Incentives, Information, and Dynamic Games: Applications in Corporations and Schools

Shuya Li

Tepper School of Business

Carnegie Mellon University

May 2022

Abstract

This dissertation contains two chapters on the application of dynamics games in delegation and in school admission. The first chapter extends the standard delegation model to a two-period setting where the bias of the agent is unknown. We formalize the intuition that discretion encourages learning in the sense that the principal is more likely to learn the bias of the agent if she delegates more actions. Moreover we analyze environments in which it is optimal for the principal to induce full separation and learn the bias with probability one. In this case, the optimal delegation set, as a function of belief, is larger in the first period than that in the second period. This implies that a dynamic interaction facilitates more discretion than an one-shot relation.

The second chapter studies dynamic school admission when exploding offers are available. In the two-period game, schools can choose when to send out offers and offers are exploding in the sense that students have to respond within the period. When the quality of the students is not perfectly known by the schools, we show that there exists an equilibrium in which schools send out offers at different times. Specifically, the less competitive school tends to send out offers earlier than their more competitive counterpart. This is because the high quality students are more likely to reject early offers from the less desirable school and remain in the market hence the more competitive school can benefit

by waiting. Our model provides a novel framework for the dynamic school admission problem and a new angel for understanding the usage of exploding offers on markets.

Acknowledgement

I would like to express my gratitude to my committee chairs Professor Ali Shourideh and Professor Onur Kesten, as well as my committee members for their guidance and support. I would also like to thank my colleague and friend Dr. Ali Polat for our inspiring conversations especially on the second chapter. I am thankful to all the kindness and efforts from Lawrence Rapp and Laila Lee. Lastly, I am deeply indebted to my parents for their continuous support and understandings.

All errors are my own.

Contents

Chapter 1

Dynamic Delegation to An Agent with Unknown Bias

1.1 Introduction

The separation of authority and information between parties with conflicts of interests is a prevailing challenge in various economic and political situations. Optimal price control depends on the production cost and efficiency of firms, most likely unknown to regulators. State governments bear cost of patients while doctors prescribe treatments. Monetary policy makers who possess more private information regarding economies and markets might be tempted to pursue overly expansionary monetary policies to boost output and employment in the short term. Venture capitalists rely on the expertise of the entrepreneurs whose interests are not perfectly aligned with theirs. Traditional incentives are either incomplete or lacking in these contexts as monetary transfers are prohibited by either regulations or public morals or are limited by the incompleteness of the financial contracts. Instead, delegation is usually suggested as a solution to these problems in theory and widely adopted in practice. Theoretical works assume that the informed party's bias is transparent and interaction one-shot. However, in practice, most relations are repeated: regulators dealing with more or less the same firms; state governments funding hospitals every year; the needs of monetary decisions coming up multiple times during the tenure of a monetary authority; venture capitalists typically participating in multiple financing rounds. In addition, it is not always the case that the bias of the informed party is known. For instance it is well-documented that monetary authorities have time-inconsistency problems, but it is naive to assume that all do and to the same degree. For doctors as well, they may care about the well-being of patients to different extents.

To incorporate these two features - dynamics and unknown bias - that are common to practical situations but understudied by the literature, we build on the standard principal-agent delegation model pioneered by Holmstrom (1978); Holmström (1984) and extend it to two periods where the bias is constant over time but unknown to the principal at the start. One action is implemented each period; the principal prefers the actions to match the states of the world whereas the agent *might* prefer the actions *higher* than the states. As in the standard delegation model, the state is privately observed by the agent each period. That is, the agent has two pieces of private information, his bias and the state.[1] The bias is constant whereas the state is independently and identically distributed over time. In the case of doctors and patients, the degree of bias remains the same with the same doctor but the conditions or ideal treatments of patients change.

The commitment power of the principal is limited as she can only commit for current periods. Specifically the principal chooses a set of actions to delegate each period and cannot commit to ignore what she has learnt from the previous

[1] Throughout we refer to the bias but not the state as *the type of the agent*.

interaction. Potentially repeated interaction provides incentives and possibilities for learning. Should the principal learn that the agent is unbiased, full discretion is optimal for future. Anticipating this, different types of agent play a signaling game with incentives to mimic the unbiased type.

In this paper, we study a two-period model where the potential bias of the agent is binary. The last period resembles a static delegation game wherein the principal delegates optimally given her posterior belief and the agent implements his most preferred action from the delegation set. Naturally, more discretion is desirable should the principal believe that the agent is less likely to be biased. As a result, the agent chooses an action in the first period, taking into account that it will affect his second-period payoff. For the first period, we start with the study of the continuation equilibria for a given interval delegation set. We construct a continuation equilibrium similar to Cho and Sobel (1990) and show that it exists and uniquely survives D1. In this D1 equilibrium, strategies are monotone in both dimensions of private information - types and states. In addition, the sets of states in which pooling or separation occurs are convex, specifically types pooling in low states and separating in high ones. To see why this is the case, the D1 refinement dictates that the only pooling action is the lowest action available. Hence, pooling becomes more costly in higher states as the bliss point of the biased type gets further away from the pooling action. For comparative statics of continuation equilibria, we show that given more discretion, i.e. a bigger delegation set, the set of pooling states shrinks and the set of separating states grows. That is to say, in expectation the principal is more likely to learn the type of the agent by delegating more actions. This formalizes the idea that discretion is an effective instrument of the principal to acquire information about the type of the agent.

However it is not known a priori that to learn the bias is optimal for the principal, especially when learning comes at a cost. Under separation, the biased type would just go ahead with his most preferred available action whereas the unbiased type might have to go out of his way - choosing an action strictly below his and the principal's most preferred action - to deter imitation. As a result, pooling might generate a higher first-period payoff for the principal especially at low states by forcing the biased type to stay low. Thus it is not ex-ante clear that how much separation is optimal. Essentially it boils down to the relative costs and benefits of learning and our model provides an interesting setup in which they intertwine.

By studying the optimal delegation problem, we identify several environments in which full separation is optimal. In the case of threshold delegation in which the principal is limited to fixing the lower bound of the delegation set weakly below the lower bound of the state space, we show that under the uniform distribution of the state and a sufficiently small bias, it is optimal to induce a fully separating continuation equilibrium in which the principal learns the type of the agent with probability one. In addition, we show that the possibility of learning in a dynamic environment motivates more discretion compared to an one-shot interaction as in Tanner (2018). Holding the same belief, the principal delegates more actions in the first period than in the second period (identical to what she would do in a static setting). These results are robust to the distribution and bias assumptions, provided that the agent is sufficiently impatient or the principal sufficiently patient.

The results provide new understanding to the additional effect of delegation and how it drives power dynamics in organizations. Intuitively, the learning effect in a dynamic setting motivates the principal to endow the agent with more freedom at beginning of their relationship, testing his loyalty. This result speaks to the dynamics of control rights allocation in venture capital contracts. It is well documented that startup founders are prone to pursue activities that are in their personal benefits as opposed to the company's interests or to protect their own private privileges at the expense of financial returns (Hellmann (1998)). The misalignment of interests in addition to unpredictable or unverifiable contingencies makes the allocation of control rights a focal point in VC contracts. Hannan, Burton, and Baron (1996) document the increasing likelihood of the founder losing the CEO position over time. Kaplan and Strömberg (2003) point out that VC control increases as the relationship progresses although the uncertainty about the venture should decrease over time. Our model provides an explanation to why control rights are taken away in spite of less uncertainty: in the early stage of the relationship, allocation of control rights is partially motivated by the possibility of learning. As venture capitalists gain more information about the venture and founders, the learning effect of discretion goes away thus less discretion is given to the startups.

As one of the few papers that study learning persistent private information without transfer, our model provides novel understandings of delegation: how discretion can be utilized to acquire information and when acquiring information is optimal. These results are robust to the threshold and distributional assumptions provided that the discount factor of the agent is sufficiently small or that of the principal sufficiently large. Moreover, we demonstrate that since discretion is used to motivate separation, i.e. learning, it might shrink over time as the principal learns more about the agent, consistent with the observations from venture capital contracts. This result also sheds new light in power dynamics in relationships with restrictions in replacing agents such as bureaucratic relations.

1.2 Literature Review

Our framework extends the standard delegation model, a rich line of literature started by Holmstrom (1978); Holmström (1984). Notably, Melumad and Shibano (1991) and Alonso and Matouschek (2008) characterize the optimal delegation set in the static environment with no uncertainty on agent's bias.

Adding uncertainty to agent's preference, Frankel (2014) looks at the max-min optimal mechanism with multiple decisions and constant bias in decisions (but not in states). His paper provides sufficient conditions for aligned delegation to be max-min optimal; under aligned delegation, every type of agent behaves as if he maximizes the principal's payoff. In terms of set-up, Tanner (2018) is probably the closest to ours. Albeit static, his model features an agent with unknown bias that is discrete and one-directional. The main result, derived under rather general conditions of preferences and distribution of states, states that a pooling contract dominates screening. That is, even if the principal learns about agent's bias, this information will turn out worthless in an one-shot environment. He suggests that learning might be more useful in a repeated interaction, which is the focus of this paper. Tanner (2018) generalizes the above-mentioned paper by extending the environment under which pooling is optimal. It also gives characterization of when screening is optimal. Our paper complements his in the exploration of delegation with unknown bias, however we do not assume that the principal has the commitment power needed for contract menus.

Our paper also relates to the literature of power dynamics in organizations. Two recent papers, J. Li, Matouschek, and Powell (2017) and Lipnowski and Ramos (2020), consider project adoptions in infinitely repeated games. In J. Li, Matouschek, and Powell (2017), the agent has private information about project availability, similarly to Aghion and Tirole (1997). In Lipnowski and

Ramos (2020), the private information concerns returns of projects. The agents have empire-building motives and always prefer adoption. Both papers find that optimal relational contracts have a flavor of *dynamic capital budget*. J. Li, Matouschek, and Powell (2017)'s equilibria eventually enter one of the absorbing states where the principal either rubberstamps the agent's recommendations or completely ignores them. On the other hand, equilibria of Lipnowski and Ramos (2020) only see the loss of power of the agent. In a similar vein, Frankel (2016) presents an environment in which (discounted) quota is optimal and allows for a finite horizon. In his model, agent's preference is unknown but only depends on actions not states. Unlike the previous two papers, Frankel (2016) adopts the mechanism design approach and endows the principal unlimited commitment power with transfers. Guo and Hörner (2017) also analyze the dynamic mechanism problem with no transfers. Their mechanism is different from the discounted quota. Instead of giving a total amount, the mechanism offers a number of units the agent can produce without question asked.

In this regard, our model departs from his by limiting the commitment power of the principal. Non-commitment, as shown by Laffont and Tirole (1987); Laffont and Tirole (1988), gives rise to novel features such as ratchet effect. Arguedas and Rousseau (2012) apply this idea to studying compliance with environmental rules. In their setting, firms might over-comply in the first period to mimic the most efficient type, hoping to reduce the probability of being monitored in later periods. This paper also looks at comparative statics of continuation equilibria and finds that a lower monitoring probability in the first period induces more separation. Different from Arguedas and Rousseau (2012), our model incorporates two dimensions of private information, one constant over time while the other not.

There are few empirical studies on discretion. Hoffman, Kahn, and D. Li (2015) is one of the firsts that offer a direct assessment of discretion in the hiring process. In their framework, managers have accesses to two pieces of information: resumes/test scores (hard information) and interviews (soft information). The company only observes the hard information. However the goals of the managers and the company are not perfectly aligned. For instance Rivera (2012) documents the important role of 'shared leisure activities' in hiring processes. It is hard to justify how such a factor affects working performances. Furthermore Hoffman, Kahn, and D. Li (2015) observe heterogeneity in exerting discretion which they interpret as the result of various levels of biases. This observation provides justification for our set-up in which the principal does not know the bias of the agent at the beginning of their relationship and can only acquire the information through interactions.

1.3 The Model: Preliminaries

1.3.1 Set Up

There are one principal (P, she) and one agent (A, he) and they interact for $T = 2$ periods. At the beginning of the game, the bias of the agent $b \in \{0, B\}$ is realized and privately observed by the agent. The bias of the agent b stays constant over the time; $b = B \in (0, 1]$ will be referred to as *the biased type* and $b = 0$ *the unbiased type*. The principal does not know the realization of the bias but only that the agent will be biased with probability p and unbiased with probability $1 - p$ where $p \in (0, 1)$.

Each period, the principal chooses a convex set of actions to delegate $[\underline{d}_t, \bar{d}_t] \equiv D_t \subseteq \mathcal{D} \subseteq \mathbb{R}$.[2] The agent is privately informed of the state of the world

[2] $\mathcal{D}$ is a compact subset of the real line. Without loss of generality, we restrict our attention to *non-redundant optimal delegation set* (Tanner (2018)). A delegation set D is *non-redundant*

$\theta_t \in \Theta = [0, 1]$ which follows an atomless distribution $f(\theta)$ and is identical and independent over time. He then implements an action from the delegation set, i.e. $d_t \in D_t$. At the end of each period the action implemented as well as the state are observed by the principal.

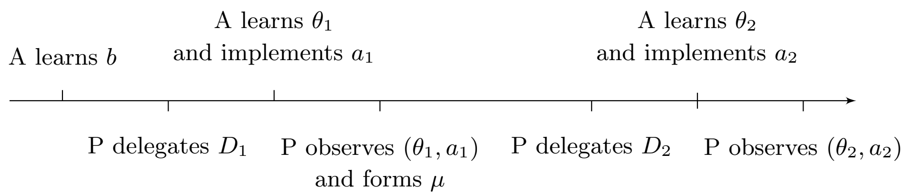

Figure 1.1: Timeline of the Dynamic Delegation Game Without Commitment

The periodic payoff of the principal is given by

$$v(a_t, \theta_t) = -(a_t - \theta_t)^2 \tag{1.1}$$

The payoff of the agent of type b is given by

$$u(b, a_t, \theta_t) = -(a_t - \theta_t - b)^2 \tag{1.2}$$

Discount factors are 1 for both parties thus the total payoff of the game is just the summation of periodic payoffs.

Trivially, $v(a, \theta)$ and $u(b, a, \theta)$ are respectively maximized at θ and $\theta + b$. We call these *the most preferred actions* at θ, denoted by $y^P(\theta)$ and $y^b(\theta)$. Sometimes, these actions might not be available from D. We call the constrained maximizers, uniquely defined by $\text{argmax}_{a \in D}\, v(a, \theta)$ and $\text{argmax}_{a \in D}\, u(b, a, \theta)$, *the most preferred available actions* at θ under D, denoted by $y^P_D(\theta)$ and $y^b_D(\theta)$.

The game is solved backwards with the solution concept PBE. First we characterize the optimal delegation set in the last period D_2^* given any posterior

if every action in the set will be optimally chosen by either type in some state.

$\mu \in [0,1]$ obtained endogenously from the first period. To analyze the first period, we start by describing agent's strategies in any continuation equilibrium Laffont and Tirole (1987). A continuation equilibrium is a PBE with a fixed first-period delegation set D_1. Given D_1, the two types play a signaling game in which the biased type has incentives to mimic the unbiased type. Lastly we study the full PBE under the restriction $\underline{d}_1 \leq 0$. One natural interpretation for this case is threshold delegation, a widely adopted practice. For instance, managers are usually endowed with budget limits but rarely with minimum investment amounts; producers are regulated by emission caps but not floors.

1.3.2 Last Period

The last-period game is played in the same way as the one-shot game. Having observed θ_2, the agent chooses his *most preferred available action* $y^b_{D_2}(\theta_2)$. The maximization problem of the principal given a posterior μ is

$$\max_{D_2 \subseteq \mathcal{D}} -\mu \int_0^1 (y^B_{D_2}(\theta_2) - \theta_2)^2 f(\theta) d\theta_2 - (1-\mu) \int_0^1 (y^0_{D_2}(\theta_2) - \theta_2)^2 f(\theta) d\theta_2 \quad (1.3)$$

Lemma 1. *The optimal delegation set in the last period is given by* $\underline{d}^*_2 = 0$ *and* $\bar{d}^*_2 \in (0,1]$ *and is decreasing in* μ*, i.e.* $\forall \mu' \geq \mu, \bar{d}^*_2(\mu') \leq \bar{d}^*_2(\mu)$.

In words, Lemma 1 says that the last period problem is effectively a threshold delegation problem and that a better posterior, in the sense that the principal believes that it's more likely the agent is unbiased, leads to more discretion in the second period.

Since the potential bias is upward, raising the lower bound above 0 would not do the principal any good in the last period. Whenever the agent moves the action downwards in the last period, he is moving it closer to the bliss point of the principal thus giving the agent more downward discretion in the last period is always beneficial for the principal. Then we show in the appendix proof that the last-period value function of the principal is submodular in belief and the upper bound. Therefore, the optimal upper bound is decreasing in belief. In addition, the last-period value function of the principal is decreasing and convex in posterior

Lemma 2. *$V_2^P(\mu)$ is decreasing and convex in μ.*

The value function of the biased type is given by

$$V_2^B(\mu) = -\int_{\max\{0,\bar{d}_2^*(\mu)-B\}}^{1} (\bar{d}_2^*(\mu) - \theta - B)^2 f(\theta) d\theta, \tag{1.4}$$

and of the unbiased type by

$$V_2^0(\mu) = -\int_{\bar{d}_2^*(\mu)}^{1} (\bar{d}_2^*(\mu) - \theta)^2 f(\theta) d\theta. \tag{1.5}$$

Apparently, both types' value functions are increasing in $\bar{d}_2^*$ thus decreasing in μ. In addition, the incremental benefit from a better posterior is always higher for the biased type than the unbiased as shown in the Lemma below.

Lemma 3. *For any $\mu' < \mu$,*

$$V_2^B(\mu') - V_2^B(\mu) > V_2^0(\mu') - V_2^0(\mu) \tag{1.6}$$

This is crucial for the monotonicity in actions in the continuation equilibrium. It is essentially equivalent to the common assumption in the signaling game literature which guarantees that a higher typer would strict prefer a higher action if the lower type weakly prefers it.

1.4 First Period: Continuation Equilibria

We start the analysis of the first period from the agent's perspective. In this section, we characterize best responses of both types in any continuation equilibrium (Laffont and Tirole (1987)) induced by any delegation set D_1. Given a realization of θ_1, the two types play a signaling game in which both want to convince the principal that they are unbiased. Following Cho and Sobel (1990), we identify the unique continuation equilibrium that survives D1; its existence is guaranteed by construction.[3] In this equilibrium, at most three cases can arise depending on the realization of the state. In low states, types pool with probability 1 and they only pool over the lower bound of the delegation set. In intermediate states, types pool and separate with positive probabilities. In high states, they separate with probability 1 and D1 selects the Riley outcome (Riley (1979)). In this case, the biased type takes his most preferred available action and the unbiased type the least costly action that achieves separation.

For the rest of this section, we omit subscription 1 and D_1 when no confusion arises.

1.4.1 Analogy to Signaling Game

To use the construction of Cho and Sobel (1990), we adapt our model into a one-shot canonical signaling game with one-dimensional private information. First we give a brief description of the game a la Cho and Sobel (1990). There are one sender and one receiver. The sender has a piece of private information, as his type, $t \in T$ where T is finite. Having observed his type t, the sender sends a message $m \in M \subseteq [0, C]$ to the receiver who then takes an action $a \in A$ from a

[3] D1 is one of the most used refinements for signaling games. Similar to intuitive criterion (Cho and Kreps (1987)), it examines the plausibility of equilibria by looking at off-equilibrium-path beliefs. According to D1, off-equilibrium-path beliefs can only put positive mass to types that are most likely to deviate from their equilibrium strategies. D1 is stronger than intuitive criterion and establishes uniqueness in our model, which is not guaranteed under intuitive criterion.

compact interval of $\mathbb{R}$. The payoff of the sender is given by $u(t, m, a)$ and that of the receiver by $v(t, m, a)$.

As for our model, the agent is the sender and the principal the receiver. The type of the sender t is equivalent to the bias of the agent in our model: $b \in B$ where B is binary thus finite. The agent takes a decision $a \in A$, analogous to the message of the sender. The principal, having observed the decision, responds with $\mu \in [0, 1]$. Given θ, the payoff of the agent is represented by

$$u(b, a, \mu|\theta) = -(a - \theta - b)^2 + V_2^b(\mu) \tag{1.7}$$

and the principal

$$v(b, a, \mu|\theta) = -(a - \theta)^2 + V_2^P(\mu) \tag{1.8}$$

In our model, the state though also unknown to the principal is not a type of the agent in the sense of a signaling game. In any continuation equilibrium the agent takes an action a given the realization of θ_1 and the principal observes θ_1 before responding with μ. Consequently, the strategies are not correlated across states, which allows that we further decompose the continuation equilibrium into signaling equilibria at θ.

Formally, we define the strategies of the agent and the equilibrium at θ. A mixed strategy of the agent is represented by $m(\cdot|b, \theta)$ which is a probability distribution over D. By an abuse of notation, we write $m(b, \theta) = a$ if type b plays a with probability 1 at state θ. The equilibrium at θ is a PBE given θ and D consisting of a triple $\{m(a|b, \theta)_{b \in B}, \mu(a, \theta)\}$ that satisfies sequential rationality and consistency.

Sequential Rationality For both b, if $m(a'|b, \theta) > 0$, then

$$a' \in \operatorname*{argmax}_{a \in D} u(b, a, \mu(a, \theta)|\theta)$$

.[4]

Consistency If $pm(a|B,\theta) + (1-p)m(a|0,\theta) > 0$, then

$$\mu(a,\theta) = \frac{pm(a|B,\theta)}{pm(a|B,\theta) + (1-p)m(a|0,\theta)}$$

Let $u^*(b|\theta)$ denote the equilibrium utility of type b at θ

$$u^*(b|\theta) = \int_{a\in D} u(b,a,\mu(a,\theta))m(a|B,\theta) \tag{1.9}$$

Abusing of terminology, we say that the equilibrium at θ is *separating* if the two types separate with probability 1 given θ; the equilibrium at θ is *pooling* if the two types pool with probability 1 given θ; the equilibrium at θ is *semi separating* if the two types pool and separate with positive probabilities given θ. For simplicity, we will drop 'at θ' and θ in payoffs, strategies, and beliefs when no confusion arises.

To refer to continuation equilibria, we will be more explicit and spell out the whole name. Given d, a continuation equilibrium is *fully separating* if the two types separate with probability 1 at all $\theta \in \Theta$. Similarly, a continuation equilibrium is *fully pooling* if the two types pool with probability 1 at all $\theta \in \Theta$.

1.4.2 Equilibrium Construction

Now we construct $\bar{u}^*(b)$ and $\bar{m}^*(\cdot|b)$ for both types. It is shown in A.1 that these consist the unique D1 equilibrium and its existence is guaranteed by construction. To achieve this, we first show that any D1 equilibria, should they exist, generate equilibrium payoffs $\bar{u}^*(b)$ for both types. Then we construct off-equilibrium beliefs consistent with D1 and prove that no deviation could be profitable for either type with $\bar{m}^*$. Finally we demonstrate uniqueness.

[4]Throughout the paper, we restrict our attention to strategies with finite support.

The construction is done inductively from the biased type. Essentially he compares his payoffs under separation and pooling. If pooling yields higher payoff, he pools whereas he separates when pooling is so costly such that even the most desirable posterior belief cannot compensate. Define

$$\bar{u}(B) = \max_{a \in D} u(B, a, 1) \tag{1.10}$$

clearly $\bar{u}(B) = u(B, y_D^B(\theta), 1) = -(y_D^B(\theta) - \theta - B)^2 + V_2^B(1)$. $\bar{u}(B)$ represents the maximum payoff of B should he separate with probability 1.

(i) If $\bar{u}(B) \geq u(B, \underline{d}, 0)$, then $\bar{u}^*(B) = \bar{u}(B)$.

(ii) If $\bar{u}(B) < u(B, \underline{d}, p)$, then $\bar{u}^*(B) = u(B, \underline{d}, p)$.

(iii) If $u(B, \underline{d}, p) \leq \bar{u}(B) < u(B, \underline{d}, 0)$, then $\exists \lambda \in (0, p]$ such that $\bar{u}^*(B) = \bar{u}(B) = u(B, \underline{d}, \lambda)$.

Case (ii) is when the biased type prefers pooling to separation; set $\bar{u}^*(0) = u(0, \underline{d}, p)$. In case (iii), the biased type pools with positive probability such that the posterior belief conditional on observing $\underline{d}$ is λ; set $\bar{u}^*(0) = u(0, \underline{d}, \lambda)$.

In case (i), the biased type prefers separation to pooling and $\bar{u}^*(0)$ is defined below. First consider a relaxed problem

$$\begin{aligned} &\max_{a \in D} u(0, a, 0) \\ &s.t. \;\; \bar{u}^*(B) \geq u(B, a, 0) \end{aligned} \tag{RP}$$

It is a relaxed problem in the sense that the unbiased type maximizes his payoff *conditional on* separation with probability 1. Since $\underline{d}$ satisfies the constraint, the feasible set is non-empty and the maximization well-defined. Let $\bar{u}(0)$ be the maximum value of RP. By setting $\bar{u}^*(0) = \bar{u}(0)$, we complete the construction of $\bar{u}^*$.

We now construct strategies $\bar{m}^*$. In case (ii), $\bar{m}^*(B) = \bar{m}^*(0) = \underline{d}$. In case (iii), $\bar{m}^*(0) = \underline{d}$; $\bar{m}^*(\underline{d}|B) = \gamma$ and $\bar{m}^*(y_D^B(\theta)|B) = 1 - \gamma$ where γ satisfies $\lambda = \frac{p\gamma}{p\gamma+1-p}$.

For case (i), $\bar{m}^*(B) = y_D^B(\theta)$ and to define $\bar{m}^*$ for the unbiased type, we write the constraint of RP explicitly

$$-(y_D^B(\theta) - \theta - B)^2 + V_2^B(1) \geq -(d - \theta - B)^2 + V_2^B(0) \tag{1.11}$$

and denote $\tilde{s}(\theta)$ as the lower decision that binds the constraint

$$\tilde{s}(\theta) = B + \theta - \sqrt{\Delta V_2^B(0,1) + (y_D^B(\theta) - \theta - B)^2} \tag{1.12}$$

It is straightforward to check that the maximizer to RP is $s(\theta) \equiv \min\{y_D^0(\theta), \tilde{s}(\theta)\}$. By setting $\bar{m}^*(0) = s(\theta)$, we finish the whole construction. It is clearly that $\bar{m}^*$ is monotone in b; the biased type always takes a higher action than the unbiased type. The continuity of $\bar{m}^*$ comes from the continuity of $\bar{u}(B)$.

Proposition 1. *Given any $D \subseteq \mathbb{R}$, there is a unique D1 equilibrium at each $\theta \in [0,1]$ in which agent of type b employs $\bar{m}^*(\cdot|b)$ and receives equilibrium payoff $\bar{u}^*(b)$.*

1.4.3 Continuation Equilibrium

As we have shown, there are three possibilities for the equilibrium at θ. For any fixed D, the state space can be partitioned into at most three sets depending on what kind of equilibrium arises at θ. Below, we show that all elements of the partition are convex thus are intervals. In particular, pooling occurs at the lowest interval, then semi-separation, and separation in the highest interval. Most importantly, we demonstrate that the set of states in which separation occurs expands and that of pooling shrinks as D grows larger. That is, more discretion increases the probability that the principal learns the bias of the agent after the first period. This result formalizes the idea that learning and discretion come hand in hand.

Lemma 4.

(a) If the equilibrium at θ is separating, so is the equilibrium at θ' for all $\theta' > \theta$.

(b) If the equilibrium at θ is pooling, so is the equilibrium at θ' for all $\theta' < \theta$.

Since pooling is only possible at the lower bound $\underline{d}$, which is further away from the most preferred available action of the biased type, i.e. $y_D^B(\theta)$ at higher θ, the cost of pooling increases in θ. Consequently, pooling only occurs at sufficiently low states.

Now denote respectively $\Theta^s \equiv \{\theta \in \Theta : \textit{equilibrium at } \theta \textit{ is separating}\}$ and $\Theta^p \equiv \{\theta \in \Theta : \textit{equilibrium at } \theta \textit{ is pooling}\}$. Lemma 4 suggests that Θ^s and Θ^p are convex. Moreover, if $\Theta^s \neq \emptyset$, then $1 \in \Theta^s$; similarly if $\Theta^p \neq \emptyset$, then $0 \in \Theta^p$. Similarly denote $\Theta^m \equiv \{\theta \in \Theta : \textit{equilibrium at } \theta \textit{ is semi-separating}\}$. It follows trivially that Θ^m is also convex since $\Theta^m = \Theta \backslash (\Theta^s \cup \Theta^p)$.

Proposition 2. *Given any $D = [\underline{d}, \bar{d}]$ and $D' = [\underline{d}', \bar{d}']$ such that $D \subseteq D'$, we have*

(a) $\Theta^s_D \subseteq \Theta^s_{D'}$;

(b) $\Theta^p_{D'} \subseteq \Theta^p_D$;

(c) For any $\theta \in \Theta^m_D \cup \Theta^m_{D'}$, $m(\underline{d}|B, \theta) > m(\underline{d}'|B, \theta)$.

The first two parts are straightforward; the last part says that the probability of pooling decreases at any state as D gets larger. Intuitively, a larger delegation set offers a (weakly) better $y_D^B(\theta)$ thus a (weakly) higher payoff under separation. Correspondingly, the cost of pooling increases. This proposition formalizes the idea that learning and discretion come hand in hand. A larger delegation set induces learning by encouraging separation and discouraging pooling. On expectation, the principal is more likely to learn the type of the agent if she endows more discretion in the first period.

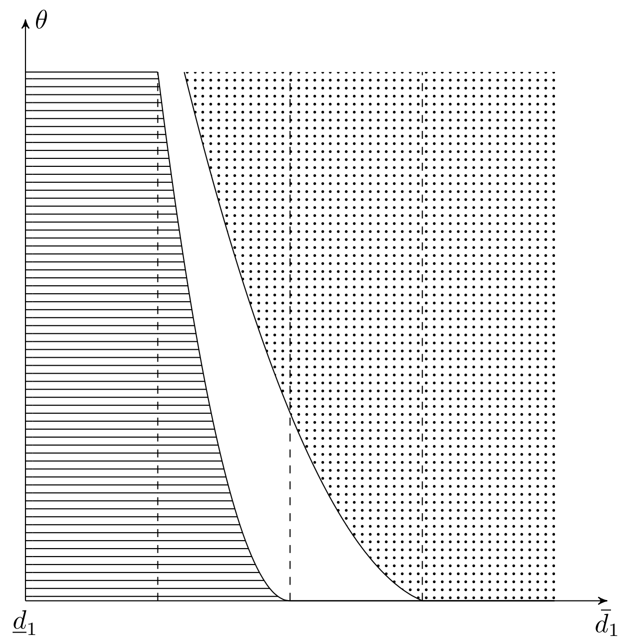

Figure 1.2: The lined area denotes the set of states with pooling and the dotted areas the set with separating. Given a $\underline{d}_1$, the separating (pooling) set expands (shrinks) as $\bar{d}_1$ rises.

1.5 Optimal Discretion

Having established that learning is feasible for the principal, in this section we investigate the optimality of learning, i.e. the delegation problem of the principal in the first period. By delegating different sets of actions, the principal could induce different continuation equilibria and probabilities of learning. More separation unambiguously leads to a higher continuation value for the second period, however learning does not come at zero cost. The first period payoff could suffer as the biased type might choose an excessively high action under separation. Essentially, how much separation is optimal depends on the relative cost and benefits of learning. To better understand the tradeoffs, we take a closer look into the principal's problem in the first period.

The principal's expected payoff at a pooling state $\theta \in \Theta_D^p$ is given by

$$-(\underline{d} - \theta)^2 + V_2^P(p) \tag{1.13}$$

at a separating state $\theta \in \Theta_D^s$ by

$$-p(y_D^B(\theta) - \theta)^2 - (1-p)(s(\theta) - \theta)^2 + pV_2^P(1) + (1-p)V_2^P(0) \tag{1.14}$$

and at a mixing state $\theta \in \Theta_D^m$ by

$$-p(1-\gamma)(y_D^B(\theta)-\theta)^2-(1-p+p\gamma)(\underline{d}-\theta)^2+p(1-\gamma)V_2^P(1)+(1-p+p\gamma)V_2^P\left(\frac{p\gamma}{p\gamma+1-p}\right) \tag{1.15}$$

Putting these pieces, we can write the objective function of the principal as the following

$$\max_D \int_{\theta\in\Theta_D^p} \left[-(\underline{d} - \theta)^2 + V_2^P(p)\right] + \int_{\theta\in\Theta_D^s} p[-(y_D^B(\theta) - \theta)^2 + V_2^P(1)] + (1-p)[-(s(\theta) - \theta)^2 + V_2^P(0)]$$
$$+ \int_{\theta\in\Theta_D^m} p(1-\gamma)[-(y_D^B(\theta) - \theta)^2 + V_2^p(1)] + (1-p+p\gamma)\left[-(\underline{d} - \theta)^2 + V_2^P\left(\frac{p\gamma}{p\gamma+1-p}\right)\right] dF(\theta) \tag{1.16}$$

By Lemma 4, we can rewrite the limits of the integrals $\Theta_D^p = [0, \bar{\theta}_D^p]$ and

$\Theta_D^s = [\underline{\theta}_D^s, 1]$ with

$$\underline{\theta}_D^s = \begin{cases} 1 & \textit{if } \Delta V_2^B(0,1) > (\underline{d} - 1 - B)^2 - (y_D^B(1) - 1 - B)^2 \\ 0 & \textit{if } \Delta V_2^B(0,1) \leq (\underline{d} - B)^2 - (y_D^B(0) - B)^2 \\ \frac{1}{2}\left(y_D^B(\underline{\theta}_D^s) + \underline{d} + \frac{V_2^B(0,1)}{y_D^B(\underline{\theta}_D^s) - \underline{d}}\right) - B & \textit{else} \end{cases}$$

and

$$\bar{\theta}_D^p = \begin{cases} 1 & \textit{if } \Delta V_2^B(p,1) > (\underline{d} - 1 - B)^2 - (y_D^B(1) - 1 - B)^2 \\ 0 & \textit{if } \Delta V_2^B(p,1) \leq (\underline{d} - B)^2 - (y_D^B(0) - B)^2 \\ \frac{1}{2}\left(y_D^B(\bar{\theta}_D^p) + \underline{d} + \frac{V_2^B(p,1)}{y_D^B(\bar{\theta}_D^p) - \underline{d}}\right) - B & \textit{else} \end{cases}$$

In words, $\bar{\theta}_D^p$ denotes the highest state in which the two types pool and $\underline{\theta}_D^s$ the lowest state in which the two types separate.

As one can see, the problem for the principal is cumbersome. By varying her choices, she not only changes the states in which the agent pools and separates but also the actions taken in each state. In addition, it is not known a priori that the principal prefers separation to pooling at any state. Separation unambiguously provides a larger expected second-period payoff due to the convexity of the value function as well as a higher expected first-period payoff from the unbiased type; yet it might come at a cost of a lower first-period payoff from the biased type. In the section below, we impose several assumptions in order to make the optimization problem more tractable. First, we limit our attention to threshold delegation, widely seen in practice. That is, the principal only has the upper bound of the delegation set at her disposal with the lower bound exogenously fixed at the lower bound of the state space. While this assumption might be innocuous in various delegation environments, it might not be the case here. Intuitively, the principal can increase payoffs from the pooling states by lifting the lower bound. Then we discuss the case where the state follows a uniform distribution. We first show that full pooling is never optimal in this setting. That is, in expectation the principal always wants to know more about the agent. Moreover, given a sufficiently small bias, it is optimal to induce a full separation equilibrium so that the principal learns the type of the agent with probability one. Under the same parameter range, we demonstrate that the optimal delegation set is bigger in the first period than that in the second period, holding the same belief. In other words, discretion decreases over time as the relationship progresses.

1.5.1 Threshold Delegation

In many economically important situations, the principal only has the upper bound of the delegation set at disposal. For instance, a manager is usually allowed to invest/borrow up to a certain limit but rarely forced to spend/borrow more than a certain amount; monopolists are regulated with price ceilings but much less often price floors; environmental laws set the maximum emission amount but not the minimum. Alonso and Matouschek (2008) also provide justifications for the prevalence of threshold delegation in a static environment. In this subsection, we fix the lower bound at zero and only allow the principal to choose the upper bound in the first period.[5]

To understand how change in discretion, i.e. the upper bound d affects principal's payoffs, we inspect the first order derivative. For any pooling state under d that will remain pooling under d', there will be no change to the principal's payoff which is also reflected by the derivative of integrated part being 0. For the state changing from pooling to mixing, the payoff difference is represented

[5] Our argument and result also apply to any $\underline{d}_1 \leq 0$.

by the following part

$$\frac{\partial \bar{\theta}_d^p}{\partial d}[-(\underline{d}-\bar{\theta}_d^p)^2+V_2^P(p)$$
$$+p(1-\gamma)(y_d^B(\bar{\theta}_d^p)-\bar{\theta}_d^p)^2+(1-p+p\gamma)(\underline{d}-\bar{\theta}_d^p)^2-p(1-\gamma)V_2^P(1)-(1-p-p\gamma)V_2^P(\lambda)]f(\bar{\theta}_d^p) \tag{1.17}$$

Rearrange to have

$$\frac{\partial \bar{\theta}_d^p}{\partial d}\{p(1-\gamma)[(y_d^B(\bar{\theta}_d^p)-\bar{\theta}_d^p)^2-(\underline{d}-\bar{\theta}_d^p)^2]$$
$$V_2^P(p)-p(1-\gamma)V_2^P(1)-(1-p-p\gamma)V_2^P(\lambda)\}f(\bar{\theta}_d^p) \tag{1.17'}$$

Recall that $\lambda = \frac{p\gamma}{1-p+p\gamma}$ and the convexity of the value function implies that $p(1-\gamma)V_2^P(1) + (1-p-p\gamma)V_2^P(\lambda) \geq V_2^P(p)$. However the sign of the first part is unclear. When the upper bound d is sufficiently small such that $d \leq \bar{\theta}_d^p$, apparently $y_d^B(\bar{\theta}_d^p) = d$ and $(y_d^B(\bar{\theta}_d^p)-\bar{\theta}_d^p)^2-(\underline{d}-\bar{\theta}_d^p)^2 \leq 0$. In this case, Equation 1.17 is non-negative.

Now for the state going from mixing to separation, the payoff change is given by the following

$$\frac{\partial \underline{\theta}_d^s}{\partial d}[-p(1-\gamma)(y_d^B(\underline{\theta}_d^s)-\underline{\theta}_d^s)^2-(1-p+p\gamma)(\underline{d}-\underline{\theta}_d^s)^2+p(1-\gamma)V_2^P(1)+(1-p+p\gamma)V_2^p(\lambda)$$
$$+p(y_d^B(\underline{\theta}_d^s)-\underline{\theta}_d^s)^2+(1-p)(s(\underline{\theta}_d^s)-\underline{\theta}_d^s)^2-pV_2^P(1)-(1-p)V_2^P(0)] \tag{1.18}$$

Rearrange to have

$$\frac{\partial \underline{\theta}_d^s}{\partial d}\{p\gamma[(y_d^B(\underline{\theta}_d^s)-\underline{\theta}_d^s)^2-(\underline{d}-\underline{\theta}_d^s)^2]+(1-p)[(s(\underline{\theta}_d^s)-\underline{\theta}_d^s)^2-(\underline{d}-\underline{\theta}_d^s)^2]$$
$$+(1-p+p\gamma)V_2^P(\lambda)-p\gamma V_2^P(1)-(1-p)V_2^P(0)\} \tag{1.18'}$$

We have the second part being non-positive as

$$
\begin{aligned}
&(1-p+p\gamma)V_2^P(\lambda) - p\gamma V_2^P(1) - (1-p)V_2^P(0) \\
\le &(1-p+p\gamma)[\lambda V_2^P(1) + (1-\lambda)V_2^P(0)] - p\gamma V_2^P(1) - (1-p)V_2^P(0) \\
= &p\gamma\lambda V_2^P(1) + (1-p)V_2^P(0) - p\gamma V_2^P(1) - (1-p)V_2^P(0) = 0
\end{aligned}
$$

The sign of the first part is again ambiguous. However when the uppder bound d is not too large such that $d \le \underline{\theta}_d^s$, $y_d^B(\underline{\theta}_d^s) = d$ and $s(\underline{\theta}_d^s) < d$, the first part is negative since both $y_d^B(\underline{\theta}_d^s)$ and $s(\underline{\theta}_d^s)$ are closer to $\underline{\theta}_d^s$ than $\underline{d}$. In this case, Equation 1.18 is non-negative.

Now for the states that will remain separating, we have

$$
\int_{\underline{\theta}_d^s}^1 -2p(y_d^B(\theta)-\theta)\frac{\partial y_d^B(\theta)}{\partial d} - 2(1-p)(s(\theta)-\theta)\frac{\partial s(\theta)}{\partial d} dF(\theta) \qquad (1.19)
$$

In particular, $s(\theta) \in \{\tilde{s}(\theta), y_d^0(\theta)\}$ where $\tilde{s}(\theta) = B+\theta-\sqrt{\Delta V_2^B(0,1) + (y_d^B(\theta)-\theta-B)^2}$. If $\Delta V_2^B(0,1) \ge B^2$, then $\forall \theta \in \Theta^s,\ s(\theta) = \tilde{s}(\theta)$. If $\Delta V_2^B(0,1) < B^2$, then $\exists \tilde{\theta}$ such that $\forall \theta \in \Theta^s \cap [\tilde{\theta}, \infty),\ s(\theta) = \tilde{s}(\theta)$ and

$$
\tilde{\theta} = d - B + \sqrt{B^2 - \Delta V_2^B(0,1)} \in (d-B, d)
$$

For sufficiently small d such that $d \le \underline{\theta}_d^s$, we have $y_d^B(\theta) < \theta$ for all $\theta \in [\underline{\theta}_d^s, 1]$. This also implies that $\underline{d} \le s(\underline{\theta}_d^s) \le \underline{\theta}_d^s$ and $s(\theta) \le \theta$ for all $\theta > \underline{\theta}_d^s$. Thus Equation 1.19 is non-negative in this case. In general, the first part, the payoff from the biased type, would increase then decrease whereas the second part, the payoff from the unbiased type, would alway increase as d goes up.

Lastly, we have the state that will remain mixing

$$\int_{\bar{\theta}_d^p}^{\underline{\theta}_d^s} -2p(1-\gamma)(y_d^B(\theta)-\theta)\frac{\partial y_d^B(\theta)}{\partial d} + p(y_d^B(\theta)-\theta)^2\frac{\partial \gamma}{\partial d} - p(\underline{d}-\theta)^2\frac{\partial \gamma}{\partial d} - p\frac{\partial \gamma}{\partial d}V_2^P(1) + p\frac{\partial \gamma}{\partial d}V_2^P(\lambda) + (1-p+p\gamma)\frac{\partial V_2^P(\lambda)}{\partial \lambda}\frac{\partial \lambda}{\partial d}dF(\theta) \tag{1.20}$$

Rearrange to have

$$\int_{\bar{\theta}_d^p}^{\underline{\theta}_d^s} -2p(1-\gamma)(y_d^B(\theta)-\theta)\frac{\partial y_d^B(\theta)}{\partial d} + p\frac{\partial \gamma}{\partial d}[(y_d^B(\theta)-\theta)^2 - (\underline{d}-\theta)^2] \quad p\frac{\partial \gamma}{\partial d}\left[V_2^P(\lambda) - V_2^P(1) + \frac{1-p}{1-p+p\gamma}\frac{\partial V_2^P(\lambda)}{\partial \lambda}\right]dF(\theta) \tag{1.20'}$$

Using the convexity of the value function and MVT for integral we have

$$\frac{\partial V_2^P(\lambda)}{\partial \lambda} \leq \frac{V_2^P(1) - V_2^P(\lambda)}{1-\lambda} = \frac{V_2^P(1) - V_2^P(\lambda)}{\frac{1-p}{1-p+p\gamma}}$$

thus the second part of Equation 1.20' is non-negative. The sign of the first part is ambiguous. However for sufficiently small $d \leq \bar{\theta}_d^p$, we have $y_d^B(\theta) < \theta$ for all $\theta \in [\bar{\theta}_d^p, \underline{\theta}_d^s]$ as well as $(y_d^B(\theta)-\theta)^2 - (\underline{d}-\theta)^2 \leq 0$. Thus in this case, Equation 1.20 is non-negative.

Taking all pieces together, we arrive at the necessary condition for optimality: the upper bound d has to be sufficiently large such that $d > \bar{\theta}_d^p$. And this leads to the following proposition.

Proposition 3. *Full pooling equilibrium is never optimal.*

Proof. Building on the above analysis, we are to show that any d to induce a full pooling equilibrium is such that $d \leq 1$ hence it cannot be optimal since the necessary condition requires $d > \bar{\theta}_d^p = 1$.

Pooling at $\theta = 1$ implies that

$$(y_d^B(1) - 1 - B)^2 \geq (1 + B)^2 - \Delta V_2^B(p, 1)$$

rearranging

$$y_d^B(1) \leq 1 + B - \sqrt{(1 + B)^2 - \Delta V_2^B(p, 1)}$$

If we can show that $1 + B - \sqrt{(1 + B)^2 - \Delta V_2^B(p, 1)} < 1$ then it implies that $y_d^B(1) = d < 1 = \bar{\theta}_d^p$ contradicting the necessary condition for optimality. To show that is equivalent to show that for any $p \in (0, 1)$, we have

$$\begin{aligned} 1 + B - \sqrt{(1 + B)^2 - \Delta V_2^B(p, 1)} \leq 1 \\ \Delta V_2^B(p, 1) \leq 1 + 2B \end{aligned}$$

The last inequality is true since $\Delta V_2^B(p, 1)$ is bounded from above by $\Delta V_2^B(0, 1)$ and we know that

$$\begin{aligned} \Delta V_2^B(0, 1) =& V_2^B(0) - V_2^B(1) \\ \leq& \int_{1-B}^{1} -(1 - \theta - B)^2 dF(\theta) - \int_0^1 -(\theta + B)^2 dF(\theta) \\ =& \int_{1-B}^{1} (2\theta + 2B - 1) dF(\theta) + \int_0^{1-B} (\theta + B)^2 dF(\theta) \\ \leq& (1 + 2B)B + (1 - B) = 1 + 2B^2 \end{aligned}$$

the first inequality comes from the fact $0 \leq \bar{d}_2^* \leq 1$ and the second inequality comes from the monotonicity of integral. Since $B \leq 1$, we know that $\Delta V_2^B(p, 1) \leq 1 + 2B^2 \leq 1 + 2B$ which completes the proof.

■

Intuitively this is a very straightforward result. Once we fix the lower bound at 0, full pooling means 0 being chosen with probability 1. Consider a marginal change in d that is just enough to move $\theta = 1$ away from pooling. Since d remains to be smaller than 1 as shown above, this marginal change unambiguously improves the payoff of the principal since $0 < d \leq 1$. Thus it is always optimal for the principal to induce at least some separation in the first period and acquire some information about the bias of the agent.

1.5.2 Uniform Distribution

Now we impose the assumption of the state being uniformly distributed to study exactly how much information the principal should acquire. We show that given a sufficiently small bias B, full separation is optimal for the principal. Moreover, holding the same belief, more discretion is given in the first period than in the second period, consistent with the observation that startups are losing control rights over time to venture capitalists.

Proposition 4. *For a sufficiently small B such that $B \leq \frac{1}{2}$ and $B^2 \geq \Delta V_2^B(0,1)$, the optimal delegation threshold induces a fully separating equilibrium. In addition, $\bar{d}_1^*(p) \geq \bar{d}_2^*(p)$ for any $p \in (0,1)$.*

Under the given parameter range, one can check readily that $\forall d \geq B - \sqrt{B^2 - \Delta V_2^B(0,1)} \equiv \tilde{d}$ induces a fully separating equilibrium. To prove the proposition, we first show that the objective function of principal can have at most two critical points in $d \in [\tilde{d}, 1+\tilde{d}]$ and the smaller one gives the maximum under full separation. Since going above $1 + \tilde{d}$ is never optimal, this finds the conditional maximum given full separation. Next, we show any d that does not induce full separation gives a lower payoff to the principal than $d = 1 - B$ which induces full separation. This completes the first part of the proof by showing that the conditional maximum is also the unconditional maximum. The second part is straightforward by showing the first order derivative is positive at $\bar{d}_2^*$.

Proof. Under full separation, the F.O.C. of the principal's objective function is given by

$$\int_0^1 -2p(y_d^B(\theta)-\theta)\frac{\partial y_d^B(\theta)}{\partial d} - 2(1-p)(s(\theta)-\theta)\frac{\partial s(\theta)}{\partial d}dF(\theta)$$
$$= \int_{\max\{0,d-B\}}^1 -2p(d-\theta) - \int_{\tilde{\theta}}^1 2(1-p)(\tilde{s}(\theta)-\theta)\frac{\partial \tilde{s}(\theta)}{\partial d}d\theta$$

where $\tilde{\theta} \in (d-B, d)$ such that $\forall \theta \geq \tilde{\theta},\ y_d^0(\theta) = \tilde{s}(\theta)$.[6]

The first part - payoff from the biased type, is concave for $d \leq 1$. For $d > 1$ however, the S.O.C. is $2p(d-1)$ thus is convex. The second part - payoff from the unbiased type is more cumbersome as $\tilde{s}(\theta) = B+\theta-\sqrt{\Delta V_2^B(0,1)+(d-\theta-B)^2}$. Now consider $d \in [\tilde{d}, 1+\tilde{d}]$ in which case $\tilde{\theta} \in [0,1]$; the derivative of its F.O.C. is given by

$$2(1-p)(\tilde{s}(\tilde{\theta})-\tilde{\theta})\frac{\partial \tilde{s}(\theta)}{\partial d}\bigg|_{\theta=\tilde{\theta}} - \int_{\tilde{\theta}}^1 2(1-p)\left[\left(\frac{\partial \tilde{s}(\theta)}{\partial d}\right)^2 + (\tilde{s}(\theta)-\theta)\frac{\partial^2 \tilde{s}(\theta)}{\partial d^2}\right]d\theta$$
$$= - \int_{\tilde{\theta}}^1 2(1-p)\left[\left(\frac{\partial \tilde{s}(\theta)}{\partial d}\right)^2 + (\tilde{s}(\theta)-\theta)\frac{\partial^2 \tilde{s}(\theta)}{\partial d^2}\right]d\theta \tag{1.21}$$

and the second order derivative of its F.O.C.

$$2(1-p)\left[\left(\frac{\partial \tilde{s}(\theta)}{\partial d}\right)^2\bigg|_{\theta=\tilde{\theta}} + (\tilde{s}(\tilde{\theta})-\tilde{\theta})\frac{\partial^2 \tilde{s}(\theta)}{\partial d^2}\bigg|_{\theta=\tilde{\theta}}\right] - \int_{\tilde{\theta}}^1 2(1-p)\left[3\frac{\partial \tilde{s}(\theta)}{\partial d}\frac{\partial^2 \tilde{s}(\theta)}{\partial d^2} + (\tilde{s}(\theta)-\theta)\frac{\partial^3 \tilde{s}(\theta)}{\partial d^3}\right]d\theta$$
$$=2(1-p)\left(\frac{\partial \tilde{s}(\theta)}{\partial d}\right)^2\bigg|_{\theta=\tilde{\theta}} - \int_{\tilde{\theta}}^1 2(1-p)\left[3\frac{\partial \tilde{s}(\theta)}{\partial d}\frac{\partial^2 \tilde{s}(\theta)}{\partial d^2} + (\tilde{s}(\theta)-\theta)\frac{\partial^3 \tilde{s}(\theta)}{\partial d^3}\right]d\theta \tag{1.22}$$

[6] To see why it is never optimal to go beyond $1+\tilde{d}$: in this case $\tilde{\theta} = 1$ and changing d would not change payoff from unbiased type in any state but will only allow the biased type more discretion on high states thus lower the payoff from the biased type. Hence setting $d > 1+\tilde{d}$ is never optimal.

Algebra shows that

$$\frac{\partial^2 \tilde{s}(\theta)}{\partial d^2} = -\frac{\Delta V_2^B(0,1)}{\sqrt{(B-d+\theta)^2 + \Delta V_2^B(0,1)}^3} < 0$$

for any θ and

$$\frac{\partial^3 \tilde{s}(\theta)}{\partial d^3} = \frac{3\Delta V_2^B(0,1)(B-d+\theta)}{\sqrt{(B-d+\theta)^2 + \Delta V_2^B(0,1)}^5} \geq 0$$

for any $\theta \geq \tilde{\theta} = d - \tilde{d} \Rightarrow B + \theta - d \geq \sqrt{B^2 - \Delta V_2^B(0,1)} \geq 0$. Therefore, Equation 1.21≤ 0 and Equation 1.22≥ 0. Thus the shapes of the F.O.C.s dictate that there are at most two critical points and the smaller critical point gives the maximum conditional on full separation; see Figure 1.3.

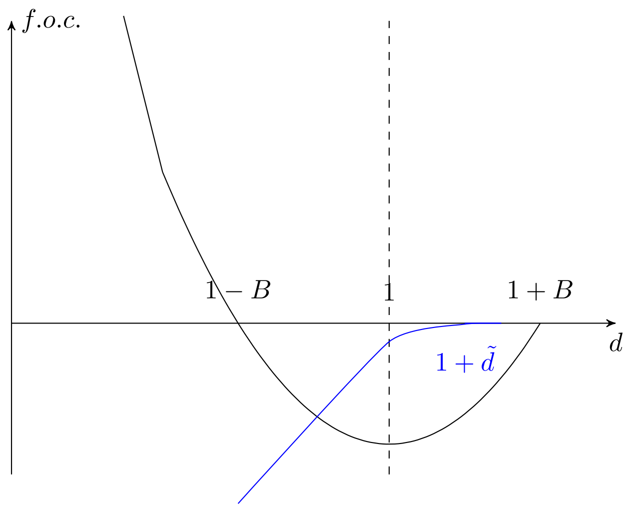

Figure 1.3: F.O.C. of payoffs from the biased type (Black) and minus of that from the unbiased type (Blue)

Next we show that any $d \in [0, \tilde{d})$ generates a lower payoff than $d = 1 - B > \tilde{d}$. To see that, we only need to compare the payoff from the biased type since a higher threshold unambiguously increases the payoff from the unbiased type. For any $d \in [0, \tilde{d})$, the maximum payoff the principal can obtain from the biased type is

$$-\int_0^{\min\{\underline{\theta}_d^s, \frac{d}{2}\}} (0-\theta)^2 d\theta - \int_{\min\{\underline{\theta}_d^s, \frac{d}{2}\}}^1 (d-\theta)^2 d\theta \tag{1.23}$$

since the biased type will either choose $y_d^B(\theta) = d$ or the pooling action $\underline{d} = 0$ when not separating and the principal prefers 0 to d if and only if $\theta \leq \frac{d}{2}$. We know that

$$\underline{\theta}_d^s = \frac{1}{2}\left(y_D^B(\underline{\theta}_D^s) + \underline{d} + \frac{\Delta V_2^B(0,1)}{y_D^B(\underline{\theta}_D^s) - \underline{d}}\right) - B = \frac{1}{2}\left(d + \frac{\Delta V_2^B(0,1)}{d}\right) - B \leq \frac{d}{2} \Leftrightarrow d \geq \frac{7B^2}{6}$$

and $\frac{7B^2}{6} < B - \sqrt{B^2 - \Delta V_2^B(0,1)}$. Moreover, Equation 1.23 is increasing in d in both cases. We can check that

$$-\int_0^{\frac{d}{2}} (0-\theta)^2 d\theta - \int_{\frac{d}{2}}^1 (d-\theta)^2 d\theta = \frac{343B^6}{864} - \frac{49B^4}{36} + \frac{7B^2}{6} - \frac{1}{3}$$

evaluated at $d = \frac{7B^2}{6}$ and

$$-\int_0^{\underline{\theta}_d^s} (0-\theta)^2 d\theta - \int_{\underline{\theta}_d^s}^1 (d-\theta)^2 d\theta$$

evaluated at $d = \tilde{d}$ are both smaller than the payoff from the biased type under full separation evaluated at $d = 1 - B$ which is given by $\frac{B^2(4B-3)}{3}$ for any B in the relevant range.

Lastly, we show that the principal optimally gives more discretion in the first period than in the second, i.e. $\bar{d}_1^*(p) \geq \bar{d}_2^*(p)$ for any belief $p \in (0,1)$. To do that, we are to show that the F.O.C. is positive at $\bar{d}_2^*(p)$. Since $\bar{d}_2^*(p) \leq 1$ for any p, it is on the concave part of the objective function thus $\bar{d}_1^*(p) \geq \bar{d}_2^*(p)$.

Conditional on full separation at the first period, the payoff from the biased type are the same for both periods. Hence we can focus on the payoff from unbiased type. Denote the derivatives of the payoffs from the unbiased type as

Φ_2 and Φ_1 respectively,

$$\Phi_2 = \; 2(1-p)\int_d^1 (\theta - d)d\theta \tag{1.24a}$$

$$\begin{aligned}
\Phi_1 =& - 2(1-p)\int_{\tilde{\theta}}^1 (\tilde{s}(\theta)-\theta)\frac{\partial \tilde{s}(\theta)}{\partial d}d\theta + \left[(1-p)(\tilde{s}(\tilde{\theta})-\tilde{\theta})^2\frac{\partial \tilde{\theta}}{\partial d}\right] \\
=& - 2(1-p)\int_{\tilde{\theta}}^d (\tilde{s}(\theta)-\theta)\frac{\partial \tilde{s}(\theta)}{\partial d} - 2(1-p)\int_d^1 (\tilde{s}(\theta)-\theta)\frac{\partial \tilde{s}(\theta)}{\partial d}d\theta \\
=& \; 2(1-p)\int_{\tilde{\theta}}^d (\theta-\tilde{s}(\theta))\frac{\partial \tilde{s}(\theta)}{\partial d} + 2(1-p)\int_d^1 \theta - d + B\left(1 - \frac{\theta + B - d}{\sqrt{\Delta V_2(0,1) + (d-\theta-B)^2}}\right)d\theta
\end{aligned} \tag{1.24b}$$

The second last equality follows from $\tilde{s}(\tilde{\theta}) = \tilde{\theta}$ and last from $\tilde{s}(\theta) = B + \theta - \sqrt{\Delta V_2(0,1) + (d-\theta-B)^2}$. Clearly $\tilde{\theta} < d$ and $\tilde{s}(\theta) \leq \theta$, the first term is non-negative. Since the second term is bigger than Φ_2, we have $\Phi_1 > \Phi_2$ at any $d \in [1-B, 1]$. This implies $\bar{d}_1^*(p) \geq \bar{d}_2^*(p)$ and completes the proof.

■

As mentioned, learning leads to a higher continuation payoff and a higher current-period payoff from the unbiased type but it might decrease the current-period payoff from the biased type. Under uniform distribution and sufficiently small biased, we show that the benefits of learning outweigh the potential cost thus it is optimal for the principal to induce a fully separating continuation equilibrium.

1.6 Discussion

1.6.1 Non-Interval Delegation

Throughout we assume the convexity of the delegation set. Practical examples of non-interval action space are rare but present, for instance, the medicare part D donut hole. In theory, interval delegation has been shown to be optimal under rather general conditions Amador and Bagwell (2013). However Tanner (2018) shows that given the possibility of screening and a sufficiently concave loss function of the principal, non-interval delegation set is optimal. Below, we briefly discuss the implication of non-interval delegation in our equilibrium construction.

We show that the unique continuation equilibrium we constructed might fail to exist or lose its uniqueness. This is because the existence and uniqueness of the equilibrium replies on the fact that under interval delegation, an arbitrarily small change in the posterior of the principal can be induced by an arbitrarily small move in action. Whereas given a non-interval delegation set, the cost of moving the posterior by a little might be too great.

Example: Non-Interval Delegation Consider this numerical example with $B = 1$ and $D = \{0, 1\}$. Clearly in any equilibrium for $\forall\theta \in [0, 1]$, the biased type would choose the higher action with probability 1 since

$$-(1-\theta-1)^2 + V_2^B(1) > -(0-\theta-1)^2 + V_2^B(0)$$
$$\Leftrightarrow \ \theta \geq \frac{\Delta V_2^B(0,1) - 1}{2}$$

where the RHS is smaller than zero. And there are two D1 equilibria depending on when the unbiased type separates. In one equilibrium, the unbiased type chooses the lower action 0 if $\theta \leq \frac{\Delta V_2^0(0,1)+1}{2}$ and pools at the higher action 1 otherwise. Under pooling, D1 dictates that $\mu(0) = 0$. In another equilibrium, everything else being the same, the unbiased type chooses the lower action 0 if $\theta \leq \frac{\Delta V_2^0(p,1)+1}{2}$. In this example, two types are pooling at the higher action instead of the lower one. This is because the "next-low" action is too costly for the unbiased to separate.[7]

1.6.2 Discounting

We have been using the assumption that there is no discounting over time for either party, i.e. $\delta^P = \delta^A = 1$. Straightforwardly, the existence and uniqueness of the continuation equilibrium will continue to hold with any $\delta \in (0, +\infty)$. However our results on optimal delegation are susceptible to discount factors. This is because by giving different weights to the current and future payoffs, the cost of learning will change accordingly for the principal. For instance, by increasing the discount factor of the agent, i.e. δ^A, separation will become more costly, which will in turn change the optimal delegation set.

Proposition 5.

1. *As $\delta^A \to +\infty$, the optimal delegation set induces full pooling and $\bar{d}_1^* = \underline{d}_1^* = E(\theta)$.*

2. *As $\delta^A \to 0$, the optimal delegation set induces full separation. In addition $\underline{d}_1^* = 0$ and $\bar{d}_1^*(p) \geq \bar{d}_2^*(p)$ for any $p \in (0, 1)$.*

The first part is straightforward as δ^A gets sufficiently large, separation become excessively costly for the principal. The cost comes from two sources. First, a large amount of discretion is needed to induce separation from the biased

[7] However it is incorrect to conclude that pooling only happens at high states. One can construct a more delicate example in which pooling occurs in extreme states, either low or high.

type as the following necessary condition shows

$$\delta^A \Delta V_2^B(p,1) \le (\underline{d} - 1 - B)^2 - (y_D^B(1) - 1 - B)^2$$

As $\delta^A \to +\infty$, eventually $[0,1] \subset D$ and the principal incurs a large amount of cost for allowing the biased type to choose upward-biased actions. Second, the unbiased type will have to skew his action downwards by a large extent for separation as shown below

$$\delta^A \Delta V_2^B(0,1) \le (s(\theta) - 1 - B)^2 - (y_D^B(1) - 1 - B)^2$$

This again incurs a cost on the principal. Since the benefit of learning is well bounded, as $\delta^A \to +\infty$ it is optimal to have a full pooling equilibrium. In a full pooling equilibrium, only 1 action will ever be chosen by either type, thus $\bar{d}_1^* = \underline{d}_1^* = E(\theta)$.

Similarly, as $\delta^A \to 0$, pooling gets excessively costly thus very little discretion is enough to induce full separation. Once we show that $\underline{d}_1^* = 0$, the rest follows the same logic as the proof of Proposition 4.

Lemma 5. *If D_1^* is optimal and induces full separation, then $\underline{d}_1^* \le 0$. Moreover, as $\delta^A \to 0$, $\underline{d}_1^* = 0$*

Proof. We prove the first part by contradiction. Suppose that there exists D that is optimal and induces full separation yet $\underline{d} > 0$. One straightforward observation is that $\underline{d} < B$ else D cannot induce a full separation. In a fully separating equilibrium, the two types separate at $\theta = 0$ and the following has to hold

$$-(y_D^B(0) - 0 - B)^2 + \delta^A V_2^B(1) \ge -(\underline{d} - 0 - B)^2 + \delta^A V_2^B(0)$$
$$\underline{d} \le B - \sqrt{\delta^A \Delta V_2^B(0,1) + (y_D^B(0) - B)^2}$$

If $B - \sqrt{\delta^A \Delta V_2^B(0,1) + (y_D^B(0) - B)^2} \le 0$, then our job is done. Suppose not then we have

$$\tilde{s}(0) = B + 0 - \sqrt{\delta^A \Delta V_2^B(0,1) + (y_D^B(0) - B)^2} > 0$$

since $\tilde{s}(0) \ge \underline{d} > 0$. Thus $s(0) = \min\{\tilde{s}(0), y_D^0(0)\} = y_D^0(0) = \underline{d}$.

Now consider D' such that $\bar{d}' = \bar{d}$ but $\underline{d}' = 0$. Apparently $D \subset D'$; by Proposition 2, D' also induces full separation. Under D', the payoff from the biased type would not change as $d' < d < B$ thus $y_D^B(\theta) = y_{D'}^B(\theta)$ for $\forall \theta \in \Theta$. As for the payoff from the unbiased type, it would increase strictly becasue $s(\theta)$ is weakly closer to θ under D' than under D for any $\theta \le \underline{d}$ and strictly for $\theta = 0$. For any $\theta > \underline{d}$, $s(\theta)$ remains the same under D' and D and this completes the proof for the first part of the lemma.

For the second part, it is straightforward with the observation that $y_D^B(0) > 0$ for optimality. Given that, $s(0) = 0$ thus actions smaller than 0 are redundant and will never be chosen.

■

Similarly, $\delta^P \to +\infty$ suggests that the principal overweights future payoffs thus optimally induces full separation. Hence, the first part of Proposition 4 holds under the same condition. The second part can be generalized following the same rationale, as stated below

Corollary 1. *As $\delta^P \to +\infty$, the optimal delegation induces full separation. In addition, $\underline{d}_1^* = 0$ and $\bar{d}_1^*(p) \ge \bar{d}_2^*(p)$ for any $p \in (0,1)$.*

These two results show that Proposition 4 is robust to a more general environment provided the discount factor of the agent is sufficiently small or that of the principal sufficiently large. Intuitively, when the discount factor of the agent is small, the incentives of imitation falls accordingly thus the cost of learning decreases. When the discount factor of the principal is large, the benefits of learning grow. In either case, we show that it is optimal for the principal to learn with probability 1, and the best way to induce full learning is to set the lower bound at zero and to give extra upward discretion, i.e. $d_1^*(p) \geq d_2^*(p)$.

1.7 Conclusion

In this paper we have shown that there exists a unique perfect Bayesian equilibrium under D1 in which the agent pools in low states and separate in high states. We formalize the idea that more discretion encourages learning in the sense that by giving the agent a bigger delegation set, the principal increases her probability of learning agent's type. Intuitively, a bigger delegation set makes it more costly for the biased type to mimic the unbiased type. Given the strategic interplay, we investigate the optimal delegation by the principal under the assumption of threshold delegation. We show that full pooling is never optimal and the principal in expectation always learns informtion about the agent's type. Moreover, given a sufficiently small bias or the agent being sufficiently impatient, it is optimal for the principal to induce full separation in the sense that she learns the bias of the agent with probability one. These results suggest that in a dynamic delegation environment, learning is not only feasible but also optimal. Last but not least, in terms of power dynamics, we show that discretion decreases over time as the motive for learning fades.

There are several directions worth exploring for future works. First, the full characterization of optimal delegation set remains to be studied. In particular, how a flexible lower bound would affect discretion and power dynamics remains an interesting question. Second, one may investigate how the commitment power of the principal affects optimal delegation. In this paper, the principal has minimal commitment power, i.e. she can not commit to offering a menu nor a contract for the entire horizon. In principle, more commitment allows better payoffs for the principal as she could always achieve at least as much as in the case with less commitment. Lastly, it would also be interesting to study the optimal amount of transparency between periods. In our set up, we assume that principal can perfectly observe the past state and action. But one can think of scenarios where this is not the case. Whether it is a strategic choice and if so what factors drive higher or lower degree of transparency in organizations? The degree of transparency could serve as a commitment device between-period for the principal as well. By muddling her observation between periods, the principal could effectively mitigate pooling incentives. All in all, our model can be easily adapted to study both information and mechanism design without transfers. Doval and Skreta (2020), a recent exploration in this realm, point out that limited commitment can be formulated as an information design problem atop the mechanism design, which should be a very exciting adventure for the future.

Chapter 2

Dynamic School Admission with Exploding Offers

2.1 Introduction

Information asymmetry causing market failures has long been studied in the economic literature. In the context of matching, incomplete information might result in miscoordination, inefficient matching outcomes, or even market failures. On the whole, a large proportion of the matching literature is devoted to understanding the static environment and to improving assignments in a one-period setup. We introduce dynamics with endogenous timing, an instrument widely used in practice, and discuss new strategies that open up when agents are provided with more time to make their decisions.

In this paper, we consider a two-period decentralized matching environment in the language of school admission[1] where schools and students that remain unmatched in the first period will go into the second period. In particular, the

[1] One can easily adopt our environment to other setups such as the labor market.

schools hold the liberty to set their paces in the sense that they can purposely remain unmatched in the first period "to observe the market for longer". The private information in our model is the quality of the student which is known by the student but not perfectly known by schools. Because of the incompleteness of information, schools might have incentives to "take their time" to gather more information by "observing the market for longer". We derive conditions of parameters under which this is indeed the case in the equilibrium. Schools endogenously choose different timings to send out offers. In particular the more competitive school waits longer than the less competitive school and utilizes the fact that high quality students who are in turn more confident in their ability are more likely to wait thus remain in the market for longer. With a formal game theoretical model, we answer the question of when schools are expected to send out offers sequentially and when to send out offers simultaneously.

Our model speaks directly to the graduate student admission process in practice. It is widely observed that different schools send out offers at different times. On one hand, less competitive schools are more prone to giving offers with a smaller window to respond, i.e. exploding offers. More competitive institutes on the other hand utilize wait-lists a lot more often. We propose that this phenomenon arises from the dynamic nature of the market and the information asymmetry about students' true qualities. Intuitively, less competitive schools need to "hurry up" to grab students that would otherwise be admitted by more competitive schools and more competitive schools are better off waiting to allow the less qualified students be filtered out first. This is consistent with our model in which the more competitive school sends offers later than its less competitive counterpart and lower quality students are more likely to accept exploding offers.

2.2 Literature Review

This paper contributes mainly to two lines of literature: decentralized school admission and games with endogenous timing.

Avery and Levin (2010) study college school admission under incomplete information. In their model, students' application strategy is not trivial. The reasons are two folds. First in their model there are two dimensions of private information - ability and preference. Students only know their preferences whereas their abilities are known by the schools. Second students can only apply to one school in the early admission process hence to some students there is the trade-off between applying to the safer school and to the better school. For early admission, schools can adopt either *early action* or *early admission*. Early action is non-binding hence students who are admitted early can still apply and go to another school should they receive a more desirable offer from the regular admission process. Early decision on the other hand is binding. Students accepted have to decide on their offers before the regular admission period and those admitted cannot apply to other schools later. In markets we are interested in such as the one for graduate students admission, there is no explicitly separated admission stages. Nonetheless, early decision which involves binding commitments works similarly to exploding offers.[2] Avery and Levin (2010) show that students could use the early admission program to signal their preferences and consequently schools would indeed favor early applicants with a lower admission threshold. Moreover, they conclude that very top-ranked schools are less enthusiastic about early admission programs which is consistent with the fact that high-ranked schools such as Harvard, Princeton, Yale, MIT, and Stanford first switched away from the binding early decision and later eliminated

[2] One major difference however is that under early decision, a student can still apply to and be accepted by the same school in the regular stage even if the student received and rejected the school's early offer. With exploding offers, this cannot happen. If a student rejected an early exploding offer, there is not another chance to receive an offer from the same school.

early admissions entirely. In their arguments, with early decision lower-ranked schools could potentially capture some highly desired students who are uncertain about their perspectives at higher-ranked schools whereas top-ranked schools are less concerned about losing top students. While our results are reminiscent to theirs, the mechanism and intuition under are vastly different. Our model allows schools to endogenously choose their timing of offers whereas in their model early admission is predetermined. Moreover, we focus on understanding how schools utilize the information not only from students but also from each other. In our framework, the more competitive school is better off purposely letting the less competitive school make offers first for a sorting rationale.

Che and Koh (2016) develop a model that captures strategic plays from colleges to avoid miscoordination in the presence of students' unknown preferences. The main model of Che and Koh (2016) is simultaneous and focuses on understanding how colleges use school-specific measures such as essays to evaluate applicants in order to avoid head-on competition with other schools. In the extension, they briefly discuss the use of wait-list and sequential play. They show that there exists no symmetric equilibrium in which both schools use threshold strategies in the first round. The intuition behind the result is that while offers are turned down by some students, the next best students might not be available any more. Therefore sending offers to very top students is very risky and might not pay off for the schools. As a result, the allocation is inefficient and unfair as top students might end up receiving fewer offers or have a higher chance of not being admitted at all than lower ranked students. Our model further complements their results by showing that in a dynamic equilibrium, the less competitive school prioritizes sending offers to the seemingly low-quality student who is most likely to accept the offer. Moreover it could occur in the equilibrium that a high-quality student receives fewer offers than a low-quality student and ends up not being admitted. Both papers conclude that inefficiency cannot simply be resolved with sequential play.

In a different context, Pan (2018) investigates how the use of exploding offers affects unravelling in a two-period decentralized labor market and shows that the use of exploding offers is necessary but not sufficient for unravelling to occur. Interestingly Pan (2018) concludes that a policy on banning exploding offers tends to be supported by high quality employers but protested against by low quality ones. While our model setup is largely different from theirs, our results bear similar spirits in the sense that in both models exploding offers are preferred by less competitive offer senders. However our model speaks differently about more competitive offer senders. In their model, the information is revealed over time and exogenously, as a result high quality firms have nothing to gain by allowing low quality firms to move early. Whereas in our model, by letting the less competitive school to send out exploding offers first, the more competitive school would enjoy a filtered pool of applicants with increased expected quality.

Methodology-wise, our model is closely related to games with endogenous timing. Games with endogenous timing are most popular in the literature of industrial organization, particularly with oligopoly and duopoly models pioneered by Hamilton and Slutsky (1990) among others. The early literature has recognized that the timing of moves is an important factor to determine profit levels and studied extensively conditions under which first-mover advantage would arise. While these are important questions on their own rights, Hamilton and Slutsky (1990) argue that they did not answer the question of how the timing is decided among different firms. To answer this question, Hamilton and Slutsky (1990) extend the standard duopoly model in two ways to incorporate endogenous timing. In essence, it can be considered that the players choose the timing of their decisions in the pre-stage of the game. If the two players choose different times then a sequential subgame occurs whereas if their timing choices coincide, a simultaneous subgame arises. In one extension, firms choose the timing of their actions but need not specify the action itself. Firms are committed to the timing they choose but choose their actions only *after* the timing choices of both are announced. This variation is thus called *observable delay*. In the other extension, if a firm wishes to choose to move early then the firm also has to specify the action itself. Consequently this variation is called *action commitment*. In this variation, a firm that decides to move early does not know if the other firm is also going to move early or late. Our model is more closely related to the *action commitment* framework but with more restricted observability between periods. Hamilton and Slutsky (1990) assume that all earlier decisions are observable which we do not. Instead, we make the assumption that schools cannot directly observe each other's action. The only thing that is directly observable is whether a student exits the market. We believe this assumption bears the closest resemblance to reality. In practice, schools rarely communicate with one another over admission and usually they only directly obtain information from applicants in the case that they withdraw from the process.[3] Besides our model differs in two other major aspects. First, essentially our game features 3 players (2 schools and 1 student) that are asymmetric. As a result, conventional duopoly analysis has no bite in our model. Second, the information asymmetry is completely novel to Hamilton and Slutsky (1990) and cannot be incorporated in a straightforward manner.

Along a similar path to Hamilton and Slutsky (1990), Mailath (1993) considers a duopoly model with endogenous timing and incomplete information. The model follows the *action commitment* variation of Hamilton and Slutsky (1990). Mailath (1993) shows that only sequential equilibria survive D1 refinement (Cho

[3]While one can argue that it is a common practice for schools to make an inquiry to students about whether they hold offers from other schools - information of such is not actually verifiable.

and Kreps (1987)). Using the *observable delay* variation, Normann (2002) shows that simultaneous play i.e. a Cournot equilibrium can be sustained under a wide range of parameters. Putting endogenous timing games to application, Kempf and Graziosi (2010) introduce the framework to public economics to study the emergence of leadership in the realm of public good provision. They show that with the presence of cross-jurisdiction spillovers, neither countries would emerge as a leader if for both countries public goods are substitutes. Whereas either country might be the leader when public goods are complements. Lee and Xu (2018) investigate how environmental externality and emission tax affect timing decisions in a duopoly market.

2.3 The Model

In this section, we present the model in the case of two schools and one student[4]. Consider the two schools $c \in \{g, b\}$ in the market as a good school and a bad school, and the student with quality high or low $\theta \in \{H, L\}$. The matching process lasts for at most two periods and the schools can send their offers at either period. In every period, each school can choose to send an offer to the student with the following restrictions:

- All offers are exploding in the sense the student has to respond within the period;
- A school cannot send an offer to the student again if a previous offer from the same school has been turned down.

Given an offer, the student can either accept or reject it. Note that the student cannot accept more than one offers and once the student accepts an offer the matching procedure comes to an end. If the student accepts no offer by the

[4] Alternatively, one can think of it as a continuum of students.

end of the second period, the procedure would end with the student remain unmatched.

The preference of the student is deterministic and common knowledge. The student always prefers the good school to the bad school. We normalize the payoff of the student being admitted into the good school to 1 and denote the payoff of being admitted into the bad school as $w \in [0, 1)$.

The preferences of the schools depend on the quality of the student $\theta \in \{H, L\}$. While both schools prefer a high-quality student to a low-quality student, only the bad school receives a positive payoff from a low quality student. The good school on the other hand receives a negative payoff from a low quality student thus would prefer no admission which gives a payoff of 0. This assumption can be interpreted in two ways. For one, different payoffs could be a result of different capacity constraints. It is not uncommon to see the more competitive schools have smaller cohorts, which is especially prominent for the top private institutes vs middle-range public universities. For another, it could be due to the good school having a larger pool of applicants. Even though our model does not explicitly model the application choice of the student, it can be justified that more competitive institutes more often than not receive more applications. Either way, given the scarcity of educational resources, the opportunity cost of admitting a low-quality student would be bigger for the good school in turn a lower net payoff. We exogenize different opportunity costs as different payoffs in our model as the followings

	H	L
Good School	V	u
Bad School	V	v

Table 2.1: Admission Payoffs for Schools

where $V > v > 0 > u$. In other words, only a high quality student is acceptable to the good school whereas both qualities are acceptable to the bad school. No admission gives a payoff of 0.

While the student knows his quality perfectly, each school only receives a noisy signal over the quality $(\hat{\theta}_g, \hat{\theta}_b) \in \{h, l\}^2$. We assume that signals are conditionally independent between the two schools and follow the information structure

	h	l
H	x	$1-x$
L	$1-y$	y

Table 2.2: Conditional Probabilities of Signals $Pr(\hat{\theta}|\theta)$

where $x, y \in [0.5, 1)$. Straightforwardly the information structure satisfies MLRP and we can think of x and y as information accuracy. A larger x or y means a more accurate signal that could be from a more revealing portfolio requirement and so on. We allow information to differ across schools because for graduate student admission, most schools tailor their admission processes by asking for different materials and students usually write school-specific essays.[5] As a result, different schools would receive different information and perhaps conduct evaluation on different ways. Therefore, it is natural to assume that different schools receive different signals.[6]

To summarize, the timing of the game goes as the following:

(0a.) Nature realizes $\theta \in \{H, L\}$ with $Pr(\theta = H) = \pi$ and $Pr(\theta = L) = 1 - \pi$.

(0b.) The student learns his quality θ and each school receives a signal $(\hat{\theta}_g, \hat{\theta}_b)$ which are conditionally independent on the true

[5] In some subjects such as physics graduate students are funded directly by their advisors which means students are even more likely to prepare tailored application materials for different schools.

[6] For simplicity, we assume the information accuracy is shared between schools. We can easily extend our model to incorporate different information accuracy such that $(x_g, y_g) \neq (x_b, y_b)$ which our main results unchanged qualitatively.

quality.

(1a.) At period 1, each school simultaneous decides whether to send an offer to the student.

(1b.) The student rejects any or accepts at most one offer received.

(1c.) The game ends if the student accepts an offer else the game proceeds to period 2.

(2a). At period 2, any school that did not send an offer simultaneously decides whether to send an offer the student.

(2b.) The student rejects any or accepts at most one offer received.

(2c.) The game ends.

In this model, we make the assumption that either school can directly observe the action of the other between periods. Put differently, should the game proceed to the second period, the schools cannot distinguish between the following two events:

- The student rejected the offer from the other school.
- The student did not receive any offer.

While the observability assumption does not change our analysis qualitatively, too much observability would give the student incentives to mimic the high-quality thus deter the filtering and in turn the dynamics. With this assumption, we can also simplify the description of the school's action space: $a_c(\hat{\theta}_c) \in \{E, L, N\}$ where $c \in \{g, b\}$. In words, each school has only one offer at disposal and the school can send the offer "early (E)", "late (L)", or "never (N)" given the signal received.[7]

In the rest of the section, we first discuss the static benchmark of the game in which there is only one period. Then we provide a full characterization of the game with two periods and give conditions under which a dynamic equilibrium can be sustained. Lastly, we compare the welfare properties of the static and dynamic equilibrium outcomes.

2.3.1 Static Benchmark

With only one period, the game is straightforward. The student would accept the best available offer and reject the other one if any. The bad school would send an offer regardless of the signal received since its payoff is positive for either type. The good school send an offer to the student if and only if the expected payoff is positive given the signal, i.e.

$$VPr(\theta = H|\hat{\theta}_g) + uPr(\theta = L|\hat{\theta}_g) \geq 0$$

Under MLRP, the strategy is monotone in the sense that if the good school sends an offer to the low signal then it would also send an offer to the high signal. Henceforward, we restrict our attention to the case where the following conditions hold

$$\begin{aligned} V\frac{\pi x}{\pi x + (1-\pi)(1-y)} + u\frac{(1-\pi)(1-y)}{\pi x + (1-\pi)(1-y)} \geq 0 \\ V\frac{\pi(1-x)}{\pi(1-x) + (1-\pi)y} + u\frac{(1-\pi)y}{\pi(1-x) + (1-\pi)y} < 0 \end{aligned} \tag{2.1}$$

[7] This formulation is reminiscent to the endogenous timing duopoly game with action commitment except that in our case the school does not need to decide on the action again should it choose to wait. This is because the action L is conditional on the student not leaving the market by the end of period 1 which is all the information the school would gather between periods.

In words, the good school would only send an offer upon receiving a high signal. This assumption ensures that at least one party would be acting on their information which is what we are interested in.

In the equilibrium of the static game, the expected payoffs are given as follows

Good School	$V\pi x + u(1-\pi)(1-y)$
Bad School	$V\pi(1-x) + v(1-\pi)y$
High Quality	$x + w(1-x)$
Low Quality	$(1-y) + wy$

Table 2.3: Static Equilibrium Expected Payoffs

and the allocation of the student given the quality-signal triple $(\theta, \hat{\theta}_g, \hat{\theta}_b)$

Hhh	Hhl	Hlh	Hll	Lhh	Lhl	Llh	Lll
g	g	b	b	g	g	b	b

Table 2.4: Ex-Post Static Outcome

From the table we can see the ex-post static outcome is not assortative which is expected under information asymmetry. And it happens with probability $\pi x + (1-\pi)y$ which increases as information accuracy improves.[8]

2.3.2 The Dynamic Equilibrium

Now we start to analyse perfect Bayesian equilibria of the game with two periods. Henceforth we use equilibrium and perfect Bayesian equilibrium interchangeably. To begin, we define the strategy spaces and beliefs for the schools and for the student.

For the schools, as defined above we have $a_c(\hat{\theta}_c) \in \{E, L, N\}$ where $c \in \{g, b\}$ and $\hat{\theta}_c \in \{h, l\}$. That is, based on the signal, the school choose to send the offer

[8] Here we use the word "assortative" loosely referring to the assignment in which a high quality student goes to the good school and a low quality student to the bad school. If we assume that $u + 1 \leq v + w$ then the assortative outcome is also efficient.

early (i.e. at the first period), late (i.e. at the second period), or never. Should a school choose to postpone their offer, the choice between $\{L, N\}$ has to be optimal given belief $\beta_c(\hat{\theta}_c) \in [0, 1]$ which denotes the Bayesian posterior of a student who remains in the market to the second period being a high quality.[9]

For the student, we go backwards. The student's strategy at period two is straightforward: The student regardless of quality accepts the best offer available if any. For simplicity of exposition, we do not model that choice explicitly. Instead we denote $s_\theta(c) \in \{A, R\}$ as the action of the student with quality θ whose *best offer* at the first period comes from school $c \in \{g, b\}$.[10] Since both types prefer the good school to the bad school, we have $s_\theta(g) = A$ for $\theta \in \{H, L\}$ in any equilibrium. However how the student responds to the bad school's offer is more convoluted because there is a trade-off between waiting with the risk of being unmatched and immediately accepting the offer from the less preferred school. We define $\alpha(\theta) \in [0, 1]$ as the belief, from the student with quality θ, of receiving an offer from the good school at period two conditional on not receiving one at the first period.

To provide a full characterization of all pure strategy equilibria, we must investigate different cases of $s_\theta(b)$ which denotes the θ-quality student's response to the bad school's offer at the first period should he not receive an offer from the good school at the same period. Straightforwardly, the student rejects the offer to wait if the expected payoff of waiting is higher than the certain payoff of w.[11] We will invoke an assumption on off-path beliefs called *isolated deviation*. Simply put, the assumption dictates that it is common knowledge that any deviation from the equilibrium play is not expected to trigger more deviation

[9]Strictly speaking the action space of the schools should be defined respectively for period 1 and period 2 such as $a_c^1(\hat{\theta}_c) \in \{E, \{L, N\}\}$ and $a_c^2(\hat{\theta}_c, \beta_c(\hat{\theta}_c)) \in \{L, N\}$. However we omit the time index and the condition of a remaining student for the simplicity of exposition.

[10]Should the student received no offer at the first period then his turn of action is skipped.

[11]For simplicity of exposition, we assume the student is risk neutral. Introducing standard risk preference that is type independent does not change our analysis qualitatively.

by the same or other players. With this assumption, we can limit our attention to only monotone strategies from the student: By MLRP we have $\alpha_H \geq \alpha_L$, thus $s_L(b) = R \Rightarrow s_H(b) = R$. That is, if a low quality student waits for the good school then must a high quality student.[12]

Below we discuss three different cases of $(s_H(b), s_L(b))$. In the two more extreme cases, we have both types taking the same action, i.e. $s_H(b) = s_L(b)$. With type-independent actions, there is no real dynamics because no information is revealed between periods. In the most interesting case where $s_H(b) = R$ and $s_L(b) = A$, a low quality student might be filtered out after the first period. To utilize this filtering effect, the good school might benefit from waiting which results in a dynamic equilibrium that we show below.

Accepting the Best Early Offer

In this subsection, we discuss the case where the student regardless of type would accept the best available offer, i.e. $s_H(b) = s_L(b) = A$. For it to be optimal against belief $(\alpha(H), \alpha(L))$, it has to be the case that $\alpha(H) \leq w$ and $\alpha(L) \leq w$.

For the good school, it is never optimal to send an early offer to a low signal given the prior. Thus the good school has two choices - to defer the offer regardless of the signal or to send an early offer to a high signal.

First we consider the case in which the good school sends no early offer regardless of signal. Straightforwardly $a_g(h) = a_g(l) = L$ cannot be sustained as part of an equilibrium as it would imply $\alpha_H = \alpha_L = 1$ contradicting the case premise. Hence we are left with $(a_g(h), a_g(l)) = (L, N)$ to consider. The best

[12] Without the assumption, one can show that under some parameter range there exist equilibria in which $\alpha_H < \alpha_L$ and $s_L(b) = R, s_H(b) = A$. For example, if the bad school only sends a late offer regardless of signal and the good school sends a late offer to a high signal only. α_H and α_L are off-path thus free to play around. One can show that in this case $s_L(b) = R, s_H(b) = A$ can be sustained as part of an equilibrium under some parameter range. We consider this kind of equilibria highly unintuitive thus adopt a widely-used assumption from repeated games to rule it out.

response of the bad school would naturally be sending an early offer regardless of signal to admit the student with probability 1. However this would in turn mean that it is optimal for the good school to deviate to sending an early offer to a high signal as a positive payoff is better than a payoff of 0. As a result, there exists no equilibrium in which the good school sends no early offer.

Now consider the case where the good school send an early offer to a high signal and a late offer to a low signal, i.e. $(a_g(h), a_g(l)) = (E, L)$. Straightforwardly this cannot be sustained as an equilibrium because this implies $\alpha(H) = \alpha(L) = 1 > w$. Second, consider the case where the good school sends an early offer to a high signal but never an offer to a low signal, i.e. $(a_g(h), a_g(l)) = (E, N)$. This would imply that $\alpha(H) = \alpha(L) = 0 < w$ which is consistent with the case premise. For the bad school, its timing would be outcome-irrelevant as $(a_b(h), a_b(l)) \in \{(E, E), (E, L), (L, E), (L, L)\}$ would all give an expected payoff of $V\pi(1-x) + v(1-\pi)y$ which is identical to its static game equilibrium payoff. In turn $(a_g(h), a_g(l)) = (E, N)$ is a best response to any of the bad school's strategy. Now let us construct beliefs $\beta_c(\hat{\theta}_c)$ for the potential equilibria:

- $(a_g(h), a_g(l), a_b(h), a_b(l)) = (E, N, E, E)$: In this case, $\beta_g(l)$ is off path hence not restricted by the Bayes' rule. However we need to make sure that for the good school (E, N) is better than (E, L) which requires that $\beta_g(l) \leq \frac{\pi(1-x)}{\pi(1-x)+(1-\pi)y}$.

- $(a_g(h), a_g(l), a_b(h), a_b(l)) = (E, N, E, L)$: In this case, we just apply the Bayes' rule for $\beta_g(l)$ and $\beta_b(l)$:

$$\beta_g(l) = \beta_b(l) = Pr(\theta = H|\hat{\theta}_g = \hat{\theta}_b = l) = \frac{\pi(1-x)^2}{\pi(1-x)^2 + (1-\pi)y^2}$$

  Both schools act upon their private information hence their posteriors would converge. We can readily check that $\beta_g(l) \leq \frac{\pi(1-x)}{\pi(1-x)+(1-\pi)y}$ as

$1 - x \leq \frac{1}{2} \leq y$.

- $(a_g(h), a_g(l), a_b(h), a_b(l)) = (E, N, L, E)$: Again in this case we can just apply Bayes' rule for $\beta_g(l)$ and $\beta_b(h)$:

$$\beta_g(l) = \beta_b(h) = Pr(\theta = H|\hat{\theta}_g = l, \hat{\theta}_b = h) = \frac{\pi x(1-x)}{\pi x(1-x) + (1-\pi)y(1-y)}$$

  Similarly schools' posterior coincident as their private information is aggregated. However we can readily check that $\beta_g(l) \geq \frac{\pi(1-x)}{\pi(1-x)+(1-\pi)y}$ thus it is not straightforward to see if (E, N) is better than (E, L) for the good school. For this to be the case the following condition has to be satisfied

$$\begin{aligned} VPr(\theta = H|\hat{\theta}_g = l, \hat{\theta}_b = h) + uPr(\theta = L|\hat{\theta}_g = l, \hat{\theta}_b = h) \leq 0 \\ V\pi x(1-x) + u(1-\pi)y(1-y) \leq 0 \\ \frac{V}{-u}\frac{\pi}{1-\pi} \leq \frac{1-y}{1-x}\frac{y}{x} \end{aligned} \tag{2.2}$$

- $(a_g(h), a_g(l), a_b(h), a_b(l)) = (E, N, L, L)$: In this case there is not update for $\beta_g(l)$ as there is no information revealed from the bad school. For the bad school, Bayes' rule dictates that

$$\beta_b(h) = Pr(\theta = H|\hat{\theta}_g = l, \hat{\theta}_b = h) = \frac{\pi x(1-x)}{\pi x(1-x) + (1-\pi)y(1-y)}$$

  and

$$\beta_b(l) = Pr(\theta = H|\hat{\theta}_g = l, \hat{\theta}_b = l) = \frac{\pi(1-x)^2}{\pi(1-x)^2 + (1-\pi)y^2}$$

To summarize for the case where the student always accepts the best available early offer, we have found the following equilibria, all of which are outcome equivalent to the static game equilibrium

- $(a_g(h), a_g(l), \beta_g(l), a_b(h), a_b(l), s_H(g), s_H(b), \alpha_H, s_L(g), s_L(b), \alpha_L)$

$$= \left(E, N, \left[0, \frac{\pi(1-x)}{\pi(1-x)+(1-\pi)y}\right], E, E, A, A, 0, A, A, 0\right)$$

- $(a_g(h), a_g(l), \beta_g(l), a_b(h), a_b(l), \beta_b(l), s_H(g), s_H(b), \alpha_H, s_L(g), s_L(b), \alpha_L)$

$$= \left(E, N, \frac{\pi(1-x)^2}{\pi(1-x)^2+(1-\pi)y^2}, E, L, \frac{\pi(1-x)^2}{\pi(1-x)^2+(1-\pi)y^2}, A, A, 0, A, A, 0\right)$$

- $(a_g(h), a_g(l), \beta_g(l), a_b(h), a_b(l), \beta_b(h), s_H(g), s_H(b), \alpha_H, s_L(g), s_L(b), \alpha_L)$

$$= \left(E, N, \frac{\pi x(1-x)}{\pi x(1-x)+(1-\pi)y(1-y)}, L, E, \frac{\pi x(1-x)}{\pi x(1-x)+(1-\pi)y(1-y)}, A, A, 0, A, A, 0\right)$$

  under the condition $\frac{V}{-u}\frac{\pi}{1-\pi} \leq \frac{1-y}{1-x}\frac{y}{x}$

- $(a_g(h), a_g(l), \beta_g(l), a_b(h), a_b(l), \beta_b(h), \beta_b(l), s_H(g), s_H(b), \alpha_H, s_L(g), s_L(b), \alpha_L)$

$$= (E, N, \beta_b(h), L, L, \beta_b(h), \beta_b(l), A, A, 0, A, A, 0)$$

  where

$$\beta_b(h) = \frac{\pi(1-x)}{\pi(1-x)+(1-\pi)y}$$

  and

$$\beta_b(h) = \frac{\pi x(1-x)}{\pi x(1-x)+(1-\pi)y(1-y)}$$

$$\beta_b(l) = \frac{\pi(1-x)^2}{\pi(1-x)^2+(1-\pi)y^2}$$

While some equilibria demonstrate some dynamics in the sense that there are offers sent in both periods, they are all payoff equivalent to the static game equilibrium. Intuitively, this is because the student's action is not type-dependent rendering the filtering effect inactive.

Always Waiting for the Best Offer

In this subsection, we consider the case in which the student regardless of type always wait for the good school's offer, i.e $s_H(b) = s_L(b) = R$. For it to be optimal for beliefs (α_H, α_L), it has to be the case that $\alpha_H, \alpha_L \geq w$.

Straightforwardly it would be optimal for the bad school to defer its offer to the second period because by doing so, the bad school would have a non-zero chance of getting some student which is strictly better than a zero chance. In turn, this means that the good school would be indifferent between sending an offer to a high signal now or later and would not send an offer to a low signal with no information update. Suppose $(a_g(h), a_g(l)) = (E, N)$. The bad school's best response is to only send offers later, i.e. $(a_b(h), a_b(l)) = (L, L)$ which gives a positive expected payoff. This is because given the student's strategy, it is clear that any offer sent in period one from the bad school would be rejected thus yielding a payoff of 0. The table below gives the bad school's expected payoffs from different strategies.

	h	l
(L, L)	$\frac{V\pi x(1-x)+v(1-\pi)y(1-y)}{\pi x+(1-\pi)(1-y)}$	$\frac{V\pi(1-x)^2+v(1-\pi)y^2}{\pi(1-x)+(1-\pi)y}$
(E, L)	0	$\frac{V\pi(1-x)^2+v(1-\pi)y^2}{\pi(1-x)+(1-\pi)y}$
(L, E)	$\frac{V\pi x(1-x)+v(1-\pi)y(1-y)}{\pi x+(1-\pi)(1-y)}$	0
(E, E)	0	0

Table 2.5: Expected Payoffs of Bad School Conditional on Signal Received

As for the student's strategy to reject bad school's offer at period one, it is off the equilibrium path but sustainable as part of the equilibrium strategy with $\alpha_H, \alpha_L \geq w$. Therefore, we can conclude that the following profile consists of an equilibrium

$$(a_g(h), a_g(l), \beta_g(l), a_b(h), a_b(l), \beta_b(h), \beta_b(l), s_H(g), s_H(b), \alpha_H, s_L(g), s_L(b), \alpha_L)$$
$$= (E, N, \beta_b(h), L, L, \beta_b(h), \beta_b(l), A, R, [w,1], A, R, [w,1])$$

where
$$\beta_b(h) = \frac{\pi(1-x)}{\pi(1-x) + (1-\pi)y}$$

and
$$\beta_b(h) = \frac{\pi x(1-x)}{\pi x(1-x) + (1-\pi)y(1-y)}$$
$$\beta_b(l) = \frac{\pi(1-x)^2}{\pi(1-x)^2 + (1-\pi)y^2}$$

all pinned down by the Bayes' rule. With a similar argument, we can also show that the following profile consists of another equilibrium

$$(a_g(h), a_g(l), \beta_g(h), \beta_g(l), a_b(h), a_b(l), \beta_b(h), \beta_b(l), s_H(g), s_H(b), \alpha_H, s_L(g), s_L(b), \alpha_L)$$
$$= (L, N, \beta_g(h), \beta_b(h), L, L, \beta_b(h), \beta_b(l), A, R, [w,1], A, R, [w,1])$$

where all posteriors of the schools pinned by the Bayes' rule are identical to their priors. In this case, no early offer is sent from either school and naturally it is outcome equivalent to the static game equilibrium.

While the first equilibrium found in this subsection has some dynamics in the sense that there are offers sent in both periods, it is outcome equivalent to the static equilibrium. As a result we do not consider the dynamic "necessary" as it does not give rise to a different outcome. The second equilibrium of this case is static in action and in outcome.

Type-Dependent Responses

Now we discuss the case where the student's action is type-dependent. By MLRP and the *isolated deviation* assumption we can limit our attention to the case in which only the low quality student would accept the bad school's offer at first period whereas the high quality student would reject and wait, i.e. $s_H(b) = R,\ s_L(b) = A$ with $\alpha_L \leq w \leq \alpha_H$.

Given the strategy of the student, waiting is to the best interest of the good school. The reason is twofold. For one, a high quality student, the only type acceptable to the good school, would for sure stay to the second period. For another, there is a non-negative chance that a low quality student would leave the market early reducing the probability of the good school getting a negative payoff. Therefore, we must have $a_g(h) = L$. To further pinpoint $a_g(l) \in \{L, N\}$, consider a case where $a_g(l) = L$. This would mean that the good school would send an offer to whomever remains to the second period which creates extra incentives to reject the bad school's early offer. In fact, given this strategy of the good school, the student, regardless of type, would get an offer from the good school by rejecting the bad school's offer. Anticipating that the student regardless of type would reject the bad school's early offer. This contradicts the premise that only a high quality student rejects the bad school's early offer hence cannot be part of the equilibrium. As a result, only $a_g(l) = N$ can be sustained as part of the equilibrium. This would impose a new condition on the parameters which we will come back to after discussing the bad school's strategy.

For the bad school, it is never optimal to not send an offer, i.e. $a_b(\hat{\theta}_b) = N$ regardless of the signal received as both types yield positive payoffs for the bad school. Therefore, there are four strategies that the bad school might take as considered below.

	Hhh	Hhl	Hlh	Hll	Lhh	Lhl	Llh	Lll
(E, E)	0	0	0	0	1	1	1	1
(E, L)	0	0	0	1	1	0	1	1
(L, E)	0	0	1	0	0	1	1	1
(L, L)	0	0	1	1	0	0	1	1

Table 2.6: Student Admitted by Bad School in Case III

$a_b(h) = a_b(l) = E$: Firstly consider the strategy in which the bad school sends an early offer regardless of the signal received. This cannot be sustained as an equilibrium because if the bad school filters both types and given the strategy of the student (that only the high-quality student would reject the early bad offer), the good school would have incentives to deviate to sending a late offer to whomever remains for the second period. By the argument above, this would break our premise that only a high quality student rejects the bad school's early offer.

$a_b(h) = E, a_b(l) = L$: Secondly, suppose that the bad school sends an early offer if the signal received is high and a late offer otherwise. We need to check whether there are incentives for the bad school to deviate. Under the prescribed strategy, the bad school would get $(\theta\hat{\theta}_g\hat{\theta}_b) = (Hll)$, (Lhh), (Llh), (Lll), which is strictly dominated by sending an early offer to a low signal and a late offer to a high signal. To see that, suppose the bad school switch strategy to $a_b(h) = L$ and $a_b(l) = E$. From the table above we can see that it would get (Hlh), (Lhl), (Llh), (Lll) that is strict improvement for the bad school as (Hlh) has a bigger measure than (Hll) and (Lhl) than (Lhh). Hence this strategy can not be part of an equilibrium.

$a_b(h) = L, a_b(l) = E$: Thirdly, assume that the bad school sends an early offer to a low signal and a late offer to a high signal. Now, we check if the bad school has incentives to deviate from the current strategy, which yields the following expected payoff

$$\pi V x(1-x) + 2(1-\pi)vy(1-y) + (1-\pi)vy^2 \tag{2.3}$$

Before we already checked that this strategy strictly dominates $(a_b(h), a_b(l)) = (E, L)$ so we are left with two other strategies. Sending an early offer regardless of the signal gives

$$(1-\pi)v(1-y)^2 + 2(1-\pi)vy(1-y) + (1-\pi)vy^2 \tag{2.4}$$

and sending a late offer regardless of the signal gives

$$\pi V x(1-x) + \pi V(1-x)^2 + (1-\pi)vy(1-y) + (1-\pi)vy^2 \tag{2.5}$$

To ensure that the bad school has no incentives to deviate, the following conditions have to hold

$$\begin{aligned} \pi V x(1-x) - (1-\pi)v(1-y)^2 \geq 0 \\ \frac{\pi V}{(1-\pi)v} \geq \frac{(1-y)^2}{x(1-x)} \end{aligned} \tag{2.6}$$

and

$$\begin{aligned} (1-\pi)vy(1-y) - \pi V(1-x)^2 \geq 0 \\ \frac{\pi V}{(1-\pi)v} \leq \frac{y(1-y)}{(1-x)^2} \end{aligned} \tag{2.7}$$

Together we have

$$\frac{1-y}{x}\frac{1-y}{1-x} \leq \frac{\pi V}{(1-\pi)v} \leq \frac{y}{1-x}\frac{1-y}{1-x} \tag{2.8}$$

As shown in the figure below, it is clear that the range of parameters under which $(a_b(h), a_b(l)) = (L, E)$ is optimal is not empty. Moreover, we can see that as information accuracy, summarized in (x, y), increases, the feasible parameter range widens. This is intuitive as a more accurate signal is more valuable thus making a good use of the information becomes more important.

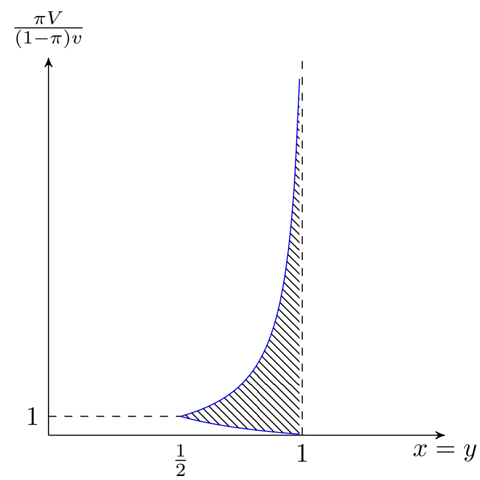

Now we need to check if the prescribed strategies of the student and of the good school would be consistent. For the student, it is straightforward that the following condition has to hold for the low quality

$$w \geq \alpha_L = 1 - y \tag{2.9}$$

and for the high quality

$$w \leq \alpha_H = x \tag{2.10}$$

where the left-hand side is the certain payoff of accepting the bad school's early offer and the right-hand side is the expected payoff of rejecting.

As for the good school, it can be readily checked that a late offer strictly dominates an early offer regardless of the signal as we discussed at the beginning of this subsection. Now we can derive the condition under which the good school would never send an offer to the low signal, i.e. $a_g(l) = N$:

$$\begin{aligned} \pi V x(1-x) + \pi V(1-x)^2 + (1-\pi)uy(1-y) &\le 0 \\ \Leftrightarrow \pi V(1-x) + (1-\pi)uy(1-y) &\le 0 \\ \Leftrightarrow \frac{\pi V}{-(1-\pi)u} &\le \frac{y(1-y)}{1-x} \end{aligned} \tag{2.11}$$

To summarize, the profile consists of a dynamic equilibrium

$$\begin{aligned} &(a_g(h), a_g(l), \beta_g(h), \beta_g(l), a_b(h), a_b(l), \beta_b(h), s_H(g), s_H(b), \alpha_H, s_L(g), s_L(b), \alpha_H) \\ = &(L, N, \beta_g(h), \beta_g(l), L, E, \beta_g(h), \beta_g(l), A, R, x, A, A, 1-y) \end{aligned}$$

where

$$\beta_g(h) = \frac{\pi x}{\pi x + (1-\pi)y^2}$$

$$\beta_g(l) = \frac{\pi(1-x)}{\pi(1-x) + (1-\pi)y(1-y)}$$

and $\beta_b(h)$ identical to the prior, if the following conditions are satisfied

$$\begin{aligned} 1-y \le w &\le x \\ \frac{\pi V}{-(1-\pi)u} &\le \frac{y(1-y)}{1-x} \\ \frac{1-y}{1-x}\frac{1-y}{x} \le \frac{\pi V}{(1-\pi)v} &\le \frac{1-y}{1-x}\frac{y}{1-x} \end{aligned} \tag{2.12}$$

and in this dynamic equilibrium the good school only sends an offer to a high signal at the second period and the bad school sends an offer to a low signal at the first period and to a high signal at the second period. At the first period, a high quality student would reject the bad school's offer whereas a low quality student would accept.

$a_b(h) = a_b(l) = L$**:** Lastly consider the strategy in which the bad school sends a late offer regardless of the signal received. This is optimal for the bad school under the following conditions:

$$\begin{aligned} \pi V(1-x) - (1-\pi)v(1-y) \geq 0 \\ \frac{\pi V}{(1-\pi)v} \geq \frac{1-y}{1-x} \end{aligned} \tag{2.13}$$

when $a_b(h) = a_b(l) = L$ generates weakly higher payoff than $a_b(h) = a_b(l) = E$ and

$$\begin{aligned} \pi V(1-x)^2 - (1-\pi)vy(1-y) \geq 0 \\ \frac{\pi V}{(1-\pi)v} \geq \frac{1-y}{1-x}\frac{y}{1-x} \end{aligned} \tag{2.14}$$

when $a_b(h) = a_b(l) = L$ generates weakly higher payoff than $(a_b(h), a_b(l)) = (L, E)$.[13] Clearly we can see that $\frac{\pi V}{(1-\pi)v} \geq \frac{1-y}{1-x}\frac{y}{1-x} \geq \frac{1-y}{1-x}$ since $y \geq \frac{1}{2} \geq 1-x$.

With the bad school filtering neither type thus providing no information update for the good school, the good school would be indifferent in sending an early offer to a high signal or a late offer. With no additional condition needed, we can assert that the good school would never send an offer to a low signal i.e. $a_g(l) = L$. Given the strategies of the schools, α_H and α_L are off path. However the *isolated deviation* assumption dictates that when $(a_g(h), a_g(l)) = (E, N)$, $\alpha_H = \alpha_L = 0$ and when $(a_g(h), a_g(l)) = (L, N)$, $\alpha_H \geq \alpha_L$. To be consistent with the case premise that $s_H(b) = R$ and $s_L(b) = A$, we construct off-path beliefs such that $\alpha_H \geq w \geq \alpha_L$.

As a result, when $\frac{V}{v}\frac{\pi}{1-\pi} \geq \frac{1-y}{1-x}\frac{y}{1-x}$ is satisfied, the following profile consists

[13] Remember that $(a_b(h), a_b(l)) = (L, E)$ strictly dominates $(a_b(h), a_b(l)) = (E, L)$. Thus generating higher payoff than $(a_b(h), a_b(l)) = (L, E)$ implies generating higher payoff than $(a_b(h), a_b(l)) = (E, L)$.

of an equilibrium

$$(a_g(h), a_g(l), \beta_g(h), \beta_g(l), a_b(h), a_b(l), \beta_b(h), \beta_b(l), s_H(g), s_H(b), \alpha_H, s_L(g), s_L(b), \alpha_L)$$
$$= (L, N, \beta_g(h), \beta_g(l), L, L, \beta_b(h), \beta_b(l), A, R, [w, 1], A, A, [0, w])$$

where $\beta_c(\hat{\theta}_c)$ for $c \in \{g, b\}$ and $\hat{\theta}_c) \in \{h, l\}$ are identical the priors. Again this equilibrium is payoff equivalent to that of a static game.

To conclude our characterization of equilibria for the two-period game, there are multiple profiles that could be sustained as an equilibrium. However, most of them give rise to an outcome that is equivalent to that of the static game. In the unique dynamic equilibrium, we have both schools endogenously choose to send out offers at different periods: the bad school sends an early offer to a low signal and a late offer to a high signal while the good school only sends a late offer to a high signal. A low quality student accepts the best offer available and leaves the market immediately whereas a high quality student waits.

Proposition 6. *There exists a dynamic equilibrium under the following conditions:*

$$\begin{aligned} 1 - y \leq w \leq x \\ \frac{\pi V}{-(1-\pi)u} \leq \frac{y(1-y)}{1-x} \\ \frac{1-y}{1-x}\frac{1-y}{x} \leq \frac{\pi V}{(1-\pi)v} \leq \frac{1-y}{1-x}\frac{y}{1-x} \end{aligned} \tag{2.15}$$

In this equilibrium, the bad school sends an early to a low signal and a late offer to a high signal while the good school only sends a late offer to a high signal. A low quality student would accept the best offer available at the first period whereas a high quality student would wait for an offer from the good school.

2.3.3 Welfare

In this section, we investigate the two possible outcomes from the two-period game. We refer the outcome equivalent to that of the static game as the *static outcome* and the other *dynamic outcome*.

Under complete information, the market would achieve assortative outcome in which a high quality student goes to the good school and a low quality student to the bad school. However the assortative outcome cannot be achieved with information asymmetry in our model as summarized in the table below.

	Hhh	Hhl	Hlh	Hll	Lhh	Lhl	Llh	Lll
Assortative	g	g	g	g	b	b	b	b
Dynamic	g	g	b	$\emptyset$	g	b	b	b
Static	g	g	b	b	g	g	b	b

Table 2.7: Allocations of the Dynamic and the Static Equilibria

Now let us investigate whether the dynamic game helps or hurts the schools and the student. For the good school, the expected payoff in the dynamic equilibrium is given by

$$\pi V x + (1-\pi)u(1-y)^2 \tag{2.16}$$

which is higher than the expected payoff from the static outcome

$$\pi V x + (1-\pi)u(1-y) \tag{2.17}$$

. For the bad school, the expected payoff from the dynamic outcome is

$$\pi V x(1-x) + (1-\pi)v(1-y)(1+y) \tag{2.18}$$

as opposed to the expected payoff from the static outcome

$$\pi V(1-x) + (1-\pi)v(1-y) \tag{2.19}$$

. The bad school is better off in the dynamic equilibrium when

$$\frac{\pi V}{(1-\pi)v} \leq \frac{1-y}{1-x}\frac{y}{1-x} \tag{2.20}$$

. Note that this is one of the necessary conditions for the existence of the dynamic equilibrium. In other words, whenever the dynamic equilibrium is possible, the bad school is better off in the dynamic equilibrium. Intuitively, if either school is made worse off with a dynamic strategy profile, then they could unilaterally deviate to achieve the static outcome. Therefore, the dynamic equilibrium can only be sustained when both schools are weakly better off.

As for the student, the expected payoff for a high quality student in the static outcome is $x + w(1-x)$ and for a low quality student $1 - y + wy$. In the dynamic outcome, the expected payoff for a high quality student is $x+wx(1-x)$ which is lower since there is a positive probability that a high quality student would end up unmatched with signal combination $(\hat{\theta}_g, \hat{\theta}_b) = (h, l)$. Furthermore, there is no signal combination that improves a high quality student's payoff in the dynamic outcome. For a low quality student, the expected payoff from the dynamic outcome is $(1-y)[(1-y) + wy] + wy$ which is also lower than that in the static outcome. This is due to the filtering effect that allows the good school to decrease the chance of admitting a low quality student in turn reducing the expected payoff of such a student. A low quality student with signal combination $(\hat{\theta}_g, \hat{\theta}_b) = (h, l)$ could get into the good school in static equilibria but not so in the dynamic equilibrium.

To summarize, both schools benefit from the dynamic outcome while both types of student suffer.

Theorem 1. *In expectation, both schools are made better off in the dynamic equilibrium whereas the student is made worse off.*

2.4 Discussion and Conclusion

2.4.1 Observability

Throughout the paper, we have assumed that the schools cannot observe each other's action between periods. In addition, they cannot observe if a student rejects an offer from the other school. Between periods, the schools can only observe a rejection of their own offer or a student withdrawing (by leaving the market). Admittedly, this is very restricted observability especially compared to previous models. As discussed earlier, there are two variations of endogenous timing games pioneered by Hamilton and Slutsky (1990), *observable delay* and *action commitment*. While our model is closely related to the action commitment framework, the degree of observability differs in our model (among other things). In Hamilton and Slutsky (1990), a late mover can observe the action from an early mover. This assumption is justifiable in the context of oligopoly where first mover firms, by releasing products into the market, reveal their actions. In addition, there are interests in understanding the role of leadership in industrial organization hence this assumption is natural. Nonetheless, we have argued that in the context of school admission which is what we focus in this paper, the assumption is less intuitive as schools rarely directly communicate with one another about their admission decisions let along making public announcements. While it is not unusual for interviewers to ask interviewees about their received offers, this information is almost not verifiable and hardly verified.

A complete relaxation of the observability assumption would upset the dynamic equilibrium we discussed. To see why, suppose the schools can observe actions taken in the first period before deciding what to do at the second period. In the dynamic equilibrium, only a high quality student would reject the bad school's early offer. Therefore, the good school would always wait and send an offer to the student who rejects the bad school's early offer. However this in turn gives a low quality student strong incentives to mimic a high quality student which would guarantee the student a late offer from the good school. In other words, a low quality student could imitate the high-quality student at no cost, rendering the filter ineffective. That said, some level of observability can be incorporated without changing our results qualitatively. For instance, consider an environment where the schools could observe a rejection with probability $p \in (0, 1)$ and with probability $1 - p$ nothing. As long as the p is sufficiently small such that a high quality student is strictly more likely to reject the bad school's early offer, our analysis still go through with mixed strategies.

2.4.2 Less Informed Student

Another assumption that we have kept throughout the paper is that the student knows the quality perfectly. This assumption can be relaxed with our results unchanged qualitatively. For example, consider the following environment in which the student receives a conditionally independent noisy signal over quality $z \in \{h, l\}$. As long as the information does not violate the monotone likelihood ratio property such that the student uses a monotone strategy in the sense that the student with a better signal is more like to wait for the best offer, our results would still hold.

2.4.3 Conclusion

To conclude, in this paper we investigate how incomplete information affects the timing of admission offers with a game theoretic model. In a two-period game, we show that under a non-empty set of parameters, there exists a dynamic equilibrium in which the less competitive school sends an early offer to a low signal student whereas the more competitive school sends a late offer to a high signal student. Only a high quality student would rather reject the early offer from the less competitive school and wait for a better offer whereas a low quality student would accept the best available early offer and exit the market. In this equilibrium, the more competitive school utilizes the filtering effect to increase its probability of admitting a high quality student later. Our result is consistent with the observation that different schools send out offers at different periods during the admission window and the widely use of exploding offers. We also provide a novel framework to examine the choice of offer timing as an endogenous variable - as far as we know this is the first paper that introduces the endogenous timing game to the school admission environment. Lastly we compare the dynamic equilibrium to the benchmark static equilibrium. For payoffs, neither equilibrium Pareto-dominates the other and both are worse off than an assortative equilibrium with compete information. We show that for the schools, the dynamic equilibrium is superior as the more competitive school decreases the probability of admitting a low quality student and the less competitive increases its chance of admission. However, this is not without a cost. For the student, the dynamic equilibrium makes him worse off in expectation. This is because he has a lower chance of being admitted to the more competitive school. As a result, our model does not make a clear suggestion for policymakers about the employment of exploding offers in the school admission context.

2.4.4 Future Avenues

While our model is relatively straightforward with two schools and two types, we believe it is a solid first step towards understanding school admission with endogenous timing and exploding offers. For the same reason, there are various interesting directions for future research. For example, instead of binary type one can consider continuous type or non-binary discrete type. However this is rather challenging technically because it might not be optimal for the schools to use a threshold strategy without additional assumptions on the information structure.

Appendix A

Appendix to Chapter 1

A.1 Existence and Uniqueness

In this Appendix, we prove Proposition 1. We show that the constructed equilibrium exists and is unique under D1 a la Cho and Sobel (1990). First we show that in any D1 equilibrium, payoffs of both types have to be $\bar{u}^*(b)$. Then we prove that with the constructed strategies, no deviations, to on- or off-equilibrium decisions, are profitable. The uniqueness comes for free since the principal only employs pure strategies and the solution to RP is unique (Mailath (1987)).

For preliminaries, we need the following claims.

Claim 1 (Monotonicity)**.** *For any $a, a' \in D$ such that $a < a'$, if*

$$-(a' - \theta)^2 + \Delta V_2^0(\mu(a'), \mu(a)) \geq -(a - \theta)^2$$

then

$$-(a' - \theta - B)^2 + \Delta V_2^B(\mu(a'), \mu(a)) > -(a - \theta - B)^2$$

Proof. The proof is by contradiction. Suppose

$$-(a' - \theta - B)^2 + \Delta V_2^B(\mu(a'), \mu(a)) \leq -(a - \theta - B)^2$$

Rearrange the two inequalities

$$\Delta V_2^B(\mu(a'), \mu(a)) \leq (a' - \theta - B)^2 - (a - \theta - B)^2$$
$$\Delta V_2^0(\mu(a'), \mu(a)) \geq (a' - \theta)^2 - (a - \theta)^2$$

and use Lemma 3, we have

$$(a' - \theta)^2 - (a - \theta)^2 \leq (a' - \theta - B)^2 - (a - \theta - B)^2 \Leftrightarrow$$
$$(a' - \theta)^2 - (a' - \theta - B)^2 \leq (a - \theta)^2 - (a - \theta - B)^2 \Leftrightarrow$$
$$-B^2 + 2B(a' - \theta) \leq -B^2 + 2B(a - \theta) \Leftrightarrow$$
$$a' \leq a$$

a contradiction.

■

Claim 2. *Fix an equilibrium in which the unbiased type takes a with positive probability and receives equilibrium payoff $u^*(0) = u(0, a, \mu)$. If $a' < a$, then*

(a) B takes decision a' with probability 0 in equilibrium;

(b) Any D1 equilibrium can be supported by beliefs such that $\mu(a') = 0$ for all $a' < a$.

Proof.

(a) Using sequential rationality again, we have $u^*(0) = u(0, a, \mu) \geq u(0, a', \mu')$. Lemma 1 implies that $u(B, a, \mu) > u(B, a', \mu')$ and sequential rationality requires that a taken by B with probability 0.

(b) If $u^*(0) \geq u(0, a', \mu')$ for any $\mu' \in [0, 1]$, we need to show $u^*(B) \geq u(B, a', \mu')$. Since $u^*(0) = u(0, a, \mu)$ and $u^*(B) \geq u(B, a, \mu)$ by sequential rationality, Claim 1 indicates that

$$u^*(0) = u(0, a, \mu) \geq u(0, a', \mu') \Rightarrow u^*(B) \geq u(B, a, \mu) > u(B, a', \mu')$$

for any $\mu' \in [0, 1]$. In this case, D1 puts no restriction on $\mu(a')$. Consider any D1 equilibria and change $\mu(a')$ to 0; the equilibria will not be upset as $u^*(0) \geq u(0, a', 0)$ and $u^*(B) \geq u(B, a', 0)$.

Similarly we can show that if $u^*(B) \leq u(B, a', \mu')$ for some $\mu' \in [0, 1]$ then $u^*(0) < u(0, a', \mu')$. In this case, D1 dictates that $\mu(a') = 0$.

■

The last important piece we need before construction is that only pooling over $\underline{d}$ survives D1.

Claim 3 (Pooling at Bottom)**.** *In any D1 equilibrium, the lowest delegated decision $\underline{d}$ is the only possible decision taken by both types with positive probabilities.*

Proof. Fix any D1 equilibrium and denote the equilibrium payoff for type b as $u^*(b)$. Assume that there exists $\underline{d} < a \in D$ over which the two types pool with positive probability. Then $u^*(0) < u(0, a, 0)$. Claim 2(b) suggests that the equilibrium can be supported by $\mu(a') = 0$ for all $a' < a$. By sequential rationality, it must be $u^*(0) \geq u(0, a', 0)$ for any $a' < a$. Together we have for any $a' < a$

$$u(0, a, 0) > u^*(0) \geq u(0, a', 0)$$

contradicting that u is continuous in the second argument.

■

To achieve the first step, fix any D1 equilibrium with equilibrium payoff to type b denoted by $u^*(b)$.

Claim 4.

(a) If the biased type pools at $\underline{d}$ *with positive probability, then the unbiased type takes* $\underline{d}$ *with probability* 1.

(b) The biased type separates with probability 1 *if and only if* $u^*(B) \geq u(B, \underline{d}, 0)$.

(c) If there is no pooling then

$$u^*(0) = \tilde{u}(0) \equiv \max_{a \in D}\{u(0, a, 0) : u^*(B) \geq u(B, a, 0)\}$$
$$u^*(B) = \bar{u}(B)$$

If there is pooling with positive probability then $u^*(B) \geq \bar{u}(B)$ *and if there is separation with positive probability then* $u^*(B) = \bar{u}(B)$.

Proof.

(a) Suppose not and the unbiased type takes $a > \underline{d}$ with positive probability. According to Claim 2(a), the biased type takes $\underline{d} < a$ with probability 0, a contradiction.

(b) If the biased type pools with positive probability, then from (a), the unbiased type takes $\underline{d}$ with probability 1. This implies that $u^*(B) = u(B, \underline{d}, \lambda) < u(B, \underline{d}, 0)$ where $\lambda \in (0, p)$. As for the other direction, note that sequential rationality requires that $u^*(B) \geq u(B, \underline{d}, \mu(\underline{d}))$. Also, Claim 2(a) implies that the biased type takes $\underline{d}$ with probability 0 if he separates with probability 1. Therefore, $\mu(\underline{d}) = 0$ when the unbiased type takes $\underline{d}$ with positive probability. Otherwise Claim 2(b) suggests that any D1 can be supported with $\mu(\underline{d}) = 0$. Together, we have $u^*(B) \geq u(B, \underline{d}, 0)$ and $u^*(0) = u(0, a, 0)$.

(c) It types separate with probability 1, then $\exists a \neq a' \in D$ such that $u^*(0) = u(0, a, 0)$ and $u^*(B) = u(B, a', 1) \geq u(B, a, 0)$. By definition, $u^*(B) \leq \bar{u}(B)$ and $u^*(0) \leq \tilde{u}(0)$. The other direction comes from sequential rationality.

If types pool with positive probability and the inequality does not hold, i.e. $u^*(B) < \bar{u}(B)$. Then there exists some $d \in D$ such that $u(B, a, 1) > u^*(B)$. Suppose the unbiased type takes a with probability $\gamma \in [0, 1]$, then equilibrium belief on a is weakly smaller than 1: $u(B, a, \mu(a)) \geq u(B, a, 1) > u^*(B)$, violating sequential rationality. The last part follows immediately as types play mixed strategies only when they are indifferent.

■

Lemma A1. *All D1 equilibria are payoff-equivalent for both types of agent. In particular, the expected equilibrium payoff for b is $\bar{u}^*(b)$.*

Proof.

(i) If types separate with probability 1, by Claim 4(b), $\bar{u}(B) = u^*(B) \geq u(B, \underline{d}, 0)$. By definition $\bar{u}^*(B) = \bar{u}(B) = u^*(B)$.

For the unbiased type, $u^*(0) = \tilde{u}(0) = \bar{u}(0)$. The first equality again comes from Claim 4(c) and the second from $u^*(B) = \bar{u}^*(B)$. Since the biased type takes $\underline{d}$ with probability 0, any D1 can be supported by $\mu(\underline{d}) = 0$: $u^*(B) = \bar{u}(B) \geq u(B, \underline{d}, 0)$ which gives $\bar{u}^*(0) = \bar{u}(0) = u^*(0)$.

(ii) If the biased type pools with positive probability that is strictly smaller than 1, then $u^*(B) = \bar{u}(B) = u(B, \underline{d}, \mu(\underline{d})) = \bar{u}(B) = \bar{u}^*(B)$. The first equality comes from Claim 4(c), the second from mixing only if indifferent, and the last from the definition. As for the unbiased type,

$u^*(0) = u(0, \underline{d}, \mu(\underline{d})) = \bar{u}^*(0)$ by Claim 4(a) and the definition as well as $u^*(B) = \bar{u}^*(B)$.

(iii) If the biased type pools with probability 1, then $u^*(B) = u(B, \underline{d}, p) \geq \bar{u}(B)$ by Claim 4(c). By definition $\bar{u}^*(B) = u(B, \underline{d}, p) = u^*(B)$ and $u^*(0) = u(0, \underline{d}, p) = \bar{u}^*(0)$.

■

Claim 5. *Any D1 equilibria can be supported by $\mu(B|a) = 1$ for all $a > m^*(0)$.*

Proof. First we show that $\forall a > m^*(0)$, if $\exists \mu \in [0, 1]$ such that $\bar{u}^*(0) \leq u(0, a, \mu)$ then $\bar{u}^*(B) < u(B, a, \mu)$. By contradiction, suppose $\exists a > m^*(0),\ \mu \in [0, 1]$ such that $\bar{u}^*(0) \leq u(0, a, \mu)$ but $\bar{u}^*(B) \geq u(B, a, \mu)$. The contrapositive of Claim 1 suggests that $\bar{u}^*(0) > u(0, a, \mu)$, contradicting the assumption.

Then D1 dictates $\mu(a) = 1$ if there exists $a > m^*(0)$ and $\mu \in [0, 1]$ such that $\bar{u}^*(B) < u(B, a, \mu)$. Otherwise, D1 puts no restriction on $\mu(a)$. Take any D1 equilibria and change $\mu(a)$ to 1. It is straightforward to see that the original equilibrium will not be upset by such a change.

■

Lemma A2. *Given any $D \subseteq \mathbb{R}$, there is a unique D1 equilibrium at each $\theta \in [0, 1]$ in which agent of type b employs $\bar{m}^*(\cdot|b)$ and receives equilibrium payoff $\bar{u}^*(b)$.*

Proof. We are to show that both types respond optimally with $\bar{m}^*$ given $\bar{\mu}^*$ where $\bar{\mu}^*$ is specified by Bayes' rule whenever possible. Otherwise $\bar{\mu}^*(a) = 0$ for all $a < m^*(0)$ and $\bar{\mu}^*(a) = 1$ for all $a > m^*(0)$. By Claims 2(b) and 5, $\bar{\mu}^*$ passes D1. To best response, the following incentive compatibility constraints have to hold

$$\bar{u}^*(B) \geq u(B, a, \bar{\mu}^*(d)) \quad \forall a \in D \tag{IC-B}$$

$$\bar{u}^*(0) \geq u(0, a, \bar{\mu}^*(d)) \quad \forall a \in D \tag{IC-0}$$

(i) First we show that no on-path deviation generates strictly higher payoffs for either types. If $\bar{u}^*(B) = \bar{u}(B)$ (separation with probability 1), then by construction $\bar{u}^*(B) \geq u(B, m^*(0), 0)$ and IC-B holds. Otherwise $\bar{u}^*(B) = u(B, \underline{d}, \bar{\mu}^*(\underline{d}))$ and the unbiased type takes $\underline{d}$ with probability 1. As a result, there is no on-path deviation and both IC-B and IC-0 trivially hold.

What is left to show is IC-0 under separation, i.e. $\bar{u}^*(0) = u(0, m^*(0), 0) \geq u(0, m^*(B), 1) = u(0, y_D^B(\theta), 1)$. By continuity, there exists $\beta \in (0, 1]$ such that if $a' = \beta\underline{d} + (1 - \beta)y_D^B(\theta)$ then $\bar{u}^*(B) = u(B, a', 0)$. Moreover a' is feasible for RP which implies $\bar{u}^*(0) \geq u(0, a', 0)$. Suppose IC-0 does not hold and $u(0, y_D^B(\theta), 1) > \bar{u}^*(0) \geq u(0, a', 0)$. Since $y_D^B(\theta) > a'$, Claim 1 dictates that $u(B, y_D^B(\theta), 1) > u(B, a', 0)$ contradicting the premise.

(ii) To complete the proof, we show that no off-path deviation is profitable either. If $\bar{u}^*(0) = u(B, \underline{d}, \mu^*(\underline{d}))$ then $\bar{u}^*(B) > \bar{u}(B) \geq u(B, a, 1)$ $\forall a > \underline{d}$. The first inequality follows from construction of $\bar{u}^*(B)$ and the second definition of $\bar{u}(B)$. The contrapositive of Claim 1 implies that $u(0, a, 1) < u(0, \underline{d}, \bar{\mu}^*(\underline{d}) = \bar{u}^*(0)$. IC-B and IC-0 are met in this case.

If $m^*(0) < m^*(B)$, RP suggests that $u(B, m^*(B), 1) = u(B, m^*(0), 0) = \bar{u}^*(B)$. Claim 1 suggests $u(B, a, 0) < \bar{u}^*(B)$ for any $a < m^*(0)$. Thus IC-B holds for any $a < m^*(0)$. For $a > m^*(0)$, $\bar{u}^*(B) = \bar{u}(B) \geq u(B, a, 1)$ which assures IC-B.

For the unbiased type, $\bar{u}^*(0) = \bar{u}(0) \geq u(0, a, 0)$ for any $a < m^*(0)$; hence IC-0 is satisfied for $a < m^*(0)$. For any $a > m^*(0)$, we have argued that $u(B, m^*(0), 0) \geq u(B, a, 1)$. By the contrapositive of Claim 1, $u(0, m^*(0), 0) > u(0, a, 1)$.

Last but not least, the uniqueness comes for free since the principal only employs pure strategies and the solution to RP is unique (Mailath (1987)).

■

A.2 Other Omitted Proofs

Proof of Lemma 1

Proof. Firstly, it is clear that $\bar{d}_2^* \in [0, 1]$. To see why, consider any $\bar{d}_2 < 0$, then both types would always choose $\bar{d}_2$ regardless of θ_2. In this case, the principal is better off by raising $\bar{d}_2$ to 0. As for $\bar{d} > 1$, it is dominated by setting it to 1.

Now look at the lower bound $\underline{d}_2$. Notice that $\bar{d}_2 \geq \underline{d}_2 > B$ is never optimal. Decreasing $\underline{d}_2$ till B induces both types to take lower actions at lower states ($\theta \leq \max\{0, \underline{d}_2 - B\}$ for biased type), which benefits the principal. Thus $\underline{d}_2^* \leq B$. Further lowering it would not induce any changes from the biased type. Nevertheless it facilitates the unbiased type to take lower actions at lower states, which again makes the principal better off. Reducing $\underline{d}_2^*$ below 0 does not make any difference as in the last period, no actions lower than 0 will be implemented by either type. Thus $\underline{d}_2^* = 0$.[1]

[1]Figure A.1 is plotted with $\bar{d}_2 > B$ but the same argument applies if $\bar{d}_2 \leq B$; in that case the biased type will always implement $\bar{d}_2$ regardless of the realized state.

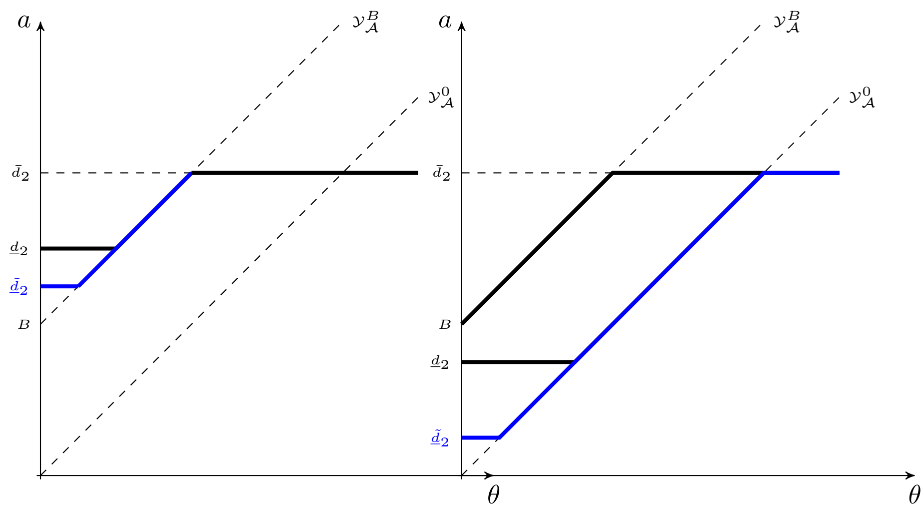

(a) Lowering $\underline{d}_2$ to B (b) Lowering $\underline{d}_2$ to 0

Figure A.1: Lower bound for the last period

As a result of the previous arguments, the only choice variable left for the principal's last period maximization 1.3 is $\bar{d}_2$. Rewrite it with $\underline{d}_2 = 0$ and $\bar{d}_2 \leq 1$

$$\max_{\bar{d}_2 \in [0,1]} -\mu \left[\int_0^{\max\{0,\bar{d}_2 - B\}} B^2 + \int_{\max\{0,\bar{d}_2 - B\}}^1 (\bar{d}_2 - \theta_2)^2 \right] f(\theta) d\theta_2 - (1-\mu) \int_{\bar{d}_2}^1 (\bar{d}_2 - \theta_2)^2 f(\theta) d\theta_2$$

To complete the proof, we are to show that the above function is submodular in (d, μ) which then implies that $\bar{d}_2^*$ is decreasing in μ. In other words, if we can show that $\forall \mu$ and $d' \geq d$, $U(d', \mu) - U(d, \mu)$ is decreasing μ, then the proof is complete. To see that, we have

$$U(d', \mu) - U(d, \mu) = \textit{constant in } \mu$$
$$+ \mu \left[- \int_{\max\{0,d-B\}}^{\max\{0,d'-B\}} B^2 - \int_{\max\{0,d'-B\}}^1 (d' - \theta)^2 + \int_{d'}^1 (d' - \theta)^2 + \int_{\max\{0,d-B\}}^1 (d - \theta)^2 - \int_d^1 (d - \theta)^2 \right] f(\theta) d\theta$$

Now we want to show that the coefficient of μ is negative for any $d' \geq d$

$$\begin{aligned}
&\left[-\int_{\max\{0,d-B\}}^{\max\{0,d'-B\}} B^2 - \int_{\max\{0,d'-B\}}^{1} (d'-\theta)^2 + \int_{d'}^{1} (d'-\theta)^2 + \int_{\max\{0,d-B\}}^{1} (d-\theta)^2 - \int_{d}^{1} (d-\theta)^2\right] f(\theta)d\theta \\
&\leq \left[-\int_{\max\{0,d'-B\}}^{1} (d'-\theta)^2 + \int_{d'}^{1} (d'-\theta)^2 + \int_{\max\{0,d'-B\}}^{1} (d-\theta)^2 - \int_{d'}^{1} (d-\theta)^2\right] f(\theta)d\theta \\
&= \left[\int_{d'}^{1} (d'-d)(d'+d-2\theta) + \int_{\max\{0,d'-B\}}^{1} (d-d')(d+d'-2\theta)\right] f(\theta)d\theta \\
&= \left[\int_{d'}^{1} (d'-d)(d'+d-2\theta) + \int_{d'}^{1} (d-d')(d+d'-2\theta) + \int_{\max\{0,d'-B\}}^{d'} (d-d')(d+d'-2\theta)\right] f(\theta)d\theta \\
&= \left[\int_{\max\{0,d'-B\}}^{d'} (d-d')(d+d'-2\theta)\right] f(\theta)d\theta \leq 0
\end{aligned}$$

Apparently, the inequality is strict if $d' > d$ and the proof is complete.

■

Proof of Lemma 2

Proof. For expositional simplicity, define $V_2^P(\mu) \equiv \max_{d\in\mathcal{D}} V_2^P(\mu, d)$ where d here denotes the upper bound of the delegation set.

1. To show that the value function for principal is decreasing in μ, consider $\mu' > \mu$. Straightforwardly, we have

$$V_2^P(\mu) \equiv V_2^P(\mu, d^*(\mu)) \geq V_2^P(\mu, d^*(\mu')) \geq V_2^P(\mu', d^*(\mu')) = V_2^P(\mu')$$

   which completes the proof for this part.

2. To show that the value function is convex, consider $\forall t \in [0, 1]$ and $p, q \in$

$[0,1]$. We have

$$\begin{aligned}V(tp+(1-t)q) =&[tp+(1-t)q]V(1,d^*(tp+(1-t)q))+[1-(tp+(1-t)q)]V(0,d^*(tp+(1-t)q))\\ =&t[pV(1,d^*(tp+(1-t)q))+(1-p)V(0,d^*(tp+(1-t)q))]\\ &+(1-t)[qV(1,d^*(tp+(1-t)q))+(1-q)V(0,d^*(tp+(1-t)q))]\\ =&tV(p,d^*(tp+(1-t)q))+(1-t)V(q,d^*(tp+(1-t)q))\\ \leq&tV(p)+(1-t)V(q)\end{aligned}$$

Thus the proof is complete.

■

Proof of Lemma 3

Proof. From Lemma 1, we know that the upper bound is decreasing in posterior. Hence under $\mu' < \mu$, $d' \geq d$. If $d' = d$, the statement is trivially true. Now consider $d' > d$. We are to show that at every state $\theta \in [0,1]$,

$$u(B, y^B_{d'}, \theta) - u(B, y^B_d, \theta) \geq u(0, y^0_{d'}, \theta) - u(0, y^0_d, \theta)$$

and the inequality is strict for θ such that $u(0, y^0_{d'}, \theta) - u(0, y^0_d, \theta) > 0$. Then the monotonicity of integral will conclude the proof.

First consider the case $u(0, y^0_{d'}, \theta) = u(0, y^0_d, \theta)$. Either $d' > d \geq \theta + B$ in which case $u(B, y^B_{d'}, \theta) = u(B, y^B_d, \theta)$ or at least one of the two bounds is smaller than $\theta + B$ in which case $u(B, y^B_{d'}, \theta) - u(B, y^B_d, \theta) > 0$.

Now suppose $u(0, y^0_{d'}, \theta) > u(0, y^0_d, \theta)$. This implies that $d < \theta$. Then $y^0_{d'} - y^0_d = \min\{\theta, d'\} - d$. Similarly $y^B_{d'} - y^B_d = \min\{\theta + B, d'\} - d$. Hence we have $y^B_{d'} - y^B_d \geq y^0_{d'} - y^0_d$ with the same starting point (i.e. d). Since the payoff function u is convex with $u(B, \cdot, \theta)$ being right of $u(0, \cdot, \theta)$, we have

$$u(B, y_{d'}^B, \theta) - u(B, y_d^B, \theta) \geq u(0, y_{d'}^0, \theta) - u(0, y_d^0, \theta)$$

■

Proof of Lemma 4

Proof. As a preliminary step, we show that

$$(\underline{d} - \theta - B)^2 - (y_D^B(\theta) - \theta - B)^2 \tag{A.2}$$

increases in θ. Take derivative at differentiable points

$$\frac{\partial[(\underline{d} - \theta - B)^2 - (y_D^B(\theta) - \theta - B)^2]}{\partial\theta} = 2\left[y_D^B(\theta) - \underline{d} - \frac{\partial y_D^B(\theta)}{\partial\theta}\left(y_D^B(\theta) - \theta - B\right)\right]$$

If $y_D^B(\theta) \in \{\underline{d}, \bar{d}\}$, then the expression becomes $2(y_D^B(\theta) - \underline{d}) \geq 0$. If $y_D^B(\theta) = \theta + B$, then the expression is $2(\theta + B - \underline{d}) \geq 0$ as $\theta + B \in D$. Since $y_D^B(\theta)$ is continuous and almost-everywhere smooth, Equation A.2 is increasing in θ.

(a) The equilibrium at θ is separating if and only if

$$\bar{u}(B) \geq u(B, \underline{d}, 0) \Leftrightarrow (\underline{d} - \theta - B)^2 - (y_D^B(\theta) - \theta - B)^2 \geq \Delta V_2^B(0, 1)$$

If the inequality holds for θ, then it must also hold for all $\theta' > \theta$ as the LHS is increasing in θ.

(b) The equilibrium at θ is pooling if and only if

$$\bar{u}(B) \leq u(B, \underline{d}, p) \Leftrightarrow (\underline{d} - \theta - B)^2 - (y_D^B(\theta) - \theta - B)^2 \leq \Delta V_2^B(p, 1)$$

We have already shown that the LHS is increasing in θ. Hence if for some θ the inequality holds, it also holds for all $\theta' < \theta$.

■

Proof of Proposition 2

Proof. For this proposition, we need two observations. First, it is straightforward to check that for any D and D' such that $D \subseteq D'$,

$$-(y^b_{D'} - \theta - b)^2 \geq -(y^b_D - \theta - b)^2 \tag{A.3}$$

for any $\theta \in \Theta$ and $b \in \{0, B\}$.

Secondly, the following inequality holds for all $\theta \in \Theta$,

$$(\underline{d} - \theta - B)^2 - (y^B_D(\theta) - \theta - B)^2 \leq (\underline{d}' - \theta - B)^2 - (y^B_{D'}(\theta) - \theta - B)^2 \tag{A.4}$$

When $\theta + B \geq \underline{d} \geq \underline{d}'$, $(\underline{d} - \theta - B)^2 \leq (\underline{d}' - \theta - B)^2$. Together with Equation A.3, we have Equation A.4. When $\theta + B < \underline{d}$, then the LHS of Equation A.4 equals 0. The RHS must be non-negative as $y^B_{D'}(\theta)$ can always be set to $\underline{d}'$, thus Equation A.4 holds.

(a) The equilibrium at θ is separating if and only if

$$(\underline{d} - \theta - B)^2 - (y^B_D(\theta) - \theta - B)^2 \geq \Delta V_2^B(0, 1)$$

It suffices to show that whenever the inequality holds at θ under D, it also holds at θ under D'. From Equation A.4, this is straightforward.

(b) The equilibrium at θ is pooling if and only if

$$(\underline{d} - \theta - B)^2 - (y_D^B(\theta) - \theta - B)^2 \leq \Delta V_2^B(p, 1)$$

It suffices to show that whenever the inequality holds at θ under D', it also holds at θ under D. Again Equation A.4 delivers what we need.

(c) By Equation A.4, we have

$$\begin{aligned} \Delta V_2^B \left(\frac{p\gamma_D(\theta)}{p\gamma_D(\theta) + 1 - p}, 1 \right) &\leq \Delta V_2^B \left(\frac{p\gamma_{D'}(\theta)}{p\gamma_{D'}(\theta) + 1 - p}, 1 \right) \\ \Leftrightarrow \quad \frac{p\gamma_D(\theta)}{p\gamma_D(\theta) + 1 - p} &\geq \frac{p\gamma_{D'}(\theta)}{p\gamma_{D'}(\theta) + 1 - p} \end{aligned}$$

for any $\theta \in \Theta_D^m \cap \Theta_{D'}^m$. It can be readily checked that the last line is equivalent to $\gamma_D(\theta) \geq \gamma_{D'}(\theta)$. Thus $m(\underline{d}|B, \theta) \geq m(\underline{d}'|B, \theta)$.

For $\theta \in \Theta_D^p \cap \Theta_{D'}^m$, $m(\underline{d}|B, \theta) = 1 > m(\underline{d}'|B, \theta)$. For $\theta \in \Theta_D^m \cap \Theta_{D'}^s$, $m(\underline{d}|B, \theta) > 0 = m(\underline{d}'|B, \theta)$.

■

Bibliography

Aghion, Philippe and Jean Tirole (1997). "Formal and real authority in organizations". In: *Journal of political economy* 105(1), pp. 1–29.

Alonso, Ricardo and Niko Matouschek (2008). "Optimal delegation". In: *The Review of Economic Studies* 75(1), pp. 259–293.

Amador, Manuel and Kyle Bagwell (2013). "The theory of optimal delegation with an application to tariff caps". In: *Econometrica* 81(4), pp. 1541–1599.

Arguedas, Carmen and Sandra Rousseau (2012). "Learning about compliance under asymmetric information". In: *Resource and energy economics* 34(1), pp. 55–73.

Avery, Christopher and Jonathan Levin (2010). "Early admissions at selective colleges". In: *American Economic Review* 100(5), pp. 2125–56.

Che, Yeon-Koo and Youngwoo Koh (2016). "Decentralized college admissions". In: *Journal of Political Economy* 124(5), pp. 1295–1338.

Cho, In-Koo and David M Kreps (1987). "Signaling games and stable equilibria". In: *The Quarterly Journal of Economics* 102(2), pp. 179–221.

Cho, In-Koo and Joel Sobel (1990). "Strategic stability and uniqueness in signaling games". In: *Journal of Economic Theory* 50(2), pp. 381–413.

Doval, Laura and Vasiliki Skreta (2020). "Mechanism design with limited commitment". In: *Available at SSRN 3281132*.

Frankel, Alexander (2014). “Aligned delegation”. In: *American Economic Review* 104(1), pp. 66–83.

Frankel, Alexander (2016). “Discounted quotas”. In: *Journal of Economic Theory* 166, pp. 396–444.

Guo, Yingni and Johannes Hörner (2017). “Dynamic Mechanisms without Money”. In.

Hamilton, Jonathan H and Steven M Slutsky (1990). “Endogenous timing in duopoly games: Stackelberg or Cournot equilibria”. In: *Games and Economic Behavior* 2(1), pp. 29–46.

Hannan, Michael T, M Diane Burton, and James N Baron (1996). “Inertia and change in the early years: Employment relations in young, high technology firms”. In: *Industrial and Corporate Change* 5(2), pp. 503–536.

Hellmann, Thomas (1998). “The allocation of control rights in venture capital contracts”. In: *The Rand Journal of Economics*, pp. 57–76.

Hoffman, Mitchell, Lisa B Kahn, and Danielle Li (2015). *Discretion in hiring (No. w21709).*

Holmstrom, Bengt Robert (1978). *On Incentives and Control in Organizations.* Stanford University.

Holmström, B (1984). “On the Theory of Delegation,” in: Bayesian Models in Economic Theory. Ed. by M. Boyer, and R. Kihlstrom. North-Holland, New York”. In.

Kaplan, Steven N and Per Strömberg (2003). “Financial contracting theory meets the real world: An empirical analysis of venture capital contracts”. In: *The review of economic studies* 70(2), pp. 281–315.

Kempf, Hubert and Grégoire Rota Graziosi (2010). “Leadership in public good provision: a timing game perspective”. In: *Journal of Public Economic Theory* 12(4), pp. 763–787.

Laffont, Jean-Jacques and Jean Tirole (1987). “Comparative statics of the optimal dynamic incentive contract”. In: *European Economic Review* 31(4), pp. 901–926.

Laffont, Jean-Jacques and Jean Tirole (1988). “The dynamics of incentive contracts”. In: *Econometrica: Journal of the Econometric Society*, pp. 1153–1175.

Lee, Sang-Ho and Lili Xu (2018). “Endogenous timing in private and mixed duopolies with emission taxes”. In: *Journal of Economics* 124(2), pp. 175–201.

Li, Jin, Niko Matouschek, and Michael Powell (2017). “Power dynamics in organizations”. In: *American Economic Journal: Microeconomics* 9(1), pp. 217–41.

Lipnowski, Elliot and Joao Ramos (2020). “Repeated delegation”. In: *Journal of Economic Theory* 188, p. 105040.

Mailath, George J (1987). “Incentive compatibility in signaling games with a continuum of types”. In: *Econometrica: Journal of the Econometric Society*, pp. 1349–1365.

Mailath, George J (1993). “Endogenous sequencing of firm decisions”. In: *Journal of Economic Theory* 59(1), pp. 169–182.

Melumad, Nahum D and Toshiyuki Shibano (1991). “Communication in settings with no transfers”. In: *The RAND Journal of Economics*, pp. 173–198.

Normann, Hans-Theo (2002). “Endogenous timing with incomplete information and with observable delay”. In: *Games and Economic Behavior* 39(2), pp. 282–291.

Pan, Siqi (2018). “Exploding offers and unraveling in two-sided matching markets”. In: *International Journal of Game Theory* 47(1), pp. 351–373.

Riley, John G (1979). “Informational equilibrium”. In: *Econometrica: Journal of the Econometric Society*, pp. 331–359.

Rivera, Lauren A (2012). “Hiring as cultural matching: The case of elite professional service firms”. In: *American sociological review* 77(6), pp. 999–1022.

Tanner, Noam (2018). “Optimal Delegation Under Unknown Bias: The Role of Concavity”. In: *FRB Boston Risk and Policy Analysis Unit Paper No. RPA*, pp. 18–1.