^2$. The independent voters are not strategic in their choice. A voter will vote Democrat if their preferred policy is less than or equal to $\frac{1}{2}$. Otherwise, they will vote Republican.

3.3 Partisan Districting Plans

In many states, the party in control of state congress can draw the districts. In this section, I will study what a districter would do if they wanted to maximize the expected number of seats that a given party will win. This isn't really an equilibrium problem. What districts other states draw, and who they elect doesn't enter into the objective in any way. We can simply solve a state's districting problem in isolation.

Example

Consider a state that has 50 percent Democrats, 25 percent Republicans, and 25 percent Independents. The intuitive outcome is that the Democrats should get 50 to 75 percent of the seats in congress and the Republicans should get 25 to 50 percent of the seats, depending on how the Independents vote. If the districts are drawn to maximize the expected number of seats for either party, the outcome will look very different. First, if the Democrats are in charge of drawing the districts, they can win all the seats in this model. All they need to do is make every district look just like the state as a whole. Each district will be 50 percent Democrat, 25 percent Republican, and 25 percent Independent. It doesn't even seem like an extreme gerrymender on the face of it. They simply make every district identical and representative of the state. However, since the Democrats are now garunteed to have at least 50 percent of the vote in every district, they will win all the congressional seats.

What if the Republicans are in charge of drawing the districts. They can use a concept called "packing" to win some of the seats. "Packing" is when you group together supporters of the opposition into a single district. If you're going to lose a district, you might as well lose big. They can simply puting all the Democrats together in the first half of the districts. Then in the other half of the districts they can employ a strategy called "cracking". "Cracking" is when you split up your own party's supporters so you can win more districts. If you're already going to win a district, each additional vote you get in that district is wasted. In the second half of the districts they can make each district an equal mix of Republicans and Independents. Since they have garunteed at least half the votes in each of these districts, they will win them all. This districting strategy is simple and it gets the

Republicans half the seats.

However, the Republicans can do even better than this. It might seem like since half the state is certainly going to vote Democrat, Republicans can never win more than half the seats, but they can. For simplicity, assume that ties are broken in the Republican's favor. Throughout the section, I will always assume that ties are broken in favor of the party drawing the districts. Otherwise they would simply need to put one additional voter from their party in each of this districts to insure a win.

Consider the following districting. In one quarter of the districts they follow a packing strategy. One quarter of the districts are made up entirely of Democrats. Then in half of the districts they use a cracking strategy. They do it a little differently than before. If you're going to win a district for sure, you might as well make sure the rest of the votes are against you. You don't want to waste any votes. So, in this half of the districts there will be an equal mix of Republicans and Democrats. The Republicans will win all of these districts. Finally, the remaining quarter of the districts are made up entirely of Independents. On average Republicans will win half of these districts.

This plan gives the Republicans five-eighths of the seats in congress on average. This is the optimal districting plan in this example. If a state used to be heavily Republican, they may have a majority of Republicans in their state congress. This districting shows how they can maintain their majority in congress even if the population in the state changes to be heavily Democrat.

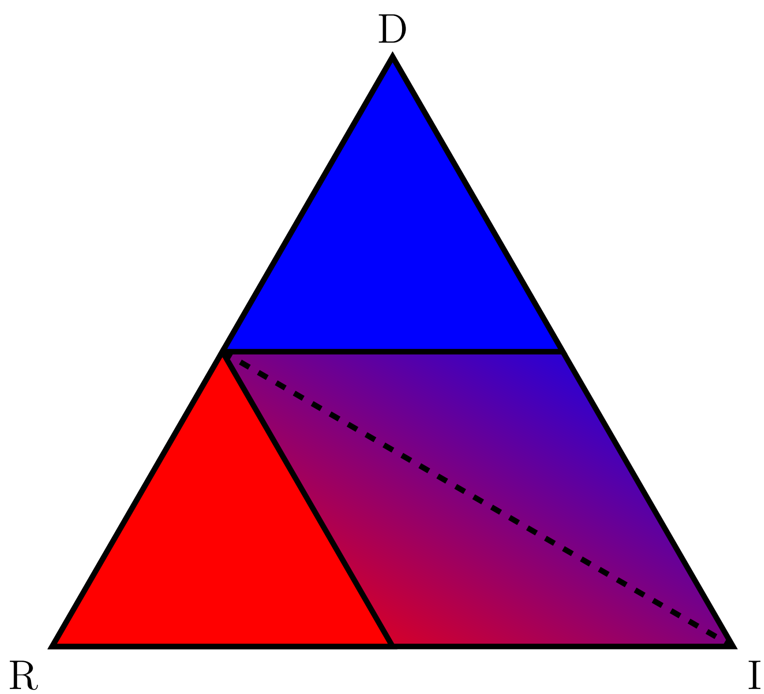

Figure 3.1: The distribution of voters in the state.

Optimal Districting

Generally, I will show that the optimal districting for maximizing the expected number of seats for a given party will have the same form as the example.

There are three types of voters. The distributions of these voters can be represented by a simplex as in figure 3.1. This problem is very similar to a Bayesian Persuasion problem. In Bayesian persuasion, the sender is looking at ways to split a up distribution (the prior beliefs) into multiple other distributions (posterior beliefs) that must average out to the first distribution (by Bayes Rule). In this gerrymandering problem, the districter is looking at ways to split up a distribution (the population of voters in the state) into multiple other distributions (district populations) that add up to the state population. Just like in Bayesian Persuasion, the optimal value can be thought of as a “concavification” on this simplex.

Suppose the districter wants to maximize the expected number of seats for the Republican party. The distributions in figure 3.1 can be broken up into a few segments. The lower left triangle in the figure (red) are distributions where the Republicans have at least 50 percent of the population. They would win these districts with certainty. The upper triangle (blue) are distributions where Democrats have at least 50 percent of the population. They would lose these districts with certainty. The rhombus of remaining distributions (purple) are distributions where either party could win depending on how the Independents vote. The distributions along the dashed line are equally likely to be won by Republicans or Democrats. The likelihood of Republicans winning transitions from one to zero as you move up and to the right through this region.

To think about concavification, picture this triangle as a 3D object with the colors representing height. The lower left triangle (red) has a height of one. The upper triangle (blue) has a height of zero. The rhombus (purple) linearly connects the two triangles. Now imagine taking a cloth, laying it on top of the object, and pulling the edges down tight. This is the concavification of the function. The height of this concavification is the maximum expected number of seats Republicans can win with the optimal districting. The points where the concavification is equal to the original function (where the cloth is touching the 3D object) are the distributions that are used in the optimal districting schemes. If the state’s population is on one of those points, it is optimal to make every district have a distribution identical to the state as a whole. If the state’s population is not one of those points, the optimal districts with break up the state into different districts that are all among those points.

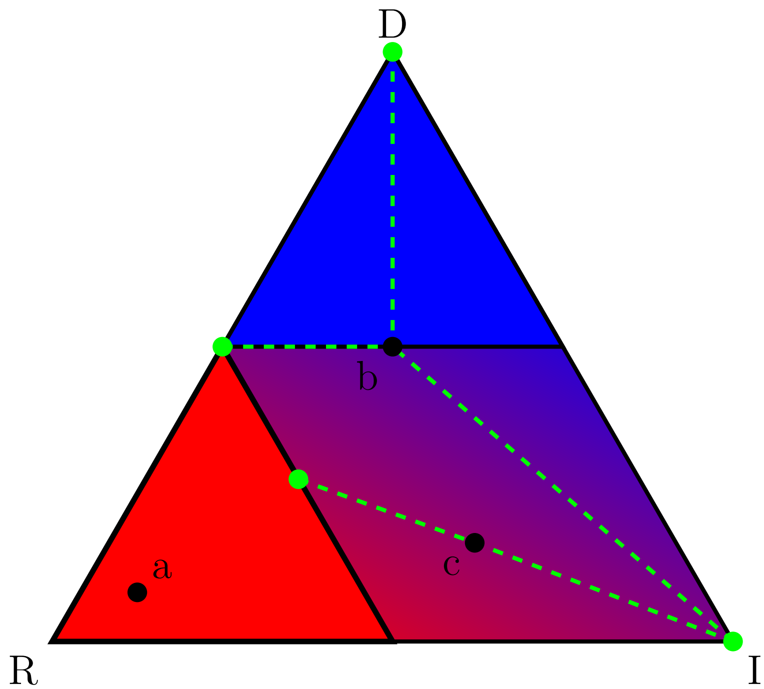

Figure 3.2: Splitting the state's population into optimal districts.

These optimal points consist of the entire lower left triangle (Republicans greater than 50 percent), and the other two corners of the simplex (Democrats have 100 percent and Independents have 100 percent). Any other distribution is dominated by a combination of these distributions. If your state is in the lower left triangle (Republicans have more than 50 percent), you don't have to do much of anything. The optimal districting will have all districts identical to the state population. The Republicans will win every district. If your state population is in the purple rhombus below the dashed line (more Republicans than Democrats), the optimal splitting is a combination of the bottom right corner (100 percent Independents) and the point on the edge of the lower left triangle (50 percent Republicans) that is straight across the prior from the corner. Republicans will win all of the latter districts and have a 50 percent chance of wining the former districts. If your state population is anywhere above the dashed line (more Democrats than Republicans) the optimal splitting is a combination of the point where the lower triangle meets the upper triangle (50 percent Republican and 50 percent Democrat), the bottom right corner (100 percent Independent), and the top corner (100 percent Democrat). Republicans will win all of the first type of districts, have a 50 percent chance in the second type of districts, an lose all of the last type of district. Each of these scenarios is graphed is shown in figure 3.2.

The optimal districting strategy can be stated simply as follows. You start with the districts you are going to win for sure. Fill as many districts as possible half full with Republicans. Fill the other half of those districts with Democrats. If there aren't enough Democrats, continue filling with Independents. If there aren't enough Independents, fill remaineder with Republicans. After filling these districts entirely, all remaining Democrats and Independents are separated into the reset of the districts. So, there are districts that are 50 percent Republican, and 50 percent Democrat or Independent (Democrat is preferred). There are districts that are 100 percent Democrat. Finally, there are districts that are 100 percent Independent.

Proposition 3. *The maximum expected fraction of seats Republicans can win is*

$$v_R = \begin{cases} 1 & \text{if } \pi_R \geq \frac{1}{2} \\ 2\pi_R + \frac{1}{2}\left(\pi_I - (\pi_R - \pi_D)\right) & \text{if } \pi_D < \pi_R < \frac{1}{2} \\ 2\pi_R + \frac{1}{2}\pi_I & \text{if } \pi_R \leq \pi_D. \end{cases} \tag{3.1}$$

This can be rewritten as

$$v_R = \min\left\{1\ ,\ 2\pi_R + \frac{1}{2}\left(\pi_I - \max\left\{0,\ \pi_R - \pi_D\right\}\right)\right\}. \tag{3.2}$$

If there are fewer Republicans than Democrats, you fill as many districts as possible with half Republicans and half Democrats. You will win all these. There are $2\pi_R$ of these districts. Then you have all the Independents in their own districts. You win half of these. There are π_I of these districts. Finally, the rest of the Democrats are in their won districts and you don't win any of those. Hence, your expected number of seats is $2\pi_R + \frac{1}{2}\pi_I$.

If there are more Republicans than Democrats, you do the same strategy. The difference is that you now have to put some Independents in the first type of district because there aren't enough districts. From the payoff you need to subtract off the Independents put in the first type of district so they aren't double counted.

This will continue until you are able to fill all the districts half way with Republicans. Then the payoff is flat at one. Any additional Republicans are superfluous.

3.4 Non-Partisan Districting Plans

In many states, the districts are not drawn by the state congress. Rather a non-patisan committee or even the judicial branch may be the ones drawing the districts. In these cases we wouldn't expect the districts to be maximizing the expected seats of either party. For this section, consider what the optimal districting strategy would be if the districter seeks only to maximize the utility of the state's citizens.

This problem needs to be solved very differently from that of the partisan districter. All that the partisan districter cared about in the previous section was an average. The non-partisan districter will care about all the outcomes. I will show you why. The non-partisan districter cares about the preferred policy of all of the citizens. If the median independent voter, m_i, turns out to be very high in a given election, this means that the citizens prefer more of a Democrat policy. Thus, the districter would like to have more Democrats elected. Whereas if m_i is very low, the districter would like to have more Republicans elected. The districter's optimal fraction of seats is different for each election outcome.

The point of different districting strategies is to induce a certain seat-vote curve. The seat-vote curve is the expected fraction of congressional seats a party wins as a function of the party's vote share in the state. Consider a state in which 40 percent are Democrats, 40 percent are Republicans, and 20 percent are Independents. You could stick all the Democrats together in the first 40 percent of the districts, all the Republicans in the next 40 percent of districts, and the Independents in the last 20 percent of districts. Then regardless of what the independents do, at least 40 percent of seats will be won be Democrats and at least 40 percent won by Republicans. This means the seat-vote curve will start with a high intercept (.4), increase only very modestly, and max out at a low height (.6). This seat-vote curve is not very responsive. Alternatively, you could make every district look demographically like the state as a whole. That is, each district would be made up of 40 percent Democrats, 40 percent Republicans, and 20 percent Independents. Then, all the districts could potentially be flipped based on what happens to the Independents. This seat-vote curve would have a minimum of 0, a maximum of 1, and be very steep in the middle region. This is a highly responsive seat-vote curve. The most intuitive idea people think of is that congressional seats should match the distribution of the voters. So, if Democrats get 54 percent of the vote they should get 54 percent of the seats in congress. This represented by a seat-vote curve that is simply a 45 degree line. The curve has a slope of one everywhere and an intercept of zero.

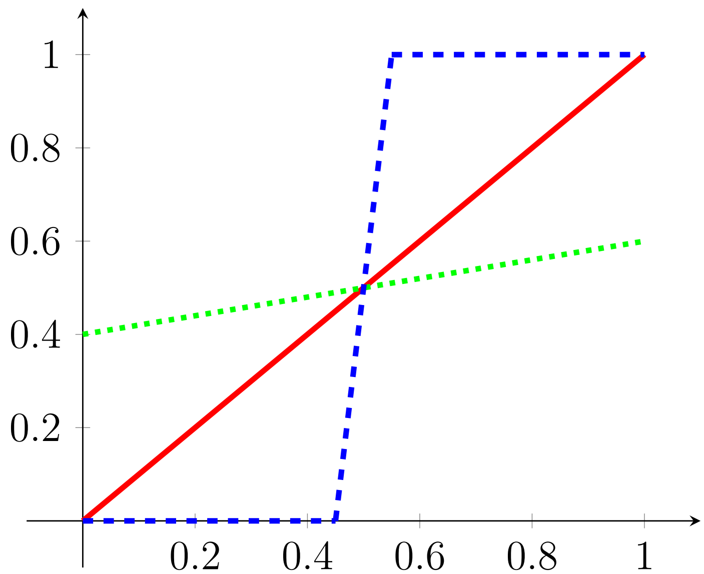

Figure 3.3: Three different seat-vote curves implemented by gerrymandering. The solid line is the "proportional" seat-vote curve. The dotted line is a less responsive seat-vote curve with safe seats for both parties. The dashed line is a highly responsive seat-vote curve.

Rather than working directly with how to split the distribution into districts, we will assume that the state simply chooses a seat-vote curve. We won't worry for right now about what districting is used to implement that curve.

Call v_i the fraction of voters that vote Democrat in state i. This is a combination of how many Democrats, Republicans, and Independents are in the state, as well as where the median Independent voter is.

$$v_i = \pi_{Di} + \pi_{Ii} \left(\frac{\frac{1}{2} - (m_i - \tau_i)}{2\tau_i} \right) \tag{3.3}$$

The seat-vote curve $S : [0, 1] \to [0, 1]$ is a mapping from the fraction of votes that Democrats won, v_i, to the fraction of representatives Democrats get in congress, $S_i(v_i)$.

After the election, the policy is chosen collectively by the representatives from all

50 states. Let N_i be the number of representatives assigned to state i. The fraction of congress comming from state i is $n_i = \frac{N_i}{\sum_{j=1}^{M} N_j}$. If each state, i, elected $S_i(v_i)$ Democrats, then congress will have a fraction $S = \sum_{i=1}^{M} n_i S_i(v_i)$. The members of congress choose the policy that maximizes their welfare. Policy is set equal to the average of all the representatives preferences. So, the chosen policy, θ is equal to $1 - S$.

This is a simulateous move game where each state chooses a seat-vote curve, $S_i(v_i) \in [0,1]\ \forall v_i \in [0,1]$, to maximize the welfare of its voters,

$$\max_{S_i(v_i)} \quad -\mathbb{E}_{m_1,m_2,\ldots,m_M}\left[\pi_{Di}\left(\theta\right)^2 + \pi_{Ri}\left(1-\theta\right)^2 + \pi_{Ii}\int_{m_i-\tau_i}^{m_i+\tau_i}\left(\theta - x\right)^2 \frac{dx}{2\tau}\right] \tag{3.4}$$

where θ, the policy, is equal to one minus the average fraction of seats chosen by the states for Democrats.

$$\theta = 1 - \sum_{i=1}^{M} n_i S_i(v_i) \tag{3.5}$$

We will find the unique Nash equilibrium of this game.

Optimal Seat-Vote Curve

Before solving explicitly for the solution of this game, let us breifly examine the desired policy of a state. If the voters of a state have individual preferences, $\tilde{\theta}$, distributed according to G_i, then the state's preferred policy solves

$$\max_{\hat{\theta}} \quad -\int\left(\tilde{\theta} - \hat{\theta}\right)^2 dG_i(\tilde{\theta}). \tag{3.6}$$

So, the state would like the policy to be equal to the average preference in the state.

$$\Rightarrow \hat{\theta}^* = \mathbb{E}_{G_i}\left[\theta\right] \tag{3.7}$$

If every voter was either a pure Democrat or pure Republican living on the extreme, $\pi_{Ii} = 0$, then the average preference is exactly equal to the vote share. This is where the idea of a proportional seat-vote curve with a slope of one comes from. However, a one percent increase in the population voting Democrat is not from one percent of the population that were staunch Republicans and are now suddenly staunch Democrats. The one percent increase in votes comes from one percent of the population that was near the center but leaning slightly Republican and is now near the center and leaning slightly Democrat. It is a much smaller shift in aggregate preference. Thus, optimal policy would shift toward the Democrats by less than one percent. This is a motivation for a flatter seat-vote curve.

In this model, the state's preferred policy is equal to the seat-vote curve derived in Coate and Knight (2007).

$$\hat{\theta}^* = \mathbb{E}_{m_i}[\theta] \tag{3.8}$$

$$= \frac{1}{2} + (\pi_{Di} - \pi_{Ri})\left(\frac{1}{2} - \tau_i\right) + 2\tau_i\left((v_i - \frac{1}{2}\right) \tag{3.9}$$

Even though the shift in desired policy is mild (slope of $2\tau_i$ in the above equation), the strategy a state must take to impliment that shift is extreme. We can see this by solving for state i's best response function in the game. Once the state sees the vote share, v_i, they no longer face any uncertainty about their own voters' preferences. This means that the optimal seat-vote curve can be solved pointwise.

$$\max_{S_i(v_i)} -\mathbb{E}_{v_{-i}}\left[\pi_{Di}\left(1 - \sum_{j=1}^{M} n_j S_j(v_j)\right)^2 + \pi_{Ri}\left(\sum_{j=1}^{M} n_j S_j(v_j)\right)^2 \right. \\ \left. + \ \pi_{Ii}\int_{m_i-\tau_i}^{m_i+\tau_i}\left(1 - \sum_{j=1}^{M} n_j S_j(v_j) - x\right)^2 \frac{dx}{2\tau_i}\ \middle|\ v_i\right] \tag{3.10}$$

Now we differentiate the equation to find the local maximum.

$$\frac{\partial W_i(v_i)}{\partial S_i(v_i)} = 2n_i \mathbb{E}\left[\pi_{Di}\left(1 - \sum_{j=1}^{M} n_j S_j(v_j)\right)\right. \tag{3.11}$$

$$\left. -\pi_{Ri}\sum_{j=1}^{M} n_j S_j(v_j) + \pi_{Ii}\left(1 - \sum_{j=1}^{M} n_j S_j(v_j)\right) - \pi_{Ii}\int_{m_i-\tau_i}^{m_i+\tau_i} \frac{x}{2\tau_i}dx \right| v_i\right] \tag{3.12}$$

$$= 2n_i\left(\pi_{Di} + \pi_{Ii}(1-m_i) - \sum_{j=1}^{M} n_j \mathbb{E}_{v_j}[S_j(v_j)|v_i]\right). \tag{3.13}$$

Setting this equation equal to zero gives us the welfare maximizing share of seats for Democrats given the realization of the Independent voters.

$$\hat{S}_i(v_i) = \frac{1}{n_i}(\pi_{Di} + \pi_{Ii}(1-m_i)) - \sum_{j\neq i} \frac{n_j}{n_i}\mathbb{E}_{v_j}[S_j(v_j)|v_i]. \tag{3.14}$$

Here the seats are a function of m_i. Of course, we would like the seats to be a function explicitly of v_i to get our seat-vote curve. We can write m_i as a function of the vote share simply by inverting equation (3.3).

$$m_i = \frac{1}{2} + \tau_i\left(\frac{\pi_{Ii} + 2\pi_{Di} - 2v_i}{\pi_{Ii}}\right) \tag{3.15}$$

Plugging this in, we can get the seats as a function of the vote share.

$$\hat{S}_i(v_i) = \frac{1}{n_i}\left(\frac{1}{2} + (\pi_{Di} - \pi_{Ri})\left(\frac{1}{2} - \tau_i\right) + 2\tau_i\left(v_i - \frac{1}{2}\right) - \sum_{j\neq i} n_j \mathbb{E}_{v_j}[S_j(v_j)|v_i]\right). \tag{3.16}$$

If the state could choose any values for their seat-vote curve, this is what they would choose. However, they are restricted to choose a share between 0 and 1. This equation does not always lie in the interval $[0,1]$. In fact, it is usually outside the interval. Since the objective is quadratic, the solution for the best response function will still be very simple. With a quadratic objective, the optimum will just be at the boundary closer to the unconstrained optimum.

Proposition 4. *The equilibrium seat-vote curve for state i is*

$$S_i^*(v_i) = \begin{cases} 0 & \textit{if } \hat{S}_i(v_i) < 0 \\ \hat{S}_i(v_i) & \textit{if } \hat{S}_i(v_i) \in [0,1] \\ 1 & \textit{if } \hat{S}_i(v_i) > 0. \end{cases} \tag{3.17}$$

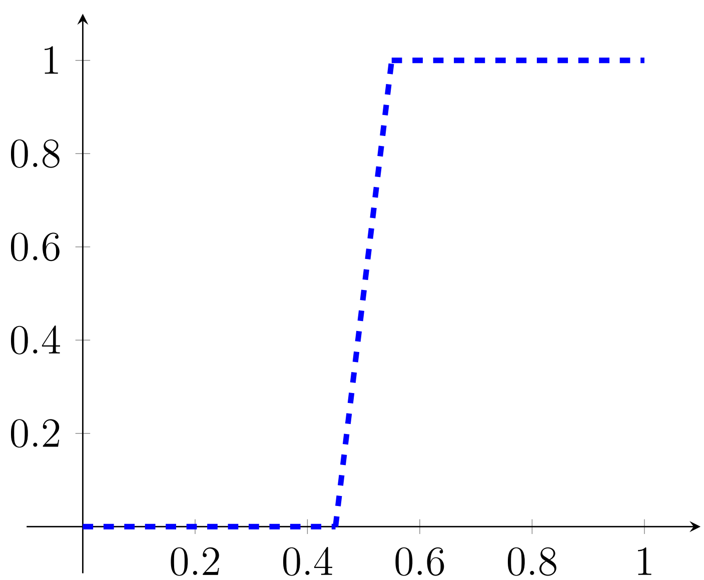

Figure 3.4: The optimal seat-vote curve.

Assume for now that m_i is drawn independently from m_j for every state j. We can see already that the seat-vote curve for state i is very steep in the middle.

$$\frac{\partial S_i^*}{\partial v_i} = \frac{2\tau_i}{n_i} \tag{3.18}$$

The first component, $2\tau_i < 1$, comes from the fact that a change in vote share represents a smaller change in average political preference in the state and the seat-vote curve should be less responsive. The second component, $\frac{1}{n_i}$, comes from the fact that state i is only a small part of the national congress and to get congress to shift by 1 percent the state needs to shift their representatives by an average of 50 percent. This is the same non-linear relationship we will see fits well in the empirical section. This is the main takeaway from this section. Equilibrium seat-vote curves are very steep and have a responsiveness proportional to the inverse of the number of representatives the state is alloted.

Equilibrium

I have not yet completely charactarized the equilibrium of the game. What we have is the best response of state i as a function of the seat-vote curves of all the other states. The equilibrium is the fixed point of all these functions. Notice in equation (3.16) that the other states' chosen curves only enter the best response in a very simple way. State i only cares about the expected fraction of seats the other states will elect from the Democratic Party, $\sum_{j \neq i} n_j E\left[S_j^*\right]$. This makes the system easy to solve.

For notational simplicity, let's call $s_i = \mathbb{E}[S_i^*(v_i)]$. Using equation (3.14) we see that the state will be in the interior portion of their seat-vote curve whenever

$$0 < \frac{1}{n_i}(\pi_{Di} + \pi_{Ii}(1 - m_i) - \sum_{j \neq i} n_j s_j) < 1. \tag{3.19}$$

Simplifying, the state will be in the highly responsive portion of the seat-vote curve as long as the median independent voter isn't too far to either extreme.

$$1 - \frac{1}{\pi_{Ii}}(n_i + \sum_{j \neq i} n_j s_j - \pi_{Di}) \ < \ m_i \ < \ 1 - \frac{1}{\pi_{Ii}}(\sum_{j \neq i} n_j s_j - \pi_{Di}). \tag{3.20}$$

Also, $S_i^*(v_i)$ will be equal to 1 whenever the median independent, m_i, is below that lower cutoff.

Now the goal is to take the expected value of equation (3.17) to get the average fraction of seats that go to Democrats in state i. There are three segments to the best response function we need to average over. On the first segment, Democrats get zero seats. So, this drops out of the equation. In the middle segment we integrate over the likelihood of each value. On the third segment, Democrats get all the seats. So, the contribution to the expectation is just the probability of being in this segment times 1. The likelihood of this is equal to the cdf of m_i at the cutoff. Since m_i is uniformly distributed on $[\frac{1}{2} - \epsilon_i, \frac{1}{2} + \epsilon_i]$, the probability that $S_i^*(v_i) = 1$ equals the following.

$$p_i = \frac{1 - \frac{1}{\pi_{Ii}}(n_i + \sum_{j \neq i} n_j s_j - \pi_{Di}) - (\frac{1}{2} - \epsilon)}{2\epsilon_i}. \tag{3.21}$$

Now we can compute the expected number of seats Democrats will win in state i.

$$s_i = \mathbb{E}[S_i^*(m_i)] \tag{3.22}$$

$$= \int_{1-\frac{1}{\pi_{Ii}}(n_i + \sum_{j \neq i} n_j s_j - \pi_{Di})}^{1-\frac{1}{\pi_{Ii}}(\sum_{j \neq i} n_j s_j - \pi_{Di})} \frac{1}{n_i}(\pi_{Di} + \pi_{Ii}(1-m) - \sum_{j \neq i} n_j s_j) \frac{dm}{2\epsilon_i} \tag{3.23}$$

$$+ \frac{1 - \frac{1}{\pi_{Ii}}(n_i + \sum_{j \neq i} n_j s_j - \pi_{Di}) - (\frac{1}{2} - \epsilon_i)}{2\epsilon_i} \tag{3.24}$$

$$= \frac{\pi_{Di} + \pi_{Ii} - \sum_{j \neq i} n_j s_j}{2\pi_{Ii}\epsilon_i} - \frac{\pi_{Ii}}{2\epsilon_i} \int_{1-\frac{1}{\pi_{Ii}}(n_i + \sum_{j \neq i} n_j s_j - \pi_{Di})}^{1-\frac{1}{\pi_{Ii}}(\sum_{j \neq i} n_j s_j - \pi_{Di})} m \, dm \tag{3.25}$$

$$+ \frac{1 - \frac{1}{\pi_{Ii}}(n_i + \sum_{j \neq i} n_j s_j - \pi_{Di}) - (\frac{1}{2} - \epsilon_i)}{2\epsilon_i} \tag{3.26}$$

$$= \frac{\pi_{Di} + \pi_{Ii} - \sum_{j \neq i} n_j s_j}{2\pi_{Ii}\epsilon_i} + \frac{n_i}{2\pi_{Ii}\epsilon_i}(1 - \frac{1}{\pi_{Ii}}(n_i + \sum_{j \neq i} n_j c_j - \pi_{Di})) \tag{3.27}$$

$$+ \frac{1 - \frac{1}{\pi_{Ii}}(n_i + \sum_{j \neq i} n_j s_j - \pi_{Di}) - (\frac{1}{2} - \epsilon_i)}{2\epsilon_i} \tag{3.28}$$

This gives a system of 50 equations and 50 unknowns. Solving for the 50 s_i terms that satisfy this system will finish the construction of the equilibrium. While this initially looks messy, we now have s_i written as a linear function of all s_j with $j \neq i$. We can write the equation simply as

$$As = b \tag{3.29}$$

where b is an $M \times 1$ vector with

$$b_i = \frac{1}{2} + \frac{2\pi_{Di} + \frac{3}{2}\pi_{Ii} - n_i^2 - n_i + n_i\pi_{Ii} - n_i\pi_{Ii}\pi_{Di}}{2\pi_{Ii}\epsilon_i} \tag{3.30}$$

and A is an $M \times M$ matrix with

$$A_{ij} = \begin{cases} \frac{\left(\frac{n_i}{\pi_{Ii}}+2\right)n_j}{2\pi_{Ii}\epsilon_i} & \text{if } i \neq j \\ 1 & \text{if } i = j. \end{cases} \tag{3.31}$$

This A matrix is always full rank. This means that there always exists a unique solution. The solution gives us the final constants we needed in the best response function. Thus, it completes the equilibrium.

3.5 Welfare

Even though the described strategies connstitute the only equilibrium, they don't give the highest total payoffs in the game. They aren't socially optimal. Let us construct the socially optimal policy if a planner was in charge of all 50 states. The objective function would look similar, but it would account for the citizens in every state.

$$\max_{\theta} \quad -\sum_{i=1}^{M} n_i \mathbb{E}_v \left[\pi_{Di}\left(\theta\right)^2 + \pi_{Ri}\left(1-\theta\right)^2 + \pi_{Ii}\int_{m_i-\tau_i}^{m_i+\tau_i} \left(\theta - x\right)^2 \frac{dx}{2\tau_i}\right] \tag{3.32}$$

We can differentiate this with respect to S to solve for the optimum.

$$\frac{\partial W}{\partial \theta} = 2\sum_{i=1}^{M} n_i \pi_{Di}(\theta) - 2\sum_{i=1}^{M} n_i \pi_{Ri}(1-\theta) - 2\sum_{i=1}^{M} n_i \pi_{Ii} \int_{m_i-\tau_i}^{m_i+\tau_i} (\theta - x) \frac{dx}{2\tau_i} \tag{3.33}$$

Evaluating the uniform integral and simplifying,

$$= 2\sum_{i=1}^{M} n_i \pi_{Di} + 2\sum_{i=1}^{M} n_i \pi_{Ii}(1 - \mathbb{E}[m_i]) - 2(1-\theta) \tag{3.34}$$

since $\sum_{i=1}^{M} n_i = 1$.

Putting the planner on the same footing as the states in the game, let them choose an entire seat-vote curve for the nation. This makes the expectation go away. Now, we can substitute in for m_i as before to get an expression in terms of vote shares, v_i. Setting this equal to zero gives the optimal policy.

$$1 - \theta^* = \sum_{i=1}^{M} n_i \left(\pi_{Di} + \pi_{Ii}(1 - m_i)\right) \tag{3.35}$$

$$= \sum_{i=1}^{M} n_i \left(\pi_{Di} + \pi_{Ii}\left(\frac{1}{2} - \tau_i \frac{\pi_{Ii} + 2\pi_{Di} - 2v_i}{\pi_{Ii}}\right)\right) \tag{3.36}$$

$$= \sum_{i=1}^{M} n_i \left(\left(\frac{1}{2} - \tau_i\right)\pi_{Ii} + (1 - 2\tau_i)\,\pi_{Di} + 2\tau_i v_i\right) \tag{3.37}$$

$$= \frac{1}{2} + \sum_{i=1}^{M} n_i \left((\pi_{Di} - \pi_{Ri})\left(\frac{1}{2} - \tau_i\right) + 2\tau_i\left(v_i - \frac{1}{2}\right)\right) \tag{3.38}$$

Assuming τ_i is the same in each state, this would be a linear function of the aggregate vote share with a slope equal to 2τ. It's also easy to impliment. Since it's linear, we would get this policy by each state having a seat-vote curve equal to

$$S_i(v_i) = \frac{1}{2} + (\pi_{Di} - \pi_{Ri})\left(\frac{1}{2} - \tau\right) + 2\tau\left(v_i - \frac{1}{2}\right). \tag{3.39}$$

with that same low slope. Notice that this is the same as equation (3.8), which is the same as in Coate and Knight (2007) (CK). This is where each state simply chooses their most preferred policy. If a state wants the policy to be 60 percent Democrat, then they just elect 60 percent Democrats.

Take the following example. Imagine that everyone in the room was to write down what temperature they would like room. Then the thermastat is set to the average of all the votes. The social optimum is obtained if everyone just honestly writes their most preferred temperature. However, you may have an incentive to write something different. If you like to the room to be 67 degrees and you think that's a little colder than most people like, you have an incentive write a vote that is much colder. You want the average of the votes to be 67. So, you might write something more extreme like 60 degrees.

In equilibrium, political districting is the same way. This mild seat-vote curve is what will maximize national welfare, but each state has an incentive to deviate. When a state gets a lot of votes for Democrats, they think their state is likely more Democrat than the average state. Then to make the policy a little more Democrat the state wants to elect a lot more Democrats. Regardless of what the other states are doing, each state has an incentive to deviate by playing a highly responsive seat-vote curve.

This is a prisoner's dilemma situation though, because in equilibrium every state is worse off. No state prefers the equilibrium to the collusive outcome where all states choose modest seat-vote curves.

The smallest states have the largest incentive to deviate. We saw in equation (3.16) that the smaller the state, the steeper they would like their seat-vote curve. In the limit, as a state grows in size its best response would approach the socially optimal curve. The state would play the social optimum if their own elected congress

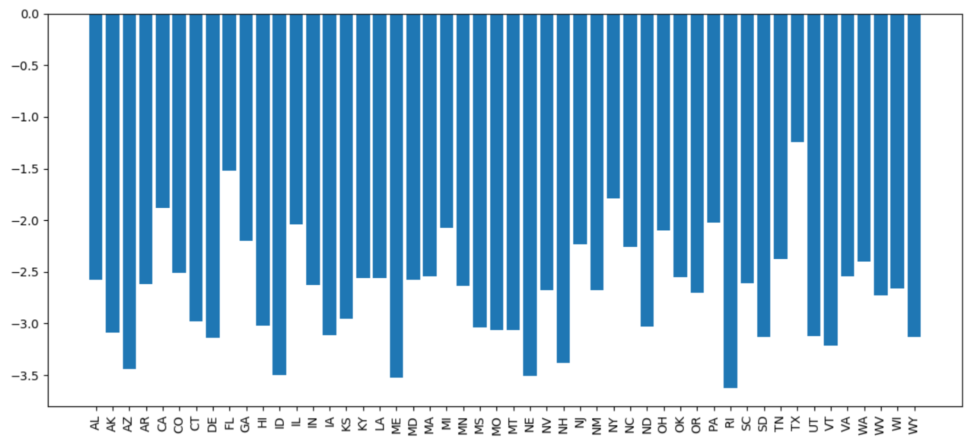

Figure 3.5: The loss each state faces in equilibrium from their expected utility in the social optimum.

members could move policy one for one. States only play the more extreme curves because the comgress members they elect are only a small part of making the policy. This becomes a larger and larger factor as a state shrinks in size.

3.6 Empirical Analysis

In this section, I validate my model by documenting a few empirical facts about seat-vote curves across states.

Estimation

First I do a simple estimation of the seat-vote curves in each state. I use data from the Cooperative Congressional Election Study. From this I aggregate across each of the 435 congressional districts what fraction report as Democrat, Republican, or Independent/Not Sure. Suppose that the self reported Democrats will vote for the Deomract candidate, the reported Republicans will vote for the Republicans candidate, and that the Independents and Not Sure respondents could vote either way.

We want to estimate the share of seats Democrats would win across all possible vote shares. To do this, I randomly draw the fraction of Independents to vote Democrat in each district 10,000 times. In each draw, we can compute how many districts in the state Democrats won and what fraction of the overall vote in the state Democrats won. Then I take the average fraction of state's seats earned for each level of the state vote share to be the seat-vote curve.

From election results in each district, we can estimate the distribution of Independents that vote Democrat and the correlation of this draw across districts within a state. In constructing the following motivational facts, I drew the fraction of Independents to vote Democrat from a Uniform distribution over $[0, 1]$ with a correlation of .7 between districts of a state. This appears to match election results fairly well.

Stylized Facts

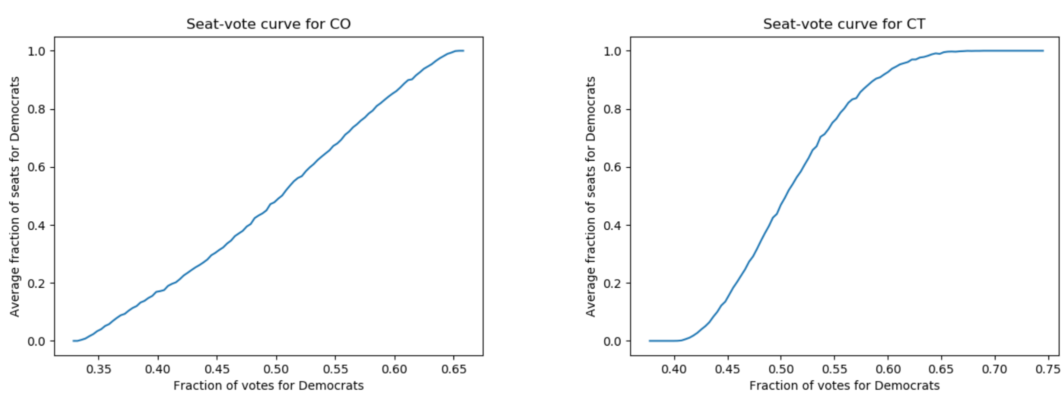

Figure 3.6: The estimated seat-vote curves for Colorado and Connecticut. Connecticut's curve is much steeper in the middle than Colorado's. Connecticut has the more responsive curve.

Figure 3.6 shows the estimated seat-vote curves for two states. Call the responsiveness of a seat-vote curve the average slope of the curve on the interval $[.45, .55]$. Over the middle ranges of votes shares, this is how much the congressional seats respond to changes in the vote on average.

The intuitive proportional seat-vote curve would be a linear function with a slope of one. We can see in the graph that the seat-vote curve in Colorado is nearly linear, but with a slope closer to three for intermediate values of the vote share. Connecticut's seat-vote curve has more of a slanted "S" shape. The is a very high slope in the intermediate values of the vote share. Between vote shares of 45 percent and 55 percent the slope averages about 6.5 or a little more than twice as steep as Colorado.

Figure 3.7 shows a scatter plot of the estimated slopes on the 45 percent to 55 percent interval for each state graphed against the number of representatives allocated to the state. Remember that the number of representatives a state has is approximately proportional to the population size of the state. The first observation is that all the slopes are well above one. The flattest state seat-vote curve is about 2.8 while the steepest curve has 10 as its average slope. Second, notice a very clear non-linear negative relationship between the responsiveness of a state's seat-vote curve and its size.

A simple regression confirms what we see in the scatter plot. *Representatives* is the number of representatives a state has. *Reps^2* is the square of *Representatives* to pick up the non-linear relationship. *Slope* is the slope of the estimated seat-vote curve. *Dem_control* is the fraction of the state that are self reported Democrats minus the fraction that are self reported Republicans. *Party_control* is the absolute value

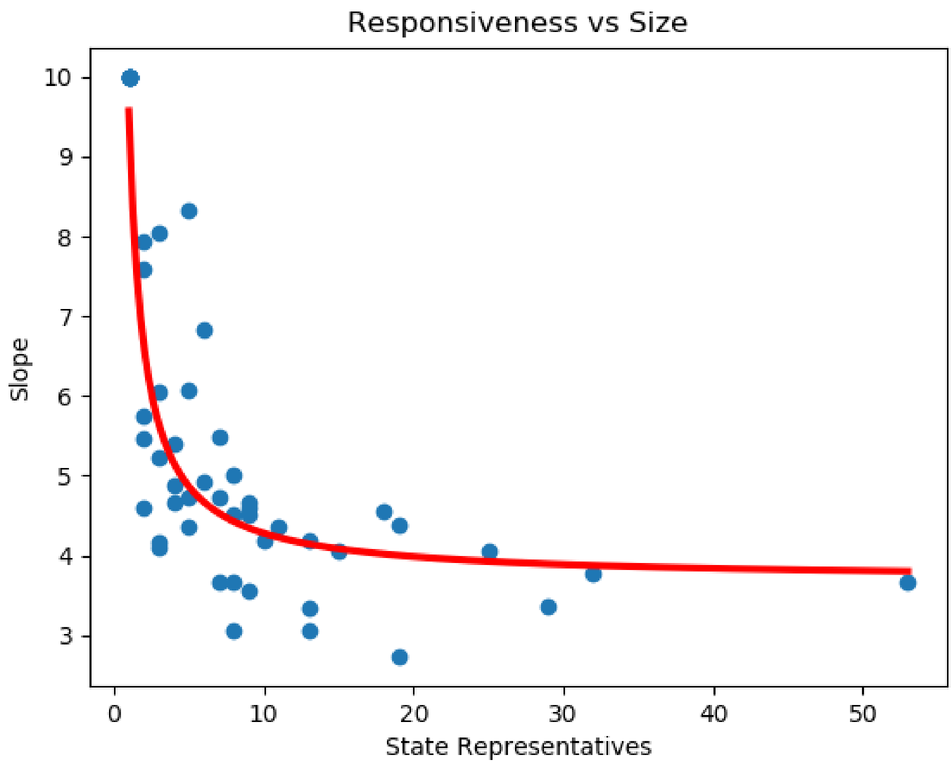

Figure 3.7: A scatter plot of seat-vote curve responsiveness verses the number of representatives allocated to a state (size).

of *Dem_control*. *n_inverse* is one divided by the number of representatives alloted to the state.

There is always a large significant relationship between *Slope* and the number of representatives. Smaller states have much steeper seat-vote curves. This relationship persist when controlling for which party is in control or how much control a party has in a state.

We can also see that the non-linearity of the relationship is picked up well by an inverse proportionality to the size of the state.

$$\frac{\partial Seats}{\partial votes} \approx a\frac{1}{n_i} + \epsilon_i \tag{3.40}$$

VARIABLES	(1) Model1	(2) Model2	(3) Model3	(4) Model3
Representatives	-0.375***	-0.335***		
	(0.0613)	(0.0649)		
Reps^2	0.00613***	0.00560***		
	(0.00136)	(0.00137)		
Dem_control		-3.805	-3.201**	-3.573**
		(2.319)	(1.316)	(1.378)
n_inverse			5.912***	5.817***
			(0.459)	(0.471)
Party_control				2.017
				(2.172)
Constant	7.777***	7.648***	3.849***	3.702***
	(0.406)	(0.407)	(0.209)	(0.262)
Observations	50	50	50	50
R-squared	0.492	0.520	0.826	0.829

Standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Table 3.1: Regression output showing the relationship between the responsiveness of a state's seat-vote curve and its size.

where n_i is the number of representatives (size) of state i.

Notice that this is the same relationship predicted by the equilibrium model.

3.7 Extension

Winner-Take-All

It may be that the optimal seat-vote curve in the model is difficult to implement in the world, requires a lot of information, or that the model assumptions don't line up exactly with your view of the world. There is a very simple seat-vote curve that can always be implemented, relies on no model assumptions, requires no information, and is very close to the optimal value. It is a winner-take-all election.

Take a seat-vote curve equal to

$$S_i(v_i) = \begin{cases} 0 & \text{if } v_i < .5 \\ 1 & \text{if } v_i \geq .5 \end{cases} \tag{3.41}$$

That is, if the Democrats get more than 50 percent of the vote in the state, all of the state's representatives will be Democrats. Otherwise, all the representatives will be Republicans.

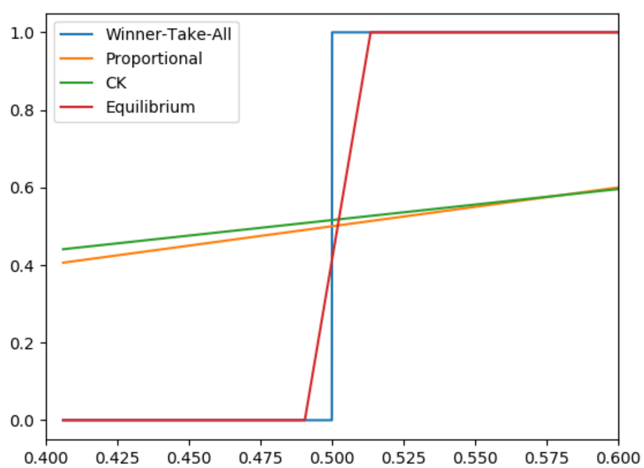

Figure 3.8: This is a graph of several seat-vote curves for the state of Minnesota. The "proportional" curve is a linear seat-vote curve with a slope of one. The "CK" curve is the socially optimal seat-vote curve. "Equilibrium" is the curve Minnesota would choose as their best response in equilibrium play. "Winner-take-all" is simply an indicator function for vote shares above 50 percent.

This doesn't require any knowledge about the voting population in the state. It is also extremely simple and easy to understand. It could be done without creating any districts. If districts are desired, all the state needs to do is randomly assign each citizen a district number regardless of geography. Each district in the United States is composed of about 700,000 people. By the law of large numbers, each district would then have pretty close to the same fraction voting Democrat in each election. This would implement the winner-take-all seat-vote curve.

While winner-take-all is not the best response, it is very close to the best response. Every state's seat-vote curve is flat at zero, then increases rapidly to one, then is flat at one. For the true best response the increase isn't infinitely steep like a winner-take-all, but the average slope is about 50. The graph shows these two curves for Minnesota.

In fact, nearly all the gains from deviating to the best response from the previous section can be obtained from deviating to a winner-take-all function.

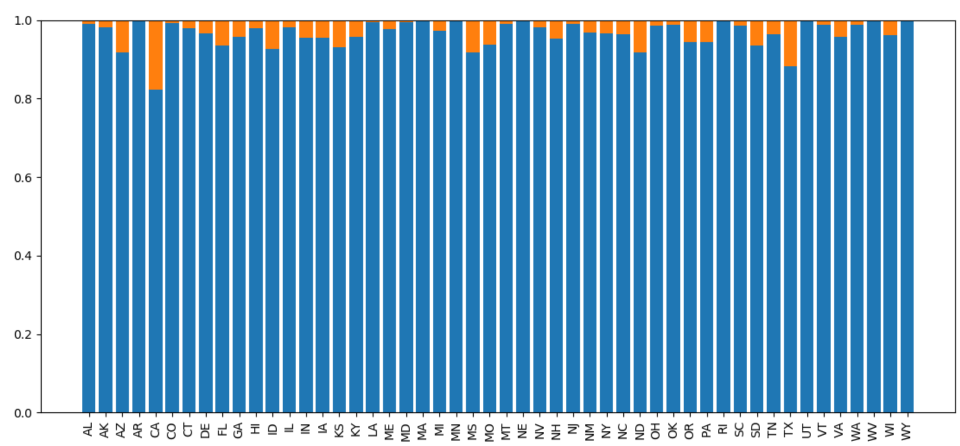

Figure 3.9: The gain from deviating from the social optimal to winner-take-all as a percentage of the gain from deviating to the best response.

In fact the strategies are almost always the same. In the best response, the fraction of Democrats elected is either 0 or 1 more than 80 percent of the time in each state already.

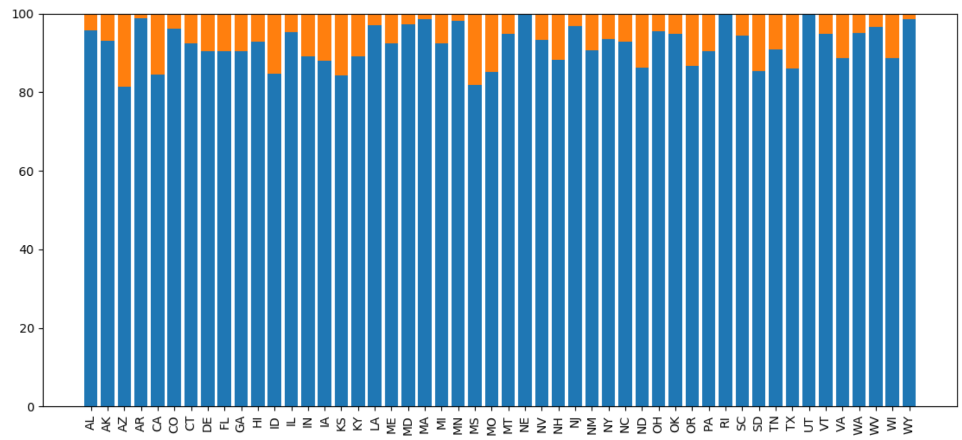

Figure 3.10: The fraction of the time the best response and a simple winner-take-all choose the same number of Democrats.

Median Member of Congress

One assumption of the model up to this point has been that the policy chosen is equal to the mean preference of all the elected congress members. One might not believe that all the members of congress are so readily willing to compromise. Perhaps, the majority party can disproportionately pull policy in their favor. In this section, I consider the opposite extreme from the rest of the paper. Suppose that the policy chosen is equal to that of the median member of congress. This would be if the majority party has complete control over choosing policy and doesn't need to compromise to please the minority. The optimal seat-vote curves will be even simpler than before but not that different intuitively.

Now, there are only two real outcomes of the game. Either the Democrats win the majority or the Republicans win the majority. State i first needs to determine which of those it prefers. State i will prefer Democrats win the majority as long as

$$\pi_{Ri} + \pi_{Ii} \int_{m_i - \tau_i}^{m_i + \tau_i} x \frac{dx}{2\tau_i} \geq \pi_{Di} + \pi_{Ii} \int_{m_i - \tau_i}^{m_i + \tau_i} (1-x)^2 \frac{dx}{2\tau_i}. \tag{3.42}$$

Computing the integral and simplifying yields a simple equation. The state will prefer a Democrat majority as long as

$$v_i \leq \frac{1}{2} + (\pi_{Ri} - \pi_{Di}) \left(\frac{1}{4\tau_i} - \frac{1}{2} \right). \tag{3.43}$$

The unique best response is for the state to give all its representatives to the Democratic Party when the vote share is above that cutoff and give all representatives to the Republican Party when it is below the cutoff. This is the case regardless of what any other state is doing. The optimal seat-vote curve for every state is a winner-take-all election. The difference from a standard winner-take-all though is that the cutoff for the Democrats to win my not be at exactly 50 percent. This is because not every vote is created equal. A firm member of the Democrat party counts for more than an independent voter that is leaning slightly Republican. Thus, the cutoff will be at a point that favors the party that has a larger political base in the state.

In this game, the socially optimal policy is not unique. Any congress with a Democratic majority is the same regardless of how strong that majority after all. The socially optimal national seat-vote curve will be any curve that is above one half if inequality 3.43 is satisfied and is below one half if inequality 3.43 is not satisfied. This could be a winner take all election with that cutoff point. This could also be the socially optimal seat-vote curve from the previous section when policy was decided by the mean congress member. This can still be implemented by every state playing the moderate "CK" seat-vote curve as before. However, every state doing a winner take all election does not implement a socially optimal national seat-vote curve.

Note that the equilibrium strategy for each state no longer depends on the size of the state. The winner-take-all election would look the same for Wyoming and Texas or for Rhode Island and California. Even if a state had complete control over the policy chosen, they would still find it optimal to choose this same seat-vote curve. That is to say, this curve is a dominant strategy for all states regardless of size and regardless of what every other state might be doing.

In truth, policy is not controlled entirely by the median congress member or by the mean congress member. It is likely something in between those two. In this model I found the optimal seat-vote curve for each state to be similar for each of the two extremes, and thus expect it to be similar for any imtermediate policy selection method.

3.8 Conclusion

Reltated Literature

The closest paper to mine is Coate and Knight (2007). In their paper, they seek the socially optimal seat-vote curve for a state. This is done for elections to the state congress instead of the national congress. In the state congress, the seat-vote curve determines all members of congress and thus has complete control over the policy. My model of a state is the same as theirs but I include it as one of 50 states that need to choose policy together. In mine the motive arises for a state to have a highly responsive curve to balance with the expected representatives from the other states.

Another paper that compliments mine well is Carson and Crespin (2004). Across states different groups are put in charge of drawing the congressional districts. They show that in states where a bipartison committee or a court is in charge of redistricting elections are more competitive than when state legislature does the redistricting. My model is closer to the motivation of a bipartison committee or court. Since having more competitive elections translates to a steeper seat-vote curve, this evidence is supportive of my model.

Bergmann et al. (2015) looks at a monopolist problem. They consider ways to divide up the population for price discrimination. Mathematically, it is very similar to how I divide up the population into districts.

Another closely related paper that deals with gerrymandering is Friedman and Holden (2008). In their paper, the districts are drawn to maximize the expected number of seats a given party wins. They find that you should sometimes “pack”, but never “crack”. In their paper you only see a signal of how each voter will vote. Different strength of signals lead to a whole spectrum of voter types.

Conclusion

Gerrymandering is a topic fiercely debated every ten years when congressional districts are redrawn. It is important because it can have a significant impact on who gets elected and ultimately on what policies are put into place. A fundamental question we should ask about gerrymandering is what seat-vote curves would we expect to see in equilibrium and what are the socially optimal seat-vote curves.

In this paper, I present a model of states choosing seat-vote curves to compete with each other over the policy that is passed. I first solve for the optimal districting strategy of one trying to maximize the expected number of seats won by a given party. I showed that the party drawing the districts can win a significantly higher fraction of seats than their share of supporters in the population. The optimal districting strategy involved both "packing" and "cracking".

I then consider a district designer that cares only about maximizing the welfare of their state's citizens. I find that in equilibrium, every state chooses an extreme seat-vote curve to disproportionately affect policy in their favor. This motive is especially strong in smaller states. In equilibrium, the slope of each state's seat-vote curve is proportional to one over the number of representatives they have.

I show that there is a deadweight loss to society in equilibrium. If every state were to commit to playing a more modest seat-vote curve, a Pareto improvement could be had. However, each state has an incentive to deviate, making this a Prisoner's Dilemma.

I also find that the optimal seat-vote curve can be approximately implimented with a very simple rule. This is a winner-take-all election for the state's representatives. If policy is chosen entirely by the majority party in congress, then this winner-take-all election becomes exactly the unique optimal rule for each state.

I then estimate the seat-vote curves for each of the 50 states. I find that seat vote curves are highly responsive. I also find that small states choose steeper seat-vote curves than larger states. The slope of the seat-vote curve in a state is approximately proportional to one over the number of representatives that state has in congress. This is the same relationship found in the equilibrium model. This relationship holds after controlling for the politcal party in power.

Chapter 4

Ratings and Reputation

Updated version at: Link

Abstract

I present a model of universities to explain patterns of grade inflation that have been observed. I model the university as an information designer hired to evaluate students for employment purposes. I find both pooling and separating equilibria. I show how the university's desire to build or maintain their own reputation interacts with their desire for the success of their students in forming the optimal grading scheme. In every equilibrium grade inflation is concentrated at the high quality universities as observed in the data. I also show how the observed patterns of grade inflation are not likely to be driven by heterogeneous students selecting into quality universities.

4.1 Introduction

Grade inflation is now a well documented phenomenon as well as a topic being discussed at major universities around the world. At today's colleges an A just doesn't mean what it used to. For decades the average GPA at universities as been steadily rising. What is more, grade inflation is not happening uniformly for all universities. Grade inflation has been more significant in the most prestigious universities. Today the average GPA at Harvard, Stanford, and Yale is over 3.6. While Middle-of-Nowhere Tech is still handing out C's regularly, the most common grade at Harvard has become an A. If the most common grade handed out is an A, there isn't much ability for the school to distinguish its best students. If the role of grades is to evaluate students and separate them for potential employers, it would seem a more uniform grading scheme would be better served.

Similar patterns have been seen for credit rating agencies of corporate bonds. The ratings of the most established agencies have grown less informative over time, while the less well recognized agencies haven't changed as much.

In this paper, I will present a model of universities as evaluators of potential employees. In this model I will show how grade inflation will be the optimal policy of the university seeking to maximize the profits they receive from tuition. The model will also generate grade inflation concentrated primarily in the best universities, while the worst universities will have grades more constant or even deflating.

The basic model is as follows. An employer would like to hire productive workers but not hire unproductive workers. Since a potential worker doesn't have any way to credibly signal their productivity on their own, they must attend a university to be evaluated. The university using assignments, projects, exams, etc. evaluates the worker and provides a grade that employers can observe. The university can commit to a grading scheme and does all this in exchange for a tuition fee from the worker. I characterize the optimal grading scheme and fee for the university.

I then extend the model to a dynamic setting where there are new workers enrolling every period. What's more, the university is not able to perfectly observe the worker's productivity. The university must now try to get workers hired so they can charge a large fee, while also build a reputation for being able to predict worker productivity well. I find pooling equilibria where all universities use the same grading scheme (conditional on their reputation), but based on the universities success rate at predicting worker success the good universities still gain reputation and the poor universities lose reputation. I also find separating equilibria. In these the good universities must first have a very strict grading scheme until they adequatly prove that they can predict worker productivity. Then they can employ a very lenient grading scheme and reap the higher profits from their good reputation. The bad universities employ a moderate grading scheme that is constant over time.

Finally, I show a version with ex ante heterogeneous potential employees. I show that the observed phenomenon in the data are not likely to be driven by the good schools imply getting higher and higher quality students.

4.2 Grading Schemes to Signal Worker Quality

There are three players in the game. There's the worker, the employer and the university. The worker is one of two types: either a porductive worker (high type), or unproductive worker (low type). The probablility that the worker is high type will be denoted by μ. The employer would like to hire a productive worker, and not hire an unproductive worker. However, the employer cannot observe the worker's type, and the worker doesn't have any creditable way to reveal their type. So, the worker pays the university to evaluate them and publish a rating. The worker pays a fee to the university, the university observes a signal of the worker's quality, and sends a message to the employer.

Setup

The timing of the game is as follows. First nature selects the worker's type, $\omega \in \{H, L\}$, but this is not observed by any player. Next, the university chooses a grading scheme and a fee. The grading scheme is a mapping from worker quality signals they might observe into messages they might send to the employer. The university can use any potentially mixed strategy on an arbitrary message space. The worker chooses whether or not to enroll in the university If the worker chooses to enroll in the university, then thier quality signal is realized and message is sent according to the grading scheme specified. After recieving the message, the employer chooses whether or not to hire the worker. The employer gets a payoff of one if they hire a productive worker and a payoff of $-\xi$ if they hire an unproductive worker. Thus the employer will choose to hire the worker if they believe the probability the worker is a high type is at least $\bar{\mu} = \frac{\xi}{1+\xi}$. Assume ξ is large enough that $\bar{\mu} > \mu$. So without any additional information, the employer would choose not to hire the worker. The worker recieves a wage of one if they are hired and a wage of zero if they are not hired.

If the university's signal was able to perfectly reveal worker type then this would be the most standard of Bayesain Persuasion games. Denote the precision of the university's signal, referred to as the university's quality, as $\theta \in [0, 1]$. Upon observing the universiy's signal, the posterior belief about the worker's type would be equal to the following.

$$\mu' = \begin{cases} \theta + (1-\theta)\mu & \text{w/ prob. } \mu \\ (1-\theta)\mu & \text{w/ prob. } 1-\mu \end{cases} \tag{4.1}$$

If $\theta = 1$, the signal is perfectly informative of the worker's type and the posterior will equal 0 or 1. If $\theta = 0$, the signal contains no information and the posterior will be equal to the prior, $\mu' = \mu$. Since the worker gets a payoff of zero if they don't attend the university, the university can set their fee equal to the expected payoff they can get for a student that enrolls. μ' isn't the signal the university receives. It is the posteriors that are induced after observing the signal. Since the distribution of posteriors averages out to the prior (is Bayes plausible), we know that there exists some signal that would induce this distribution of posteriors (see Kamenica and Gentzkow (2011)). It is more convienient to work directly with the posteriors than the signal itself. So I will leave the signal unspecified.

The object of interest is the optimal grading scheme for the university. I am going to charactarize this in the context of perfect Bayesian equilibrium.

Imperfect Signals

First take $\theta = 1$. This problem is just like the example in the Kamenica and Gentzkow (2011). The university's payoff is equal to the fee they can charge the worker. Therefore, the university would like to choose a grading scheme to maximize the enrolled worker's expected payoff, then charge a fee equal to that expected payoff. There is a simple grading scheme to achieve this. If the university sees that the worker is good, they will pass them. If the university sees that the worker is bad, they will randomize by sometimes passing and sometimes failing them. The university will pass as many workers as possible will still making sure that everyone they pass gets the job. This means the university increases the fraction of bad workers they pass until the employers posterior after observing a pass is exactly equal to $\bar{\mu}$. This gives the usual concavification payoff. The fee they can charge is then equal to $v(\mu) = \frac{\mu}{\bar{\mu}}$.

With $\theta < 1$ the optimal grading scheme is very similar as long as

$$\theta \geq \theta_{min} = \frac{\bar{\mu} - \mu}{1 - \mu}. \tag{4.2}$$

A grading scheme can induce any distribution of posteriors that second order stochastically dominates the distribution of posteriors obtained by fully revealing their information (see Kolotilin et al. (2017)). When the university can learn perfectly what the worker's type is, the full information distribution of posteriors puts all weight on zero or one. This distribution is second order stochastically dominated by any distribution of beliefs that has the same mean. This is why Kamenica and Gentzkow (2011) has only the condition that the posteriors must have a mean that matches the prior. Here, the full information distribution of posteriors is $\theta + (1 - \theta)\mu$ or $(1 - \theta)\mu$ with frequencies μ and $1 - \mu$ respectively. The university choosing a grading scheme is equivalent to choosing a distribution of posteriors that second order stochastically dominates that full information distribution. Since the full information distribution has all its mass on two points, the second order stochastic dominance condition is simple. A distribution with the same mean will second order stochastic dominate the full information distribution if and only if its support is contained in the interval $[(1 - \theta)\mu, \theta + (1 - \theta)\mu]$. That is, the imperfect information on the part of the university simply puts bounds on the distribution of posteriors they can induce. The only addition requirement beyond beliefs being a martingale, is that the university can not convince the employer of the state beyond their own level of conviction.

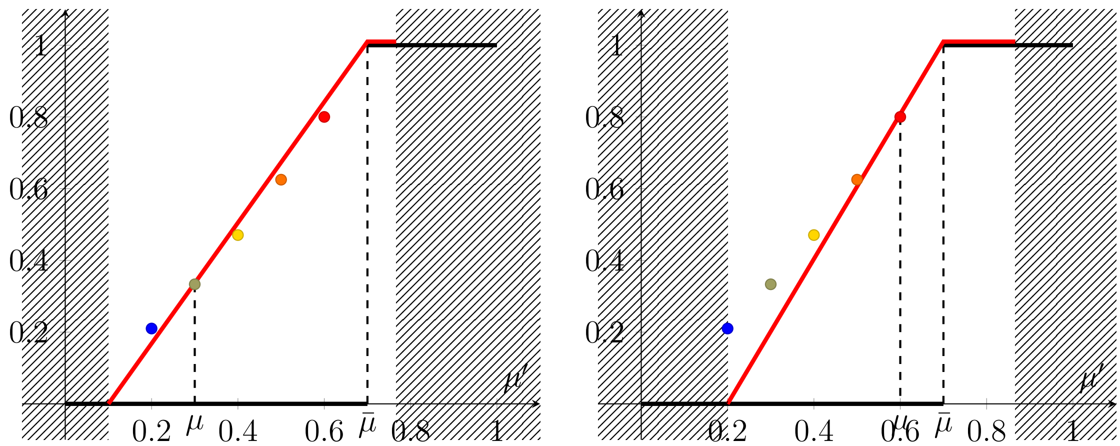

(a) $\mu = .3$ (b) $\mu = .6$

Figure 4.1: Persuasion with imperfect information ($\theta = \frac{2}{3}$).

The optimal grading scheme will have the same general form as before. The university will pass the worker whenever they get a good signal, and when the university gets a bad signal that will randomize between passing and failing the worker. They will pass just enough bad workers so that the posterior of the employer after observing a pass is equal to $\bar{\mu}$. Call the frequency with which they can pass workers after getting a bad signal $\pi(\theta)$. First see that the posterior belief of the employer after observing a pass is

$$\frac{\mu}{\mu + (1-\mu)\pi}(\theta + (1-\theta)\mu) + \left(1 - \frac{\mu}{\mu + (1-\mu)\pi}\right)(1-\theta)\mu. \tag{4.3}$$

The first term of this equation is the updated probability that the university got a good signal times the probability the worker is good conditional on the university getting a good signal. The second term is the updated probability that the university got a bad signal times the probability the worker is good conditional on the university getting a bad signal. To get the optimal frequency with which the university passes bad workers, you set this equal to $\bar{\mu}$ and solve for π.

$$\pi^*(\mu) = \frac{\mu}{1-\mu}\left(\frac{\theta}{\bar{\mu}-(1-\theta)\mu} - 1\right) \tag{4.4}$$

When $\theta = 1$, this is the usual equation.

$$\pi^*(\mu) = \frac{\mu}{1-\mu}\frac{1-\bar{\mu}}{\bar{\mu}} \tag{4.5}$$

When $\theta = \frac{\bar{\mu}-\mu}{1-\mu}$, the university is not able to pass any bad workers and still have their students hired. Then $\pi^*(\mu) = 0$.

The better the university is, the more often they can pass bad workers. Regardless of θ the posterior after a pass is equal to $\bar{\mu}$. The difference is that the posterior after a fail is increasing in θ. The loss in payoff to the university comes from how often a good worker goes to the university and gets misclassified as bad. Each time that type of misclassification happens, that's a worker everyone agrees should get hired but probably doesn't get hired. Additionally, if all those workers were added into the pool of workers that pass then more bad workers could also be added in to keep the poster at $\bar{\mu}$. In the concavification picture, we can see that the loss in payoff comes from the left endpoint moving to the right.

If $\theta < \frac{\bar{\mu}-\mu}{1-\mu}$, then there is no grading scheme the university is able to employ that will get any of their graduates hired. The university is just viewed by the employer as too low of quality and the employer doesn't trust their judgement enough to hire their graduates. In this case, the university and the worker will both get a payoff of zero regardless of the grading scheme chosen.

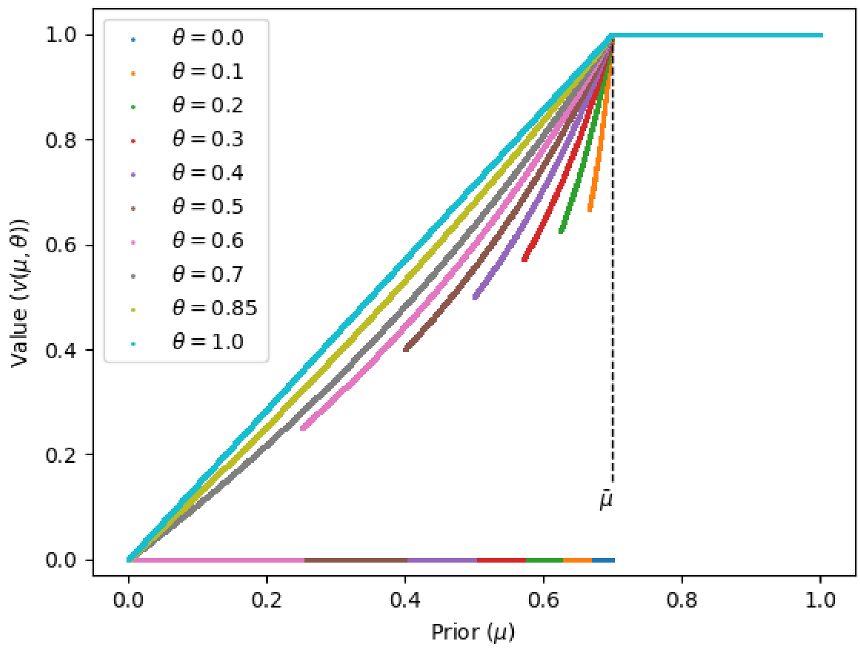

Figure 4.2: The value to the university as a function of their quality.

Uncertain School Quality

Suppose that the quality level of the university, θ, is not publicly observed. Given a grading scheme π, the employer's posterior belief depends on the university's quality. Notice from equation (4.3) that the posterior is a linear function of θ. So, if x is the distribution of θ in the mind of the employer, the posterior belief

$$\frac{\mu}{\mu + (-\mu)\pi}(\mathbb{E}_x[\theta] + (1-\theta)\mu) + \left(1 - \frac{\mu}{\mu + (1-\mu)\pi}\right)(1 - \mathbb{E}_x[\theta])\mu. \tag{4.6}$$

will only depend on the expected value of θ. Also, note that all university types receive a good or bad signal with the same frequency. This means that no updating of x is needed upon observing the university's message (given the grading scheme π).

Now the payoff of all players, including the university, depend only on the expected value of θ and not the actual value. This means the equilibrium will be exactly the same as in the case where the school's quality was known, with that known level being equal to $\mathbb{E}_x[\theta]$. It isn't even necessary for the university to know their own quality level, since their optimal strategy and payoff only depend on the the average quality level.

The average quality level is a public good for all universities now. If at the beginning of the game universities chose a quality level facing some cost function, there would be an underprovision of quality compared to the socially optimal level. This could justify some regulations we see. A university must be accredited before they can give out degrees and credit rating agency must have their methodology filed with the SEC. This underprovision of quality would be expected in any advice giving industry, where the firm cannot know the state perfectly. We see similar legislation arise elsewhere. Financial advisors have a CFA, doctors have a medical license, lawyers pass a bar exam, etc. As we'll see in the next section, the only way such legislation isn't necessary is if the advice giver can establish an individual reputation. However, in many instances this may not be possible. It is difficult for the customer to find the success rate of an individual doctor and so forth.

4.3 School Reputation

In this section, I will show you how the university can establish an individual reputation to circumvent the public goods problem. We will need to add some dynamics to the model to make this possible. In a one period model, there is nothing a high quality university can do. There is no time for the employers to learn of the university's quality by their success rate. Also, there cannot be a seperating equilibrium where high and low type universities use different grading schemes to reveal themselves. This is because the type, θ, is not directly payoff relevent to the university. When a bad university plays the same action as a good university, they get the same payoff. There is no way to distinguish. In a model like Spence (1973) a seperating equilibrium arises because the different types of students have different costs for attending school. Here, the grading scheme is costless for universities of all types.

If we make this a dynamic game where new workers come to the university every period, both claims above are reversed. Over time, the employer can learn the university's quality by observing their success rate at predicting which workers are good. This learning is hindered somewhat by the fact that universities are not perfectly revealing everything they know. Also, there can be seperating equilibria where high and low quality universities choose different grading schemes. I will explain each of these in turn.

Pooling Equilibria

I will set up the dynamic game here. It is the same as the one period game except that new workers will arrive every period. I will first characterize the pooling equilibria where high and low type universities use the same grading scheme. Even though the strategies played are the same, employers still learn the university's type through observing success rates over time. Thus, the universities are still able to build a good reputation. Then, I will characterize separating equilibria.

Setup

At the beginning of the game the university's type, θ, is drawn and remains fixed throughout the game. This type is unobserved by the employer. The employer's prior over this type is x_0.

In each period, the university chooses a grading scheme and a fee. Then a worker with productivity $\omega \in \{0, 1\}$ (μ is the probability that $\omega = 1$) chooses whether or not to enroll in the university. If the worker is enrolled, the university sees a signal of the worker's quality (the same way as before) and sends a message to the employer according to their grading scheme. The employer sees the message and chooses whether or not to hire the worker. Then, the worker's productivity (ω) is observed and all players receive their period payoffs. The employer gets a payoff of 1 for hiring a productive worker, $-\xi$ for hiring an unproductive worker, and 0 for not hiring a worker. The university's payoff is equal to the fee they choose if the worker enrolls. The worker gets a payoff of 1 if they are hired and 0 if they are not hired minus the fee they pay to the university if they enroll.

The employer then updates their beliefs about the university's quality. In the next period, another worker is independently drawn and the same game is played again. There is a discount factor of δ between periods.

It isn't important in this section whether the worker or the university know θ. All university types receive a good signal with the same frequency. This means that given the grading scheme, the worker's probability of passing and getting the job doesn't depend on θ. So the worker doesn't care about going to a good university, they only care about going to a university with a good reputation. Also, in this section we are looking at pooling equilibria. You can think of them as pooling equilibria of the game when the university does know their own type, or they would be the only equilibria of the game when the university doesn't know their own type.

Equilibria

The first obvious equilibrium is simply the repeated Nash. The university could play the grading scheme corresponding to

$$\pi^*(\mu, x) = \frac{\mu}{1-\mu}\left(\frac{\mathbb{E}_x[\theta]}{\bar{\mu} - (1 - \mathbb{E}_x[\theta])\mu} - 1\right). \tag{4.7}$$

Even though the universities are playing the same strategies and it is just the repeated static Nash, there is still learning and reputation. Consider the probabilities of each possible outcome of a period respectively: university says the worker is good and they actually are good, university says good but the worker is bad, university says bad and the worker is bad, and university says bad but their actually good.

$$p^{gg}(\pi, \theta) = \mu\left(\mu + (1-\mu)\pi + (1-\mu)(1-\pi)\theta\right) \tag{4.8}$$

$$p^{gb}(\pi, \theta) = (1-\mu)\left(\mu + (1-\mu)\pi - \mu(1-\pi)\theta\right) \tag{4.9}$$

$$p^{bb}(\pi, \theta) = (1-\mu)(1-\pi)(1-(1-\theta)\mu) \tag{4.10}$$

$$p^{bg}(\pi, \theta) = (1-\mu)(1-\pi)(1-\theta)\mu \tag{4.11}$$

We can take the derivatives of these probabilities with respect to θ.

$$\frac{\partial p^{gg}}{\partial \theta} = \frac{\partial p^{bb}}{\partial \theta} = -\frac{\partial p^{gb}}{\partial \theta} = -\frac{\partial p^{bg}}{\partial \theta} = \mu(1-\mu)(1-\pi) \tag{4.12}$$

There are a few things we can notice here. First, the probability of a good message and the probability of a bad message are constants in θ. Thus, the message itself doesn't reveal anything about the univerisity's quality. The learning only happends after the worker's type is revealed. Next see that p^{gg} and p^{bb} are increasing functions of θ while p^{gb} and p^{bg} are decreasing functions of θ. Regardless of the grading scheme, $\pi \in [0, 1)$, the good universities are more likely to correctly identify the productivity of the worker. However as π increases, the dependence on θ goes to zero. The more leanient the grades are, the harder it is to learn not just of the worker quality but also of the university quality.

From this, we can derive the posterior beliefs about university quality after each period outcome. Here x_i denotes the probability that the university is type θ_i. After the university correctly identifies a good worker:

$$x_i^{gg} = x_i \frac{\mu + (1-\mu)\pi + (1-\mu)(1-\pi)\theta_i}{\mu(1-\mu)\pi + (1-\mu)(1-\pi)\mathbb{E}_x[\theta]}. \tag{4.13}$$

After claiming a worker is good they are actually bad:

$$x_i^{gb} = x_i \frac{\mu + (1-\mu)\pi - \mu(1-\pi)\theta_i}{\mu + (1-\mu)\pi - \mu(1-\pi)\mathbb{E}_x[\theta]}. \tag{4.14}$$

After correctly reporting a bad worker:

$$x_i^{bb} = x_i \frac{1 - (1-\theta_i)\mu}{1 - (1 - \mathbb{E}_x[\theta])\mu} \quad \forall \pi^x \neq 1. \tag{4.15}$$

After incorrectly claiming a worker is bad:

$$x_i^{bg} = x_i \frac{1 - \theta_i}{1 - \mathbb{E}_x[\theta]} \quad \forall \pi^x \neq 1 \text{ and } \theta \neq 1. \tag{4.16}$$

We get that the universities are building up their reputation over time. As their reputation grows, the university experiences grade inflation. The bad universities have reputation levels falling over time. At first grades fall too, but soon the university hits the lower bound of reputation where they aren't passing any bad students. For all reputation levels below that, the grading scheme is constant. GPAs are fairly constant over time at the bad universities but rising at the good universities, consistent with the data.

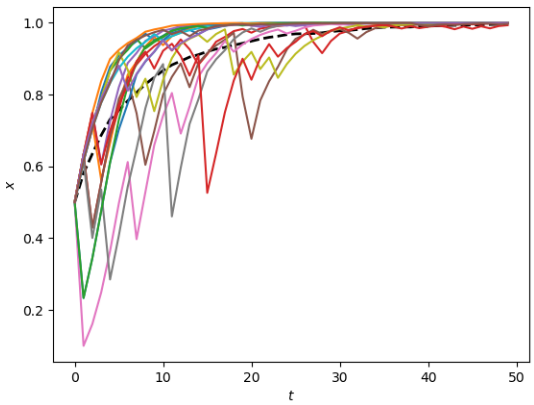

(a) Reputation of the high type over time.

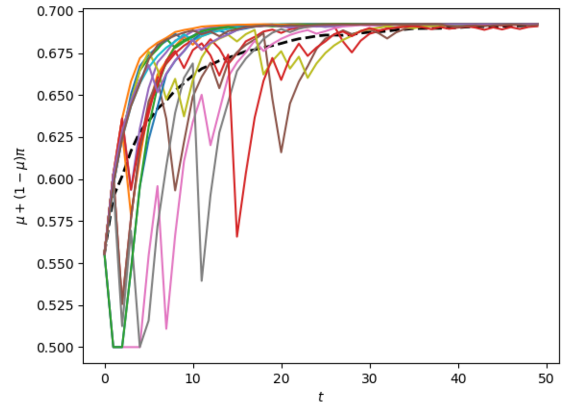

(b) Fraction of students passed over time by the high type.

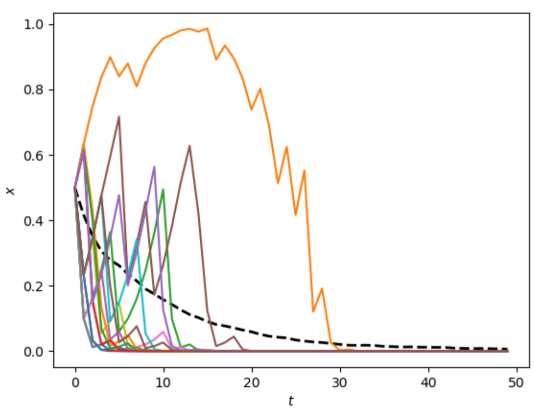

(c) Reputation of the low type over time.

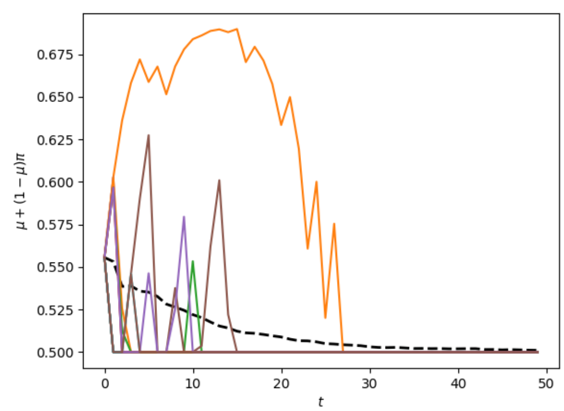

(d) Fraction of students passed over time by the low type.

Figure 4.3: Simulations of reputation and pass rate at the university.

When the university has a good reputation and employs a lenient grading scheme, learning slows down. The grades are less informative about worker productivity, so getting the grade "right" or "wrong" doesn't mean as much about the university's quality. Conversely when the reputation is low, the grading scheme is very rigorous. This means that if a worker passes or fails it is because the university thinks they are good or bad respectively. Any worker incorrectly identified is due to a mistake by the university. Mistakes are then very informative about the university quality. So, a bad university that lucks into a high reputation can maintain that reputation for a while and only slowly be revealed. On the other hand a good university that lucks into a bad reputation is able to turn things around very quickly and bring their reputation back up.

Seperating Equilibrium

Unlike the one period model, there can be separating equilibria in the multi-period model. In this section, suppose there are only two quality levels for the university, $\theta \in \{\theta_H, \theta_L\}$. Separating equilibria can only exist if θ_L is sufficiently high. If $\theta_L \leq \frac{\bar{\mu}-\mu}{1-\mu}$, then a university revealed as the low type will get zero profits forever. Assuming a university can't have negative fees, the low type university will always be better off pretending to be a high type.

Call π_L the optimal grading scheme for a university known to be low type, $\theta_L \geq \frac{\bar{\mu}-\mu}{1-\mu}$.

$$\pi_L = \frac{\mu}{1-\mu}\left(\frac{\theta_L}{\bar{\mu}-(1-\theta_L)\mu} - 1\right) \tag{4.17}$$

In the separating equilibria, the low type universities will play π_L every period. The high type universities will play a history dependent grading scheme, $\pi(h^t)$. If the university were to play any off-path action, the employer would believe them to be the low type. I will now describe the conditions on $\pi(h^t)$ required to make this constitute an equilibrium.

First, we need the low type university not to want to pretend to be the high type for any number of periods. This requires that

$$\pi_L \left(\frac{1-\delta^\tau}{1-\delta} \right) \geq \mathbb{E} \left[\sum_{t=0}^{\tau-1} \delta^t \pi(h^t) \mid \theta = \theta_L \right] \qquad \forall \tau = 1, 2, \ldots \tag{4.18}$$

It needs to be that the for any number of periods, from the start of the game, the low type would be better off playing π_L than playing the high type's strategy in expectation. Taking $\tau = 1$ this necessarily implies that the high type will play a more revealing grading scheme at first than the low type, $\pi(h^0) < \pi_L$. In any separating equilibrium, the high type will be initially playing a more stringent (revealing) grading scheme.

The other necessary condition is that the high type would be better of playing $\pi(h^t)$ than playing π_L for the rest of the game after any history. This will require the following.

$$\mathbb{E} \left[\sum_{t=\tau}^{T} \delta^{t-\tau} \pi(h^t) \mid \theta = \theta_H, h^\tau \right] \geq \pi_L \left(\frac{1-\delta^{T-\tau}}{1-\delta} \right) \qquad \forall \tau, h^\tau \tag{4.19}$$

At any point in time, τ, no matter what has happened up until that point, it must be better for the the high type to stay the course by playing $\pi(h^t)$ than switch to π_L. Since we know that the high type is playing a lower π in the first periods than the low type, for this condition to hold when $\tau = 0$ it must be that the high type is playing a higher π in the later periods. We can already see that a separating equilibrium will have grade inflation for the good universities but a constant grading scheme for the bad universities.

If we take $\tau = T$ in equation (4.18) and we take $\tau = 0$ in equation (4.19), the two equations seem to nearly make a contradiction.

$$\mathbb{E} \left[\sum_{t=0}^{T} \delta^t \pi(h^t) \mid \theta_H \right] \geq \pi_L \frac{1-\delta^T}{1-\delta} \geq \mathbb{E} \left[\sum_{t=0}^{T} \delta^t \pi(h^t) \mid \theta_L \right] \tag{4.20}$$

In fact, this can still be consistent (even if the inequalities were strict). The separating equilibria still look a little different from those in other models though. There will never be an equilibrium like, "the high type must play a low π for two periods and then they can play π_H forever," or anything similar. Anything such rule for $\pi(h^t)$ will have the same problem that the one period model had. It will always have the same payoff for both types. If the high type finds it profitable to do that strategy, the low type will also. With any rule like that, the left most term in inequality (4.20) and the right most term in inequality (4.20) will be equal. The only difference between those two terms in the inequality is in the expectation. We saw earlier that the probability of each outcome (a message and worker productivity realization together) is dependent on the university quality, even though the message distribution isn't. Thus the only way to get separation between the two terms of the inequality is to make $\pi(h^t)$ depend on both the message distributions and the outcomes of previous periods. Instead of the high type having to play a low π for two periods, they may need to play a low π until correctly identifying two workers' productivity for example. The correct identification is more likely for the high type with any given grading scheme. Thus, the low type university would need to play the low π for longer on average. This is the only way to get different "costs" that depend on the university's quality.

We also know that neither type will want to deviate to anything other than π_L or $\pi(h^t)$. At any other deviation, the employer assumes the university is low quality. We have already established that π_L is the best thing the university can play if they are known to be low quality.

Together with the grading scheme being such that workers that pass will get hired,

$$\pi(h^t) \leq \frac{\mu}{1-\mu}\left(\frac{\theta_H}{\bar{\mu}-(1-\theta_H)\mu} - 1\right), \tag{4.21}$$

inequalities (4.18) and (4.19) are sufficient conditions for $\pi(h^t)$ to be the strategy of the high type in a separating equilibrium where the low type plays π_L.

Example

In this section, I will show an example of such a separating equilibrium for given parameter values. Consider a two period model. Let the worker be either productivity level with equal probability, $\mu = \frac{1}{2}$. Also let $\xi = 3$ so the required posterior for the worker to get hired is $\bar{\mu} = \frac{3}{4}$. Then the maximal grading scheme for a university of type θ in equation (4.4) simplifies.

$$\pi^{\theta} = \frac{\theta - \frac{1}{2}}{\theta + \frac{1}{2}} \tag{4.22}$$

Suppose also that the high type university is able to perfectly predict the worker's productivity, $\theta_H = 1$.

In equilibrium the low type will play π_L both periods. The high type will play $\pi = 0$ in the first period. If the university playing $\pi = 0$ in the first period correctly identifies the worker productivity, then in the second period they will play $\pi = \pi_H$. If the university is incorrect in the first period, in the second period they will play π_L. The low type is able to employ some persuasion by passing some bad workers each period. The high type will take a lower payoff in the first period by employing no persuasion. The high type simply reveals the worker's productivity perfectly without passing any bad workers. Then, after correctly identifying the first worker's productivity and proving they are a high type, the university can use a large amount of persuasion in the second period. The low type will unwilling to copy the high type's action because there is a chance they would misidentify the period one worker's productivity and be revealed as a low type.

First, we'll look at the condition to ensure the high type doesn't deviate to the low type's strategy.

$$\mu + \delta\left(\mu + (1-\mu)\pi_H\right) \geq (1+\delta)\left(\mu + (1-\mu)\pi_L\right) \tag{4.23}$$

This will hold as long as the payoff to the low type is sufficiently low.

$$\pi_L \leq \frac{\delta}{1+\delta}\pi_H \tag{4.24}$$

Using $\theta_H = 1$ and $\delta = 1$, this becomes a condition on θ_L.

$$\frac{\theta_L - \frac{1}{2}}{\theta_L + \frac{1}{2}} \leq \frac{1}{6} \tag{4.25}$$

$$\theta_L \leq .7 \tag{4.26}$$

So this type of equilibrium is can only be supported if θ_L is low enough.

Now we need to make sure that the low type doesn't want to deviate to playing the high type's strategy. This requires that playing π_L in both periods is better than playing $\pi = 0$ in the first period and playing π_H in the second period if the first period identification was correct and playing π_L in the second period if the first period identification was incorrect. Call $p^{m\omega}$ the probability of sending message m and the worker's type being ω, with the grading scheme $\pi = 0$ and university quality $\theta = \theta_L$, $p^{m\omega} = p^{m\omega}(0, \theta_L)$. The condition is now the following.

$$\begin{aligned}(1+\delta)\left(\mu + (1-\mu)\pi_L\right) \;\; \geq \;\; \mu \; + \; & \delta\left(p^{gg} + p^{bb}\right)\left(\mu + (1-\mu)\pi_H\right) \\ + \; & \delta\left(p^{gb} + p^{bg}\right)\left(\mu + (1-\mu)\pi_L\right)\end{aligned} \tag{4.27}$$

This can be simplified to be only the additional gain from each strategy.

$$\pi_L \;\geq\; \delta\left(p^{gg} + p^{bb}\right)\left(\pi_H - \pi_L\right) \tag{4.28}$$

The benefit of playing the low type strategy is getting π_L of the bad workers hired in the first period while the high type isn't getting any. The benefit of playing the high type strategy is the extra bad workers you can get hired in the second period, $\pi_H - \pi_L$, if you're able to pass the first period test of correctly identifying a worker. Rearrange the inequality to get the restriction on π_L.

$$\pi_L \;\geq\; \frac{\delta\left(p^{gg}+p^{bb}\right)}{1+\delta\left(p^{gg}+p^{bb}\right)}\pi_H \tag{4.29}$$

Using the probability equations above (equation (4.8)), we get the following.

$$p^{gg}+p^{bb}=1\;-\;2\mu(1-\mu)(1-\theta) \tag{4.30}$$

Now we can plug this into our restriction on π_L along with $\theta_H = 1$, $\mu = \frac{1}{2}$, and $\delta = 1$.

$$\pi_L \;\geq\; \frac{1-\frac{1}{2}\left(1-\theta_L\right)}{2-\frac{1}{2}\left(1-\theta_L\right)}\frac{1}{3} \tag{4.31}$$

Putting in the equation for π_L and solving for θ_L, we get $\theta_L \geq .68$. We find that θ_L needs to be sufficiently high that the low type doesn't want to give up their first period gains for a gamble.

Together with the restriction to keep the high type from deviating, we have now found that if $\pi_L \in [.68, .7]$ this constitutes a separating equilibrium. Observe the grading inflation pattern in this equilibrium. Take $\theta_L = .7$. Exactly one half of the workers have a high productivity each period. The poor quality universities are passing $\frac{7}{12}$ of their workers every period. The bad university lets a few workers skate by when they don't deserve it. The high quality university only passes $\frac{1}{2}$ of the workers in the first period, but in the second period they pass $\frac{2}{3}$ of the workers. At first the high quality university only passes worker that truely deserve it, but after some time they start letting more bad workers skate by than the low quality university did.

4.4 Heterogeneity

Some will certainly argue that good schools give out better grades because they have better students. This would imply that grade inflation at good schools is being driven by the good schools getting better and better students. First since the employer knows a school is better and will account for that and what workers will attend in equilibrium in forming their posterior beliefs, that really isn't a reason for higher grades. However, let's now extend the model a bit to add in ex ante heterogeneous workers and see some other reasons this won't lead to the grading patterns observed in the data.

Say there are a mass of workers and the employer (and university) have different priors about different workers' productivity even before they attend the university. The workers have different grades from high school, SAT scores, extra curricular activities, recommendations, etc. There is a unit mass of workers and the priors over worker productivity is uniformly distributed on the unit interval, $\mu \nabla U[0, 1]$. There will be only one period. The university will choose only one grading scheme, π, and one fee, f. Then each worker with their individual likelihood of being productive (prior, μ) will choose whether or not to enroll in the university. For every worker enrolled at the university, the university will receive a signal of productivity and a message will be sent to the employer according to the grading scheme. Call the set $\mathcal{E} \subset [0, 1]$ the set of workers that choose to enroll in the university. The university's payoff is then

$$v = \int_0^1 f \mathbf{1}_{\mathcal{E}} \, dx \tag{4.32}$$

or f times the measure of the set $\mathcal{E}$.

The university's grading scheme will be of the same form as earlier. The university will pass all the workers it receives a good signal about, and also pass π fraction of the workers it receives a bad signal about. Note that given the grading scheme and the fee, if a worker at prior μ finds it profitable to enroll then a worker at prior $\mu' > \mu$ will also find it profitable to enroll (except of course workers with prior $\mu' > \bar{\mu}$ who will choose not to enroll regardless of the fee.) This means that the set of students who choose to enroll will always be a simple interval, $\mathcal{E} = [\hat{\mu}, \bar{\mu}]$ where $\hat{\mu}$ is such that

$$\hat{\mu} + (1 - \hat{\mu})\pi = f. \tag{4.33}$$

We can think of the university's problem as one of choosing $\hat{\mu}$ then back out π and f from that. π will be the highest π that can still get the $\hat{\mu}$ workers the job after a good signal and f will equal the probability of the $\hat{\mu}$ worker passing.

The value to the university from worker $\hat{\mu}$ is then

$$v(\hat{\mu}) = \hat{\mu} + (1 - \hat{\mu})\pi^*(\hat{\mu}) \tag{4.34}$$

$$= \frac{\hat{\mu}\mathbb{E}[\theta]}{\bar{\mu} - (1 - \mathbb{E}[\theta])\hat{\mu}}. \tag{4.35}$$

The overall value to the university is found by then multiplying by the measure of workers that enroll.

$$v = \frac{\hat{\mu}\mathbb{E}[\theta]}{\bar{\mu} - (1 - \mathbb{E}[\theta])\hat{\mu}}(\bar{\mu} - \hat{\mu}) \tag{4.36}$$

The university will choose $\hat{\mu}$ to maximize v. Differentiate v with respect to $\hat{\mu}$.

$$\frac{\partial v}{\partial \hat{\mu}} = \frac{\bar{\mu}\mathbb{E}[\theta] - 2\mathbb{E}[\theta]\hat{\mu}}{\bar{\mu} - (1 - \mathbb{E}[\theta])\hat{\mu}} + \frac{(1 - \mathbb{E}[\theta])(\bar{\mu} - \hat{\mu})\hat{\mu}\mathbb{E}[\theta]}{(\bar{\mu} - (1 - \mathbb{E}[\theta]))^2} \tag{4.37}$$

The second derivative is will show that we a finding a local maximum. Set this derivative equal to zero and solve for the optimal $\hat{\mu}$.

$$\hat{\mu}^* = \bar{\mu}\frac{1 - \sqrt{\mathbb{E}[\theta]}}{1 - \mathbb{E}[\theta]} \tag{4.38}$$

The lower end of workers that enroll is a decreasing function of the university's reputation. If the university is known to be no good ($\mathbb{E}[\theta] = 0$), then they set $\hat{\mu}^* = \bar{\mu}$. So, no workers enroll. As reputation increases, the university lowers $\hat{\mu}^*$. When the university is known to be the most knowledgable ($\mathbb{E}[\theta] = 1$), the lower bound of workers approaches $\hat{\mu}^* = \frac{\bar{\mu}}{2}$. Exactly half of the workers looking to go to school choose to enroll.

We can now back out the optimal grading scheme of the university.

$$\pi^* = \pi_\theta(\hat{\mu}^*) \tag{4.39}$$

$$= \frac{\hat{\mu}^*}{1-\hat{\mu}^*}\left(\frac{\mathbb{E}[\theta]}{\bar{\mu} - (1-\mathbb{E}[\theta])\,\hat{\mu}^*} - 1\right) \tag{4.40}$$

Now insert the equation for $\hat{\mu}^*$ we just solved for and simplify.

$$\pi^* = \frac{\sqrt{\mathbb{E}[\theta]} - \mathbb{E}[\theta] - \bar{\mu}\left(1 - \sqrt{\mathbb{E}[\theta]}\right)}{1 - \mathbb{E}[\theta] - \bar{\mu}\left(1 - \sqrt{\mathbb{E}[\theta]}\right)} \tag{4.41}$$

The optimal π is an increasing function of the university's reputation.

These equations illustrate story of why the good universities experience grade inflation that isn't based on them having better students. We see that $\hat{\mu}^*$ is a decreasing function of $\mathbb{E}[\theta]$ while π^* is an increasing function of $\mathbb{E}[\theta]$. The story this model presents is that back in the day, before Harvard (any high quality university) established their reputation, only the very best students went to Harvard. Also for the students that did go to Harvard, it was very difficutly to pass and get good grades. However, Harvard's reputation grew over time. Now everyone knows that Harvard is excellent. So now more students go to Harvard bringing their average student quality down (but still not nearly all students). Also Harvard can now be more lenient in their grading. It is easier to get a passing grade at Harvard now than it once was.

4.5 Conclusion

In this paper, I presented a model of universities as being hired by a worker to evaluate them for potential employers. I show how reputation concerns can enter the problem of an information designer with imperfect inforamtion. I believe these concerns can enter into all standard situations of a strategic player sharing advice or information. The model gives an explanation for grade inflation, particularly as concentrated among the best universities.

There are many papers on related topics. I would like to mention only a few that I found essential to this paper. Another paper with a model seeking to explain grade inflation is Boleslavesky and Cotton (2015). In their model grade inflation is driven by the university investing more in the quality of their education. My paper is more after the models of Spence (1973) and various extensions. The role of the university here is to help the employers distinguish the productive and unproductive workers. I use a framework and techniques recent developed by information design papers such as Kamenica and Gentzkow (2011) and Kolotilin et al. (2017). Lastly, my separating equilibrium conditions resemble those that were well studied by Kaya (2009).

Chapter 5

Conclusion

This is the conclusion of my dissertation on application of information design to finance, politics, and education. A large recent interest in information design problems has led to a wealth of new techniques. I hope that this dissertation has shown how these techniques can be applied in novel and interesting ways to important questions across all fields of economics. I believe there are numerous avenues yet to be explored where these types of methods will be beneficial. It may be looking at insider trading regulation as a problem of a designer controlling the information content in the market, or barganing with incomplete information in major league baseball contracts. There are many projects I will yet persue in this vein, and I hope others follow suit. Thank you.

Bibliography

Aumann, R. J. and Maschler, M. B. (1995). *Repeated Games with Incomplete Information*. MIT Press, Cambridge, MA.

Bergmann, D., Brooks, B., and Morris, S. (2015). The limits of price discrimination. *American Economic Review*, 105(3):921–957.

Boleslavesky, R. and Cotton, C. (2015). Grading standards and education quality. *American Economic Journal: Microeconomics*, 7(2):248–279.

Carson, J. and Crespin, M. (2004). The effect of state redistricting methods on electoral competition in the united states house of representatives races. *State Politics and Policy Quarterly*, 4(4):455–469.

Coate, S. and Knight, B. (2007). Socially optimal redistricting: A theoretical and empirical exploration. *Quarterly Journal of Economics*, 122(4):1409–1471.

Crawford, V. P. and Sobel, J. (1982). Strategic information transmission. *Econometrica*, 50(6):1431–1451.

De Meyer, B. (2010). Price dynamics on a stock market with asymmetric information. *Games and Economic Behavior*, 69(1):42–71.

De Meyer, B. and Saley, H. M. (2003). On the strategic origin of brownian motion in finance. *International Journal of Game Theory*, 31(2):285–319.

Ely, J. (2017). Beeps. *American Economic Review*, 107(1):31–53.

Ely, J., Frankel, A., and Kamenica, E. (2015). Suspense and surprise. *Journal of Political Economy*, 123(1):215–260.

Fama, E. (1970). Efficient capital markets: A review of theory and empirical work. *Journal of Finance*, 25(2):383–417.

Foster, F. D. and Viswannathan, S. (1996). Strategic trading when agents forcast the forcasts of others. *Journal of Finance*, 51(4):1437–1478.

Friedman, J. N. and Holden, R. T. (2008). Optimal gerrymandering: Sometimes pack, but never crack. *American Economic Review*, 98(1):113–144.

Glosten, L. R. and Milgrom, P. R. (1985). Bid, ask and transaction prices in a specialist market with heterogeneously informed traders. *Journal of Financial Economics*, 14(1):71–100.

Hanson, R. (2003). Combinatorial information market design. *Information Systems Frontiers*, 5(1):107–119.

Hanson, R. (2007). Logorithmic market scoring rules for modular combinatorial information aggragation. *Journal of Prediction Markets*, 3(1):61–63.

Hellwig, M. (1980). On the aggregation of information in competitive markets. *Journal of Economic Theory*, 22(3):477–498.

Holden, C. and Subrahmanyam, A. (1992). Long-lived private information and imperfect competition. *Journal of Finance*, 47(1):247–270.

Hörner, J. and Skrzypacz, A. (2016). Selling information. *Journal of Political Economy*, 124(6):1515–1562.

Kamenica, E. and Gentzkow, M. (2011). Bayesian persuasion. *American Economic Review*, 101(6):2590–2615.

Kaya, A. (2009). Repeated signaling games. *Games and Economic Behavior*, 66(2):841–854.

Kolotilin, A., Mylovanov, T., Zapechelnyuk, A., and Li, M. (2017). Persuasion of a privatley informed receiver. *Econometrica*, 85(6):1949–1964.

Kyle, A. (1985). Continuous auctions and insider trading. *Econometrica*, 31(6):1315–1336.

Lipnowski, E. and Ravid, D. (Forthcoming). Cheap talk with trasparent motives. *Econometrica*.

Mertens, J.-F. and Zamir, S. (1977). Maximal variation of a bounded martingale. *Israel Journal of Math*, 27(3-4):252–276.

Orlov, D., Skrzypacz, A., and Zryumov, P. (2019). Persuading the priciple to wait. *Working Paper*.

Ostrovsky, M. (2012). Information aggregation in dynamic markets with strategic traders. *Econometrica*, 80(6):2595–2647.

Renault, J., Solan, E., and Vieille, N. (2017). Optimal dynamic information provision. *Games and Economic Behavior*, 104:329–349.

Samuelson, P. (1965). Proof that properly anticipated prices fluctuate randomly. *Industrial Management Review*, 6(2):41–49.

Shiller, R. (1981). Do stock prices move too much to be justified by subsequent changes in dividends? *American Economic Review*, 71(3):421–436.

Spence, M. (1973). Job market signaling. *Quarterly Journal of Economics*, 87(3):355–374.

Van Bommel, J. (2003). Rumors. *Journal of Finance*, 58(4):1499–1520.

Zhong, W. (2020). Optimal dynamic information acquisition. *Working Paper*.

Appendix

Proof of Theorem 3.1

Here I give the remaining details needed for the proof of theorem 1. Take $\mu < \mu^*$. We know that each period beliefs jump down to zero or go up to an interior point that is indifferent.

$$P(\mu) - P(0) = P(\mu_{buy}) - P(\mu) + \delta V(\mu_{buy}) \tag{1}$$

Putting the value function in, we can solve for the posterior induced after buying.

$$\mu_{buy} = P^{-1}\left(\frac{2P(\mu) - (1-\delta)P(0)}{1+\delta}\right) \tag{2}$$

from prior μ.

Beliefs being a martingale pins down what the probability of buying and selling must be.

$$p_{sell} = \frac{P^{-1}\left(\frac{2P(\mu)-(1-\delta)P(0)}{1+\delta}\right) - \mu}{P^{-1}\left(\frac{2P(\mu)-(1-\delta)P(0)}{1+\delta}\right)}. \tag{3}$$

Taking limits gives,

$$\lim_{\delta\to 1} \mu_{buy} = \mu \tag{4}$$

and therefore

$$\lim_{\delta\to 1} p_{sell} = 0. \tag{5}$$

Link the discount factor to the length of each period with $\Delta t = 1 - \delta$. We now need to show that $\lim_{\delta \to 1} \frac{p_{sell}}{\Delta t}$ converges to a finite positive number. That number is the arrival rate of the Poisson process.

$$\frac{p_{sell}}{\Delta t} = \frac{P^{-1}\left(\frac{2P(\mu)-(1-\delta)P(0)}{1+\delta}\right) - \mu}{P^{-1}\left(\frac{2P(\mu)-(1-\delta)P(0)}{1+\delta}\right)(1-\delta)} \tag{6}$$

Clearly, the numerator and demoninator are both going to zero as δ goes to one. We need to use L'Hôpital's rule.

First the derivative of the numerator.

$$\frac{\partial top}{\partial \delta} = P^{-1'}\left(\frac{2P(\mu)-(1-\delta)P(0)}{1+\delta}\right)\left(\frac{P(0)}{1+\delta} - \frac{2P(\mu)-(1-\delta)P(0)}{(1+\delta)^2}\right) \tag{7}$$

The limit of this is non-zero.

$$\lim_{\delta \to 1} \frac{\partial top}{\partial \delta} = \frac{1}{2} P^{-1'}\left(P(\mu)\right)\left(P(0) - P(\mu)\right) \tag{8}$$

Now the derivative of the denominator.

$$\begin{aligned}\frac{\partial bottom}{\partial \delta} = P^{-1'}\left(\frac{2P(\mu)-(1-\delta)P(0)}{1+\delta}\right)\left(\frac{P(0)}{1+\delta} - \frac{2P(\mu)-(1-\delta)P(0)}{(1+\delta)^2}\right)(1-\delta) \\ - P^{-1}\left(\frac{2P(\mu)-(1-\delta)P(0)}{1+\delta}\right)\end{aligned} \tag{9}$$

This limit is even simpler.

$$\lim_{\delta \to 1} \frac{\partial bottom}{\partial \delta} = -\mu \tag{10}$$

We now have our arrival rate.

$$\begin{aligned}\lambda(\mu) &= \frac{p_{sell}}{\Delta t} &(11)\\ &= \frac{P^{-1'}(P(\mu))}{2\mu}(P(\mu) - P(0)) &(12)\\ &= \frac{P(\mu) - P(0)}{2P'(\mu)\mu} &(13)\end{aligned}$$

The last equality follows from the inverse function theorem.

This gives a Poisson term on beliefs equal to

$$-\mu_t dN_t \tag{14}$$

where N_t is a standard Poisson process with arrival rate λ. The drift must be such that beliefs are a martingale.

$$drift dt = \mathbb{E}[\mu_t dN_t] \tag{15}$$

$$= \mu_t \lambda(\mu_t) dt \tag{16}$$

Beliefs follow

$$d\mu_t = \mu_t \lambda(\mu_t) dt - \mu_t dN_t. \tag{17}$$

Since the price is a differentiable function of beliefs, this gives us the process for prices.

$$dP(\mu_t) = P'(\mu_t)\mu_t \lambda(\mu_t) dt - (P(\mu_t) - P(0)) dN_t \tag{18}$$

Putting in our equation for $\lambda(\mu_t)$ gives the result.

$$dP(\mu_t) = \frac{1}{2}(P(\mu_t) - P(0)) dt - (P(\mu_t) - P(0)) dN_t \tag{19}$$

A symmetric arguement holds for $\mu > \mu^*$.

Proof of Proposition 3.2

The maximal time to full information revelation is the amount of time it take beliefs to drift to μ^* if the Poisson jump doesn't arrive. After that, beliefs immediately move to the boundary. Take $\mu_0 \leq \mu^*$. We saw from theorem 1 that

$$dP(\mu_t) = \frac{1}{2}(P(\mu_t) - P(0)) - (P(\mu_t) - P(0)) dN_t. \tag{20}$$

When the Poisson jump doesn't arrive, the price follows a smooth differentiable path. This gives a linear differential equation for prices as a function of time.

$$\frac{dP(\mu_t)}{dt} = \frac{1}{2}(P(\mu_t) - P(0)) \tag{21}$$

This implies that the price must be an exponential function in time.

$$P(\mu_t) = ce^{\frac{1}{2}t} - P(0) \tag{22}$$

We get the constant from the initial condition that the price at time 0 equals $P(\mu_0)$.

$$P(\mu_t) = (P(\mu_0) - P(0))e^{\frac{1}{2}t} - P(0) \tag{23}$$

We then solve for the time at which $P(\mu_t) = P(\mu^*)$. This gives the result.

$$t^{max} = 2\log\left(\frac{P(\mu^*) - P(0)}{P(\mu_0) - P(0)}\right) \tag{24}$$

If we start with a high initial belief, $\mu_0 > \mu^*$, the only change in the equations is the use of $P(1)$ in place of $P(0)$.

Proof of Theorem 3.2

Here I give the remaining details needed for the proof of theorem 2. Once you write the problem as choosing posteriors subject to Bayes' plausibility, the problem becomes

$$\begin{aligned} V(\mu) = \max_{\bar{\epsilon},\underline{\epsilon}} \ & |P(\mu+\bar{\epsilon}) - P(\mu)|\frac{\underline{\epsilon}}{\bar{\epsilon}+\underline{\epsilon}} + |P(\mu-\underline{\epsilon}) - P(\mu)|\frac{\bar{\epsilon}}{\bar{\epsilon}+\underline{\epsilon}} \\ & + \delta\left(V(\mu+\bar{\epsilon})\frac{\underline{\epsilon}}{\bar{\epsilon}+\underline{\epsilon}} + V(\mu-\underline{\epsilon})\frac{\bar{\epsilon}}{\bar{\epsilon}+\underline{\epsilon}}\right). \end{aligned} \tag{25}$$

First note that for small values of $\bar{\epsilon}$

$$|P(\mu+\bar{\epsilon}) - P(\mu)| \approx |P'(\mu)|\bar{\epsilon} \tag{26}$$

and the same for $\underline{\epsilon}$.

Now take a second order approximation to $V(\mu')$.

$$V(\mu+\bar{\epsilon}) \approx V(\mu) + V'(\mu)\bar{\epsilon} + \frac{1}{2}V''(\mu)\bar{\epsilon}^2 \tag{27}$$

Then

$$V(\mu+\bar{\epsilon})\frac{\underline{\epsilon}}{\bar{\epsilon}+\underline{\epsilon}} + V(\mu-\underline{\epsilon})\frac{\bar{\epsilon}}{\bar{\epsilon}+\underline{\epsilon}} = V(\mu) + \frac{1}{2}V''(\mu)\bar{\epsilon}\underline{\epsilon} \tag{28}$$

since the first order term cancels out.

Now (25) becomes

$$(1-\delta)V(\mu) = \max_{\bar{\epsilon},\underline{\epsilon}} \; 2|P'(\mu)|\frac{\bar{\epsilon}\underline{\epsilon}}{\bar{\epsilon}+\underline{\epsilon}} + \frac{1}{2}\delta V''(\mu)\bar{\epsilon}\underline{\epsilon} \tag{29}$$

subject to the posterior remaining between zero and one.

The first term is positive, so for any given level of the product $\bar{\epsilon}\underline{\epsilon}$ you would like to minimize the denominator $\bar{\epsilon}+\underline{\epsilon}$. The way to minimize the sum of two variables given a fixed level of their product is always $\bar{\epsilon}=\underline{\epsilon}$. Call it ϵ.

$$(1-\delta)V(\mu) = \max_{\epsilon} \; |P'(\mu)|\epsilon + \frac{1}{2}\delta V''(\mu)\epsilon^2 \tag{30}$$

Taking the derivative of the right hand side and setting it equal to zero yields

$$\epsilon^* = -\frac{|P'(\mu)|}{\delta V''(\mu)}. \tag{31}$$

Note that this is positive because $V(\mu)$ is concave.

Putting the solution for ϵ^* into the objective gives a second order differential equation.

$$V''(\mu)V(\mu) = -\frac{|P'(\mu)|^2}{2\delta(1-\delta)} \tag{32}$$

Define

$$\hat{V}(\mu) = V(\mu)\sqrt{2\delta(1-\delta)}. \tag{33}$$

This gives a simpler differential equation.

$$\hat{V}''(\mu)\hat{V}(\mu) = -|P'(\mu)|^2 \tag{34}$$

While this differential equation is not generally solvable for any function $P(\mu)$ we can see that $\hat{V}(\mu)$ is constant in δ.

This implies that $V''(\mu)$ is going to minus infinity at rate $\frac{1}{\sqrt{1-\delta}}$. Link the discount factor to the width of each time interval, $\Delta t = \frac{1-\delta}{\delta}$.

We now have beliefs following a binomial model,

$$\mu' - \mu = \begin{cases} \sigma(\mu)\sqrt{\Delta t} & \text{with probability } \frac{1}{2} \\ -\sigma(\mu)\sqrt{\Delta t} & \text{with probability } \frac{1}{2} \end{cases} \tag{35}$$

where

$$\sigma(\mu) = \frac{\sqrt{2}|P'(\mu)|}{\hat{V}''(\mu)}. \tag{36}$$

This gives convergence of beliefs to a Brownian motion.

$$d\mu_t = \sigma(\mu)dB_t \tag{37}$$

Itô's Lemma gives the price process.

$$dP(\mu_t) = \frac{1}{2}P''(\mu_t)\sigma^2(\mu_t)dt + P'(\mu_t)\sigma(\mu_t)dB_t \tag{38}$$