The Pennsylvania State University
The Graduate School

ESSAYS IN LIMITED COMMITMENT

A Dissertation in
Economics
by
Vasundhara Mallick

Submitted in Partial Fulfillment
of the Requirements
for the Degree of

Doctor of Philosophy

August 2024

The dissertation of Vasundhara Mallick was reviewed and approved by the following:

S. Nageeb Ali
Professor of Economics
Dissertation Advisor
Chair of Committee

Nima Haghpanah
Associate Professor of Economics

Miaomiao Dong
Assistant Professor of Economics

Chloe Tergiman
Associate Professor of Risk Management

Barry Ickes
Professor of Economics
Department Head, Department of Economics

Abstract

This dissertation comprises two chapters.

The first chapter, "The Art of Waiting", co-authored with Ece Teoman, studies delegated project choice without commitment: a principal and an agent have conflicting preferences over which project to implement, and the agent is privately informed about the availability of projects. We consider a dynamic setting in which, until a project is selected, the agent can propose a project, and the principal can accept or reject a proposed project. Importantly, the principal cannot commit to his responses, and cannot implement a project unless it is proposed. In this setting, the agent has an incentive to hold back on proposing projects that the principal favors so that the principal approves a project favored by the agent. Nevertheless, the principal achieves his commitment payoff in an equilibrium of the game in the frequent-offer limit. This high payoff equilibrium showcases the art of waiting and contrasts with Coasian logic: by giving proposer power to the agent, the principal makes it credible to reject his dispreferred projects until later in the game giving the agent an incentive to propose principal-preferred projects earlier on. We apply these results to the economics of organization. In particular, these results suggest that to curb a manager's *empire building* plans, eliciting proposals from her "bottom-up" might be better than issuing "top-down" commands.

The second chapter, "How Markets Disrupt Mediated Trade", studies markets with adverse selection and the degree to which intermediaries can foster efficient trade. I consider a setting in which a seller and buyer have interdependent values. Without any intermediation, the Lemons Problem guarantees that only the lowest type trades in any equilibrium. I consider an intermediary who brokers trade between the seller and the buyer by using a screening mechanism. When this is the *only* channel for trade, more efficient outcomes are possible in equilibrium, where higher types trade with positive probability. My main result, however, concludes that once the seller can *also* sell her asset without going through the intermediary, market failures re-emerge: trade of assets above the lowest quality shuts down in *both* the decentralized and mediated market. This paper shows that intermediation might be rendered completely ineffective when assets cannot be exclusively traded through the intermediary.

Table of Contents

List of Figures

List of Tables

Acknowledgments

I am deeply grateful to my advisor S. Nageeb Ali, for being a tremendous source of guidance, support, and encouragement. To say that Nageeb has been an amazing advisor would be an understatement. Not only is he always very generous with his time, I think he has, at times, spent more time thinking about my research than I have! And neither of my papers would be fit for human consumption without him; he suffered through several abominable drafts of my papers so that others don't have to. He has also always been a source of pep talks and positive energy; he genuinely cares about his students being happy overall. He has tremendously influenced the way I think, read papers, approach research, and exposit ideas; his contribution to my journey as a researcher is immesurable.

My other committee members, Nima Haghpanah and Miaomiao Dong, were a constant source encouragement and support over the years. Miaomiao is one of the most clear thinkers I have come across and and has always been extremely generous with her time and feedback. Every meeting with Nima is so clarifying; his amazing incisive questions deeply enrich my research questions whenever I talk to him. And his course on mechanism design is probably the best I have taken at Penn State.

Over the years, I have benefited immensely from conversations with Kalyan Chatterjee, Berk Idem, Yuhta Ishii, Vijay Krishna, Kala Krishna, Wenhao Li, Kippeum Lee, Xiao Lin, Amrutha Manjunath, Orville Dustin Mondal, Ran Shorrer, Ron Siegel, Darshana Sunoj, Chloe Tergiman, and Christelle Zozungbo. I would also like to thank my PhD cohort for their support and friendship over the years, especially Wenhao Li, Darshana Sunoj, Kapilan Mahalingam, and Carlos Salamanca. Ece Teoman is an amazing co-author and I have learnt a lot by working with her. Rohit Lamba has always been exceptionally kind, and without his encouragement, I would probably have abandoned the idea that ultimately became my job market paper. I would also like to thank Krista Winkelblech for her patience in helping me with various administrative stuff since the day of my arrival at Penn State. She really made my life here so much easier. I am also indebted to the various restaurants of State College for keeping me, a person with disastrous cooking skills, alive for all these years.

Finally, I would like to thank my family, who are an integral part of who I am. My parents, Rachna and Saurabh, never pushed me to do an MBA, and therefore, this dissertation would not exist without them. They have always believed in me, and have always encouraged me to make choices for my own happiness, even ones that are risky, and for that, I'm very grateful. My brother Aditya is one of the smartest people I know

and I'm always inspired and amazed by how creative he is. Our dog Bruno, who passed way towards the end of my second year, was the best, most noble dog ever and brought so much joy and unconditional love to our lives for so many years. I miss him so much and I wish he was here. Given the amount of uncertainty a PhD involves, the unconditional love and support of my family had truly been an anchor for me and none of this would have been possible without them.

Chapter 1
The Art of Waiting

1.1 Introduction

This paper considers a principal-agent problem with the following two features: (i) the agent knows what actions or "projects" are feasible and the principal does not, and (ii) the interests of the two parties are not aligned. Such principal-agent problems abound. Consider the interaction between a CEO (principal) of a firm and a manager (agent). The manager may be better informed about what actions the firm can undertake, and unlike the CEO, is motivated by empire building. In such cases, the CEO cannot blithely assume that the manager selects actions purely for shareholder interests. Another example is that of an antitrust authority deciding which mergers to approve: It only wants to approve those mergers that enhance efficiency or consumer welfare, but firms would like to propose only those mergers that increase industry profits. In such settings, what should the principal do?

These issues have been studied in the literature on *project selection* problems, initiated by the seminal work of Armstrong and Vickers [2010] and Nocke and Whinston [2013]. The dominant approach presumes that the principal can commit to which projects he would accept in a one-shot interaction. But in many settings, the principal may be unable to commit, particularly if projects can be proposed across several rounds. If the agent does not propose any project that the principal deems acceptable, the principal may then infer that such projects are infeasible and capitulate. Anticipating this reaction, the agent may then wish to hold back on proposing projects that the principal finds acceptable. How well can the principal do and can he obtain his commitment payoff?

We investigate this question in a dynamic framework. The agent is privately informed about which projects are feasible at time 0. In each round $t \in \{0, 1, 2, \ldots\}$, the agent can propose a project that is feasible or stay silent; if a project is proposed, the principal can accept or reject it. This process continues until a project is accepted—in which case, players obtain payoffs from that selected project— or no proposed project is *ever* accepted, so all players obtain payoffs from the status quo. We consider the frequent-offer limit of this model, a sequence of games where the period length vanishes.

In such a setting, one may anticipate that the principal would suffer a significant loss of payoffs relative to the (static) commitment benchmark: after all, given the logic sketched above, the principal may capitulate when no acceptable project is proposed. Moreover, our extensive form endows the agent with *proposal power*: The principal is effectively giving up bargaining or proposal power, so even complete-information intuitions suggest that the principal will do poorly in the dynamic game. Our main result, informally stated, is

> **Theorem**. *In the frequent-offer limit, the principal attains his commitment payoff in an equilibrium of the game.*

The key idea underlying our main result is that endowing the agent with the right to propose, along with restricting the principal's action to accepting or rejecting a proposed project, circumvents the principal's commitment problem. Our high payoff equilibrium stipulates that the agent and the principal wait for many rounds before respectively proposing and accepting any project that the agent greatly prefers to other projects and the principal disprefers (but prefers to the status quo). We show that this behavior is sequentially rational, even at histories where the principal attributes probability 1 to the agent only having such projects. If the agent proposes such projects earlier than specified, the principal believes that the agent must have other projects at her disposal, and rejects the proposal. Such "punishment through beliefs" incentivizes the agent not to propose such projects earlier than stipulated and as such solves the principal's commitment problem. Because the agent anticipates such delays to get her preferred projects approved, she is willing to propose, as in the commitment benchmark, feasible projects that the principal prefers (and she may disprefer).

We observe that it is essential that the agent has proposal power for the commitment problem to be solved. By contrast, if the principal were the one making all the offers, Coasian forces would take over, resulting in him granting full discretion to the agent. It is also crucial that the principal cannot implement projects unless the agent proposes them. Otherwise, at a history where the principal believes that only his dispreferred projects are available, he would implement them himself.

We view this finding to be of more than just theoretical interest. It suggests a gain to organizations from allowing the agents—be managers, or employees—to propose projects rather than issuing "top-down" commands from the principal. One may envision that such "bottom-up" organizational structures emerge to motivate the agent, as implemented projects follow their initiative as proposed in Aghion and Tirole [1997]. Our work also suggests that antitrust authorities, venture capital boards, and grant funding agencies may gain from allowing the agent to be the one to propose projects flexibly rather than constraining the agent to propose only certain projects.

We also explore what may happen when the principal and the agent interact in ways other than the agent always proposing and the principal accepting or rejecting these proposals. We study a general class of *delegation protocols* and compare these under commitment; i.e. if the principal could commit to a strategy within a protocol. We show that the delegation protocol we study does as well as any delegation protocol under commitment, and always achieves the payoff from the best static stochastic mechanism. When there are only two possible projects, we show that this commitment benchmark is always attained in an equilibrium; with more than two possible projects, some additional assumptions are needed to achieve the commitment benchmark without commitment.

The rest of the paper is organized as follows. We first present the related literature. Section 2 introduces the setup, describes the sequential delegation game, and establishes a commitment benchmark. We present our main result in Section 3 which exhibits the main forces and the intuition behind them in the cleanest way. We establish a commitment benchmark as an upper bound for the principal's payoff and our main result shows that the commitment payoff is always attained in an equilibrium of the game. In Section 4, we include the general analysis for N projects. There, we work in a class of delegation protocols and show that our sequential delegation game is an optimal protocol under commitment. We also extend our main result beyond two projects under some regularity conditions. Section 5 concludes the paper.

1.1.1 Related Literature

Our paper studies a delegation problem in a project selection setup. This problem is studied by Armstrong and Vickers [2010] in a static setup. In their work, the formal authority to choose the project lies with the agent but the principal can restrict the set of projects the agent can choose from. This is equivalent to the principal being able to commit to a deterministic mechanism. In contrast, our game is dynamic, with the agent making proposals, and we use the optimal *stochastic* mechanism as the commitment benchmark. We show that this commitment payoff can be attained in an equilibrium of our game.[1]

In Aghion and Tirole [1997], the uncertainty is about the payoffs from projects, rather than their feasibility. Both *formal* authority to make decisions, and *real* authority, where the agent's proposals are accepted by the principal, can incentivize the agent to exert effort to learn the payoffs. Although our setup is quite different, the agent having *real* authority in their paper is similar to the agent having the control over proposals in our model. In our work, the principal can only implement what the agent proposes, so even though the he has the ultimate authority to make decisions, agent has significant *real* control over the decision.

Our setup is one that involves *hard evidence*, since the the agent's private information is about the *feasibility* of decisions, and she can only propose available projects. This is an important way in which our paper differs from some of the broader literature on delegation starting with Holmström [1984], and more recently, Alsonso and Matouschek [2008]. They also study a joint decision problem where the principal has the formal authority to take decisions but the agent possesses private information relevant to decision-making. In these models, if the principal can commit to a decision rule as a function of the agent's report, then the formal allocation of authority is irrelevant. This is because any type can imitate the report of another. So, the optimal mechanism can be implemented by constrained delegation, where the principal delegates decision-making to the agent, but restricts her to choose from a *delegation set*.

In contrast to this, in our setup, even if the principal can commit, the optimal static mechanism wouldn't, in general, be equivalent to choosing a delegation set. This is because in our commitment benchmark, since the agent can only include available projects in her report, this helps convey her private information more effectively, rather than if she simply chooses an available project from a delegation set. So the principal making decisions as a function of the agent's report results in more effective communication of agent's private information.

An alternative approach in the project selection literature is to model the interaction as a cheap talk game where the principal lacks commitment power, as in Che, Dessein, and Kartik [2013] and Balzer and Schneider [2019]. Che, Dessein, and Kartik [2013] finds that in the presence of a bias regarding the outside option, the agent tends to propose unconditionally better projects for the principal to secure the approval for implementation.

[1] Nocke and Whinston [2013] study a similar problem in a static setup, in the context of mergers. An antitrust authority can commit ex-ante to its merger-approval rule. However, this is not a direct static counterpart of our setup. There are multiple firms (agents) here, and given the set of permitted and feasible mergers, the implemented merger is the result of a bargaining process among firms.

As its dynamic extension, Balzer and Schneider [2019] finds that the dynamic interaction allows for different equilibrium outcomes. It characterizes a mixing equilibrium where the agent randomizes between pandering and not, and a waiting equilibrium similar to ours in spirit where the agent waits to persuade the principal that the unconditionally worse project has a better payoff realization. This is the paper that closest to ours. The key difference is that our work establishes a commitment benchmark and for the case of two projects (and for N projects, under some conditions), characterises the equilibria that achieves the commitment payoff.

Our paper also relates to the literature on mechanism design with hard evidence, starting with Green and Laffont [1986] and Bull and Watson [2007]. In these papers, the setting is endowed with an evidentiary structure under which different mechanisms are compared and Revelation Principles are established. We start with a type dependence of evidentiary actions in delegation protocols and establish an evidentiary structure. Static mechanisms with this evidentiary structure serve as an upper bound for any outcome that can be implemented with any mechanism with this particular type dependence of actions. The Revelation Principle from these papers does not hold directly in ours because the agents's action also restricts what the principal can choose, and evidentiary actions can be taken at multiple nodes, which these papers do not allow. The paper closest to our mechanism design analysis is Deneckere and Severinov [2008] which has a Revelation Principle that allows for evidentiary actions at multiple nodes. However, this result does not follow directly in our setting either. If the agent is the proposer, the principal cannot implement a project she does not propose, so evidentiary actions have significance beyond providing verifiable information.

1.2 The Model

A principal (he) and an agent (she) jointly choose a project to implement. A project is a pair of payoffs $(\alpha, \pi) \in \mathbb{R}^2_{++}$, where α is the agent's payoff from implementing the project, and π is the principal's. There are two *possible* projects, denoted by $\mathcal{N} = \{g, b\}$: a good project g with payoffs (α_g, π_g) and a bad project b with payoffs (α_b, π_b). The players are expected utility maximizers, and have conflicting preferences over the projects: $\pi_g > \pi_b > 0$ and $\alpha_b > \alpha_g > 0$ (Figure 1.1).

The challenge is that not every *possible* project may be *available* to implement, and only the agent knows which projects available: her *type* represents the set of available projects, and is her private information. The agent has four possible types:

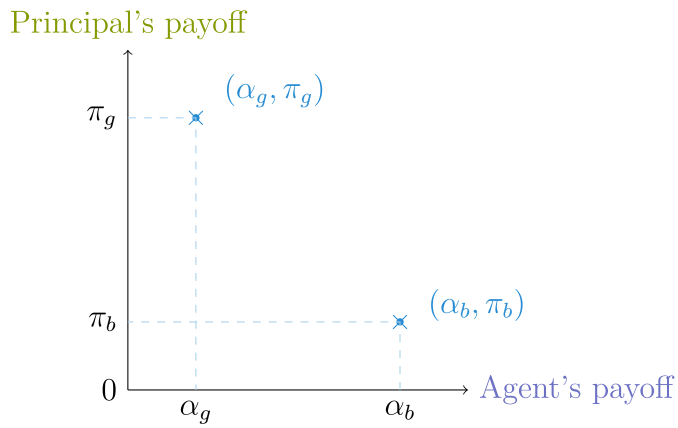

Figure 1.1. The project space with the principal-preferred good project (α_g, π_g) and the agent-preferred bad project (α_b, π_b).

- $E = \varnothing$, the *empty* type with no available projects;
- $G = \{g\}$, the *good* type with only the good project available;
- $B = \{b\}$, the *bad* type with only the bad project available;
- $M = \{g, b\}$, the *mixed* type with both projects available.

The set of all possible types is denoted by $\mathcal{S} \equiv 2^{\mathcal{N}} = \{E, G, B, M\}$.[2]An element of $\mathcal{S}$, or a possible type, is denoted by S. The agent's type is drawn from $\mathcal{S}$ and μ_S denotes the probability of type S.

We now describe how the principal and the agent solve the joint decision problem. We refer to their interaction as the *sequential delegation game*, which proceeds as follows. Time is discrete and the principal and the agent have a common discount factor $\delta \in (0, 1)$. In each period $t = 0, 1, 2, \ldots$, the agent makes a proposal and the principal responds to this proposal.

The set of actions, or proposals available to the agent, is type-dependent. If the agent is of type S, then at any time period t, the agent can make proposals from the set $A_S = \{\{i\} | i \in S\} \cup \varnothing$. This means that in any t, the agent of type S can either propose exactly one available project $i \in S$, or stay silent.[3] For the principal, if at any t, project $i \in \mathcal{N}$ is proposed by the agent, he can either *accept* i, or *reject* it. If the agent was silent and did not propose anything, then the only possible action for the principal is to *reject*.

[2]This is assuming that μ has full support on $2^{\mathcal{N}}$; this is not essential for our results.
[3]For type E, $A_E = \varnothing$, so staying silent is the only option.

If the principal rejects a proposed project i, or if the agent is silent, both players obtain a payoff of 0 in that period and the game proceeds to the next period. If the principal accepts the proposal of project i at time t, the game ends: the principal's payoff is $\delta^t \pi_i$ and the agent's is $\delta^t \alpha_i$ (Figure 1.2). We focus on the case of $\delta \to 1$, which we interpret it as the frequent-offer limit of the game.

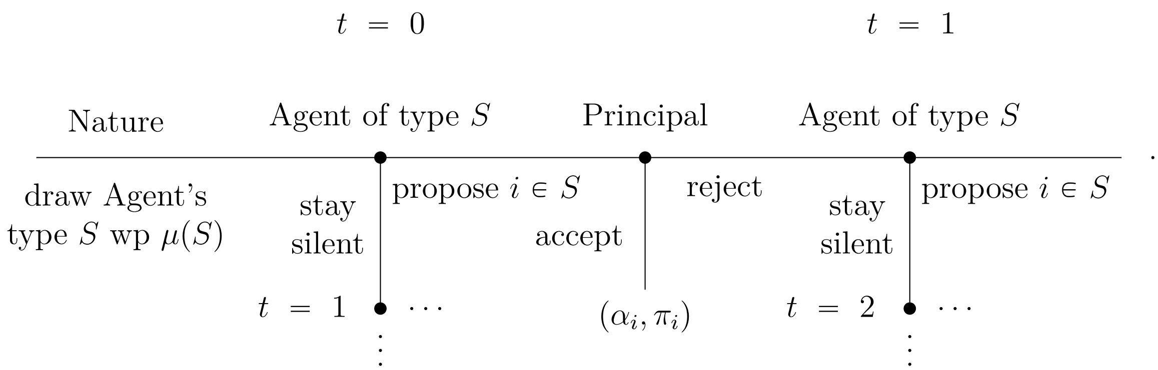

Figure 1.2. Timeline of the sequential delegation game.

Thus, in the sequential delegation game, for any time period t, the set of all possible histories at the beginning of period t is $\mathcal{H}_t = (\mathcal{N} \cup \varnothing)^t$. This captures the fact that if we are at t, for any $t' \leqslant t-1$, we can have two cases: (i) a project $i \in \mathcal{N}$ was proposed and rejected, or (ii) Agent was silent, and nothing was proposed, so there is nothing for the principal to accept. This is denoted by $\varnothing$. An element of $\mathcal{H}_t$ is denoted by h_t. If the agent is of type S, her strategy maps any history to a probability distribution over $\{\{i\}|i \in S\} \cup \varnothing$. A principal's strategy maps any history, and a proposal of project i at that history to a probability of accepting i. If agent is silent, then principal's only possible action is to reject. Our equilibrium concept is Perfect Bayesian Equilibrium; both players play sequentially rationally and the principal's beliefs about the agent's type are updated according to Bayes' rule whenever possible.

1.3 The Benefits of Giving Up Control

This section presents our main result. We study how *well* the principal can do in an equilibrium of our sequential delegation game, despite being uninformed and lacking proposal power. We first establish a commitment benchmark that acts as an upper bound for what the principal could achieve if he could commit to a strategy in the sequential delegation game. Then, our main result shows that this commitment payoff is, in fact, attainable in an equilibrium of the sequential delegation game, even though the principal has no commitment power. Finally, we highlight the importance of *giving up control* over making proposals; in a game where the principal makes proposals, it may not be possible to attain the commitment payoff.

1.3.1 Commitment Benchmark

We first define a class of static, stochastic mechanisms with type-dependent message spaces, that we refer to as mechanisms hereafter. In Section 4.1, we prove that this class of mechanisms is indeed an upper bound to what the principal can achieve if he can commit to a strategy in the sequential delegation game. For now, we will take this fact as given.[4] In a mechanism in this class, the message space is type dependent, and the set of messages that a type S of the agent can send is defined to be $M(S) = 2^S$. So, each type can report only subsets of her available projects as a message in the mechanism.[5] A mechanism is a tuple (M, q), where $M = \bigcup_{S \in \mathcal{S}} M(S)$ is the set of all possible messages, and $q : \mathcal{S} \to \Delta(S \cup \varnothing)$ is the outcome function. Therefore, when S is reported, only projects *in* S can be implemented, or no project at all, as captured by $\varnothing$. If no project is implemented, the players obtain the status quo payoff, which is zero for both the principal and the agent. For any type $S \in \mathcal{S}$ and any project $i \in S$, q_{Si} represents the probability of implementing project i when the type S is reported.[6] We define a mechanism to be *incentive compatible* (IC) if no type finds it optimal to report a strict subset.

Every mechanism determines an allocation, which is a vector $\{q_{Si}\}_{S \in \mathcal{S}, i \in S}$. A *feasible* allocation is one where for each type S, we have: $q_{Si} \in [0, 1]$ for each $i \in S$ and $\sum_{i \in S} q_{Si} \leqslant 1$. The principal-optimal mechanism maximizes the principal's payoff by choosing implementation probabilities for projects in each type, subject to feasibility and incentive compatibility constraints.

[4] It is in fact a *tight* upper bound in the sense that for any mechanism, there exists a strategy in our sequential delegation game such that commitment to this strategy gives the principal the same expected payoff as the mechanism as $\delta \to 1$. So, this upper bound is equivalent to commitment within the sequential delegation game. We show this in Section 4.1. There, we also show that this class of mechanisms actually acts as an upper bound for a very general class of interactions between the principal and the agent, not just the one where the agent makes all the offers.

[5] This includes the empty set.

[6] So, $1 - \sum_{i \in S} q_{Si}$ is the probability of not implementing any project when S is reported.

$$\max_{q_{Gg},q_{Bb},q_{Mg},q_{Mb}\in[0,1]} \quad \mu_G q_{Gg}\pi_g + \mu_B q_{Bb}\pi_b + \mu_M q_{Mg}\pi_g + \mu_M q_{Mb}\pi_b$$

$$\text{subject to} \qquad \textstyle\sum_{i\in S} q_{Si}\alpha_i \geqslant \sum_{i'\in S'} q_{S'i'}\alpha_{i'} \qquad (\mathrm{IC}_{SS'})$$

$$q_{Mg} + q_{Mb} \leqslant 1$$

where $\mathrm{IC}_{SS'}$ denotes the IC constraint for type S to not report type $S' \subseteq S$. The second constraint is just the feasibility constraint for the implementation probabilities when the mixed type is reported. Before we solve for the principal-optimal mechanism, we make a few simplifying observations about some properties that must be true for an optimal mechanism.

Observation. No type finds it profitable to report the empty type, so IC_{GE}, IC_{BE}, and IC_{ME} are all redundant.

Recall that corresponding to any report, the mechanism only implements project *in* the report. So, the payoff from reporting the empty type is zero; it has no project, so no project is implemented when it is reported. Since each project has strictly positive payoffs for both players, each type of the agent gets a weakly higher payoff by reporting her own type, than she does by reporting the empty type. Hence, the IC constraints for the other types to not report the empty type, IC_{GE}, IC_{BE}, and IC_{ME}, are all redundant.

Observation. In an optimal mechanism, we must have $q^*_{Mg} + q^*_{Mb} = 1$.

When the mixed type is reported, the probabilities of implementing the good and the bad projects must sum up to 1 in an optimal mechanism. Suppose not, i.e. $q^*_{Mg} + q^*_{Mb} < 1$. Then, we can increase both q^*_{Mg} or q^*_{Mb} slightly, and have new implementation probabilities $(q^{**}_{Mg}, q^{**}_{Mb}) > (q^*_{Mg}, q^*_{Mb})$ and $q^{**}_{Mg} + q^{**}_{Mb} < 1$. It must be the case that the IC constraints involving the mixed type still hold, as

$$q^{**}_{Mg}\alpha_g + q^{**}_{Mb}\alpha_b > q^*_{Mg}\alpha_g + q^*_{Mb}\alpha_b \geqslant q_{Gg}\alpha_g$$

$$q^{**}_{Mg}\alpha_g + q^{**}_{Mb}\alpha_b > q^*_{Mg}\alpha_g + q^*_{Mb}\alpha_b \geqslant q_{Bb}\alpha_b.$$

From these new implementation probabilities, the principal obtains a strictly higher payoff. So, $q^*_{Mg} + q^*_{Mb} < 1$ cannot be part of an optimal mechanism.

Observation. The incentive compatibility constraint for the mixed type to not report the good type, IC_{Mg}, is redundant.

We know from the previous observation that when the mixed type is reported, the probabilities of implementing projects sum up to 1. It implies that the payoff of the mixed type, from reporting truthfully, will be at least α_g. On the other hand, her payoff from reporting the good type is at most α_g as $q_{Gg} \in [0, 1]$. Thus, the mixed type is always weakly better off by reporting truthfully than by pretending to be the good type, and IC_{MG}, is redundant.

Observation. In any optimal mechanism, we must have $q_{Gg} = 1$.

Since IC_{MG} is redundant, and there is no type other than the nixed type that can report the good type, therefore when the good type is reported, an optimal mechanism must implement the good project with certainty.

Given the above *observations*, the problem of finding the optimal mechanism reduces to that of choosing q_{Mg} and q_{Bb} to maximise the principal's expected payoff, subject to IC_{MB}. This is because the other IC constraints are redundant, $q_{Gg} = 1$, and $q_{Mg} + q_{Mb} = 1$, so choosing q_{Mg} and q_{Bb} pins down the optimal mechanism.

The lone IC constraint, IC_{MB}, represents the trade off that the principal faces in implementing the good project with positive probability from the mixed type. If $q_{Mg} > 0$, the principal will have to set $q_{Bb} < 1$, so that the mixed type doesn't imitate the bad type. We now define two mechanisms, and it turns out that one of them is always *an* optimal mechanism.

Definition 1. *The pooling mechanism implements the bad project from the agent's bad and mixed types:* $q^*_{Gg} = 1, q^*_{Bb} = 1, q^*_{Mg} = 0, q^*_{Mb} = 1$.

Definition 2. *The separating mechanism implements the good project from the mixed type and the bad project from the bad type with an interior probability:* $q^*_{Gg} = 1, q^*_{Bb} = \frac{\alpha_g}{\alpha_b}, q^*_{Mg} = 1, q^*_{Mb} = 0$.

It is easy to see that both these mechanism are IC. In the *pooling* mechanism, the outcome when the type is mixed, is same as the outcome when the type is bad. In both cases, the bad project is implemented with probability one. So, mixed type is *pooled* with the bad. On the other hand, in the *separating* mechanism, the outcome when the type is mixed is different from the outcome when the type is bad. In one case, the good project is implemented with probability one, and in the other, the bad project is implemented with probability $\frac{\alpha_g}{\alpha_b}$. So, this mechanism *separates* the mixed and bad types.

Before stating our result about the optimal mechanism, we establish some notation. Let $\lambda = \frac{\mu_M}{\mu_B}$ be the likelihood ratio of mixed type compared to bad type, and $\lambda^* = \frac{(1-\frac{\alpha_g}{\alpha_b})\pi_b}{(\pi_g - \pi_b)}$.

Proposition 1. *Either the pooling or the separating mechanism is always optimal.*

a) When $\lambda < \lambda^*$, the pooling mechanism is optimal.

b) When $\lambda > \lambda^*$, the separating mechanism is optimal

c) When $\lambda = \lambda^*$, any mechanism with $q_{Gg} = 1$, $q_{Mg} + q_{Mb} = 1$, and $q_{Mg}\alpha_g + q_{Mb}\alpha_b = q_{bB}\alpha_b$ is optimal. In particular, both the pooling and separating mechanisms are optimal.

The proof of the above result is in the appendix. To see the intuition behind the result, recall that we only need to worry about IC$_{MB}$. We can show that this reduces the problem to: (i) whether the principal wants to *separate* the mixed and bad types, as in the separating mechanism, or (ii) *pool* them, as in the pooling mechanism. Consider the case where $\lambda = \frac{\mu_M}{\mu_B} \to \infty$. Then it is *as if*, between the mixed and the bad types, only the mixed type exists. In this case, it would be optimal to set $q_{Mg} = 1$, and therefore, $q_{Bb} = 0$ bu IC$_{MB}$. On the other hand, if $\frac{\mu_M}{\mu_B} = 0$, i.e only the bad type exists, then it is optimal to set $q_{Bb} = 1$. So it makes sense that if λ is high enough, the separating mechanism would do better than the pooling mechanism. This is precisely λ^*, the threshold in our result.

Now that we have solved for the optimal mechanism, we turn our attention back to the sequential delegation game where the principal cannot commit to his responses. In particular, we are interested in exploring how well the principal can do in equilibria of the game when he is constrained by sequential rationality and whether he can obtain his commitment payoff.

1.3.2 Implementing the Commitment Benchmark as Equilibrium

In the sequential delegation game, the agent's proposals are like *reports* of which projects are available, somewhat like the mechanism. However, unlike the mechanism, the principal cannot commit to his responses to the agent's proposals here. This therefore limits how effectively he can elicit the relevant information from the agent, and we might expect a gap between the commitment payoff and what the principal can achieve in an equilibrium of the game.

In the absence of commitment power, one might think that proposal power could help the principal; he could provide the necessary incentives by effectively restricting the choices of the agent. However, in our game, the principal lacks the ability to make proposals as well! This translates into a reduced level of control over which project is implemented, as he cannot implement something that the agent has *not* proposed. He can only accept or reject a project proposed by the agent. The principal seems to be at every possible disadvantage here.

Our main result however, contradicts the above intuition, and establishes that there is always an equilibrium of the sequential delegation game where the principal attains his commitment payoff in the frequent-offer limit. In particular, there is an equilibrium where delay emerges as a costly signalling device for the agent and allows for the separation of the mixed and the bad type. This equilibrium attains the optimal commitment payoff when the optimal mechanism is separating. We will also argue that the lack of proposal power in fact helps with this *signalling through delay*. We now state our main result where, for each form that the optimal mechanism takes, we describe an equilibrium that attains the payoff from this optimal mechanism in the frequent offer limit.

Theorem 1. *There is always an equilibrium of the sequential delegation game in which the principal's payoff approximates his commitment payoff in the frequent-offer limit, as* $\delta \to 1$*. On-path behavior in the equilibria that attain the commitment benchmark is as follows.*

a) (Pooling) When the pooling mechanism is optimal, $\lambda \leqslant \lambda^*$*, the pooling equilibrium attains the principal's commitment payoff: Each type of the agent proposes her favorite available project at* $t = 0$*. The principal accepts any proposal at* $t = 0$*.*

b) (Separating) When the separating mechanism is optimal, $\lambda > \lambda^*$*, the separating equilibrium approximates the principal's commitment payoff:*

- *The agent's good and mixed types propose the good project at* $t = 0$ *and the bad type stays silent until* $t^*(\delta) := \min\{t : \alpha_g \geqslant \delta^t \alpha_b\}$*, at which point she proposes the bad project.*
- *The principal accepts the good project at* $t = 0$ *and the bad project at* $t^*(\delta)$*.*

The details of the strategies and beliefs that constitute the pooling and the separating equilibria are in the Appendix. Here, we focus on the more interesting case; that of the separating equilibrium (Figure 1.3). We first argue that the on path behaviour we described indeed attains the principal's payoff from the separating mechanism. In this equilibrium, on the equilibrium path, the mixed type and the good type both propose the good project g at $t = 0$ and it is accepted. So, since g is implemented, i.e. proposed and accepted without delay, this replicates the implementation probabilities of $q_{Gg} = q_{Mg} = 1$. In the mechanism, the bad project b is implemented with an interior probability of $q_{Bb} = \frac{\alpha_g}{\alpha_b}$, and in the equilibrium is is implemented (proposed and accepted) with a delay, at $t^*(\delta)$. By definition of $t^*(\delta)$, we have that as $\delta \to 1$, $t^*(\delta) \to \frac{\alpha_g}{\alpha_b}$. Thus, as $\delta \to 1$, the principal's payoff is the separating mechanism is attained by the separating equilibrium.

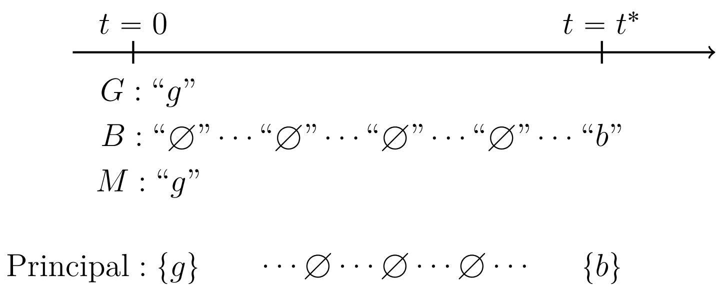

Figure 1.3. The timing of the proposals and accepted projects on the path of separating equilibrium of the sequential delegation game.

We now informally describe the separating equilibrium, and provide some intuition behind the key forces that hold this equilibrium together. The principal's strategy involves accepting the g whenever proposed, and rejecting b whenever it is proposed before the threshold $t^*(\delta)$. In fact, if b is proposed before $t^*(\delta)$, the principal's strategy is to reject not just this proposal, but *any* future proposal of b as well. More precisely, fix a history h_t at the beginning of period t. Suppose there is a $t' < t$ such that along h_t, b was proposed at t', and $t' < t^*(\delta)$. Then, if b is proposed at period t following h_t, it is rejected. If b is proposed at $t^*(\delta)$, and has never been proposed before this, then it is accepted.

This essentially means that a type that has b *cannot* get the principal to accept it before $t^*(\delta)$. Rather, if such a type proposes b before $t^*(\delta)$, then the principal will never accept b at any future time period. So, the mixed type faces a choice: it can get g accepted right away, at $t = 0$, or wait till $t^*(\delta)$ to get b implemented. By definition of $t^*(\delta)$, it (weakly) prefers to propose g at $t = 0$. The bad type has no option but to wait by staying silent till $t^*(\delta)$. At this point he proposes his only project, and it is accepted.

Thus, delay emerges as a costly signalling device in equilibrium; it is used by the bad type to signal that she indeed only has the bad project. But why does the principal reject b at *any* history that involves b being proposed before $t^*(\delta)$? This is because of his off-path beliefs. At any such history, he believes that it is the mixed type with probability one. Thus, the agent gets punished by the extremal off-path beliefs of the principal, if he ever proposes the bad project before $t^*(\delta)$.

However, ex-ante, it is not clear why this *punishment through beliefs* should be possible. Even if the principal attaches probability one to the mixed type, why does he find it optimal to always reject b given this belief? The agent still has control of the proposals after all, and the principal cannot implement something she doesn't propose. In this case, even if the principal *knows* the agent has both projects, its not obvious that he can make the agent propose g. The agent can just keep proposing b.

The intuition here is that in the complete information counterpart of our game, where it is common knowledge that the agent is of mixed type, there exists an equilibrium where, on path, the agent proposes g at $t = 0$. Consider the following strategies of the principal and the agent: the principal, irrespective of history, rejects b, and accepts g. The agent, irrespective of history, proposes g. In particular, at any history h_t, if the agent proposes g and it is rejected, then at this off-path history $h_{t+1} = (h_t, b)$, the agent's strategy is to propose g. It is easy to see that no party has a profitable one-shot deviation. For the principal, at any time period, if b is proposed, by rejecting it, he expects g to be proposed in the next period, which he would then accept. So, if he is sufficiently patient, it is optimal to reject b. For the agent, at any time period, if she proposes b, it would be rejected, and she would propose g in the next period, which would be accepted. So, her payoff from this deviation is $\delta\alpha_g$. If she doesn't deviate and proposes g in the current period, it is accepted and she gets α_g. Thus, the strategies constitute an SPE.[7]

It is *this* equilibrium that our analysis leverages. Consider a history h_t, where b was proposed and rejected at $t' < t^*(\delta)$ along this history. In the separating equilibrium, if the agent's type is mixed, her strategy is to propose g following such a history. So, if b is proposed at $t < t^*(\delta)$, the principal believes it is the mixed type with probability one, and therefore expects a proposal of g in the next period if he rejects b. This makes rejection of b sequentially rational for the principal at *any* such history, and holds this equilibrium together.

We have therefore argued that off-path beliefs can be used to exploit a complete information equilibrium, and separate the mixed type from the bad. We now argue that there is another force: *lack of proposal power*, that makes this leveraging this complete information equilibrium possible. Consider an alternative to our sequential delegation game, where, in each period, the principal makes an offer in form of a restriction set,

[7]The *discreteness* of the offer space is important here. Consider the setting from Rubinstein [1982], but with one party making all the offers. Then, in the unique equilibrium, this party captures the entire surplus. However, as Van Damme, Selten, and Winter [1990] shows, any split can be supported if the offer space is discrete. A similar reasoning is at play here.

which is a subset R of $\mathcal{N}$ that the agent is allowed to choose from. [8] If the agent's type is S, she can either implement a project $i \in R \cap S$, in which case the game ends, or not implement anything and reject R altogether. In this case the game moves to the next period, and the principal offers another restriction set.[9]

The complete information counterpart of this new game, where it is common knowledge that the agent's type is mixed, *also* has an equilibrium where, along the equilibrium path, the good project is implemented at $t = 0$. The strategy of the principal is to set $R = \{g\}$ at any history, and the agent's strategy is to always accept any project that's allowed. It is easy to verify that these strategies constitute an SPE.

However, it turns out that in the new principal-offer game, this complete information equilibrium cannot be exploited! When the separating mechanism is optimal, there is *no* equilibrium of this game that attains the optimal commitment payoff. This is surprising, since we might think that the principal having control of proposals means that he can exert greater influence over what is implemented. But there is a trade-off between control and sequential rationality here.

To understand the intuition behind this, suppose the principal attempts to replicate the separating mechanism by setting $R = \{g\}$ at all $t < t^*(\delta)$. At $t^*(\delta)$, he sets $R = \{g, b\}$. So the mixed type indeed finds it optimal to accept g at $t = 0$. But this strategy of the principal is not sequentially rational; at $t = 0$, if g is not accepted, the principal *knows* it is the bad type, and therefore would permit b in the very next period. The sequential rationality constraint actually prevents him from waiting till $t^*(\delta)$ before allowing the bad project. Also, since non-acceptance of g is on-path, there is no scope for the principal to form extremal off-path beliefs, and leverage the equilibrium of the complete information game. The limited action set of the agent therefore hurts the principal, and makes *punishment through beliefs* harder to achieve. In fact, Li [2022] establishes that this Coasian force is prevalent in any equilibrium of this alternative game, and the unique equilibrium outcome is the principal setting $R = \{g, b\}$ at $t = 0$.

So, there are various forces that work here and it is their combination that helps with the attainment of commitment payoff in an equilibrium of our game.

[8] R can be $\varnothing$, in which case, the agent has no choice and the game proceeds to the next period.

[9] We show in Section 4.1 that the commitment benchmark we established also serves as an upper bound for what the principal can achieve by committing to a strategy in this alternative game.

1.4 A General Model with N projects

When there are two possible projects, we showed that the payoff from the static commitment benchmark can always be achieved in an equilibrium of the sequential delegation game. However, when we move beyond two projects, it is not clear if the commitment payoff can be achieved in equilibrium.

There are potentially two reasons why this could happen: (i) there is a wedge between the static stochastic mechanism and what can be achieved by commitment to a strategy in the dynamic game, or (ii) a sequential rationality consideration that impedes the principal from achieving his dynamic commitment payoff in an equilibrium of this game. We want to differentiate between these two forces. If commitment payoff in the sequential delegation game cannot be achieved, is it because of a limitation of the extensive form, as in (i), or is it because the principal cannot commit to his strategy, as in (ii).

Even with two projects, we might wonder if the principal can do better if the principal and the agent used an alternative extensive form to solve the joint decision problem of implementing a project. We therefore also want to explore the merits and demerits of the sequential delegation game, relative to other extensive forms. We want to understand if our sequential delegation game has any *limitations* that other extensive forms don't. But to answer this question, we need to define precisely what the general space of extensive forms is, and how we are comparing them.

In this section, we consider a general setting where there are $N > 2$ possible projects, and the interaction between the principal and the agent may take one of several possible forms. The set of all *possible* projects is now $\mathcal{N} \equiv \{1, 2, \dots, N\}$ where project i corresponds to (α_i, π_i). As before, only the agent knows which projects are *available*; her *type* is $S \subseteq \mathcal{N}$ representing the set of available projects. The agent's type is drawn from $\mathcal{S} \equiv 2^{\mathcal{N}}$ according to the probability distribution $\mu : \mathcal{S} \to [0, 1]$. We refer to the different extensive forms we consider as *delegation protocols*.[10]

We first define a general class of delegation protocols. We then show that the class of mechanisms that we defined in Section 3.1 serves as a general commitment benchmark in this setting. For this, we prove a Revelation Principle and show that any outcome that can be achieved in any extensive form can be achieved by a direct, static, and stochastic mechanism with type-dependent message spaces. Furthermore, we show that

[10]We can think of them as capturing the different institutional settings that the principal and the agent might interact in. The sequential delegation game is one possible delegation protocol. Another possible protocol is the principal specifying the set of projects from which the agent is permitted to choose in each period.

for the sequential delegation game, this upper bound is tight. More precisely, for any static mechanism, there exists a commitment strategy of the principal in the sequential delegation game that attains the same payoff as $\delta \to 1$. So, the sequential delegation game imposes no constraints *beyond* sequential rationality. Finally, we provide an example of a protocol where, even with the ability to commit to a strategy, the principal is not always able to achieve the optimal commitment payoff from a static mechanism. This highlights that some protocols may impose constraints beyond sequential rationality.

1.4.1 Protocols

We define a delegation protocol to be an extensive form game that specifies the proposer, and what they are allowed to offer at any history. Formally, at any history h_t, the the protocol specifies the proposer, $P(h_t)$ and the set of permissible offers, $\mathcal{O}(h_t) \subseteq 2^{\mathcal{N}}$, so any offer $O(h_t)$ is a subset of $\mathcal{N}$.[11] When an offer is made by the proposer, the other party responds by either accepting a project in the offer or rejecting the offer altogether. A history is a sequence of offers that have been rejected. [12]

The set of actions feasible for the agent at any history is type-dependent; if she is the proposer, she can only include available projects in her offer, and if the principal is the proposer, she can only accept an available project. Formally, if, at h_t, the proposer is the agent, and the agent's type is S, then we must have $O(h_t) \subseteq S$. On the other hand, if the proposer is the principal, the offer $O(h_t)$ can be any subset of $\mathcal{N}$ that's in $\mathcal{O}(h_t)$, but the agent can only accept a project in $O(h_t) \cap S$.

A delegation protocol is therefore simply a dynamic game with type-dependent action space for the agent at any history. If any project from an offer is accepted, the game ends, and players get their discounted payoffs; otherwise the game proceeds to the next period. Our equilibrium concept is Perfect Bayesian Equilibrium; both players play sequentially rationally and the principal's beliefs about the agent's type are updated according to Bayes' rule whenever possible.

[11] An offer can be $\varnothing$, this can be interpreted as the party making the offer permitting no project to be chosen by the other, like the agent choosing to stay silent in the sequential delegation game. If the offer is $\varnothing$, then the party responding to the offer has no option but to reject it, as there is no project to accept.

[12] Tying this definition back to our sequential delegation game, it is the protocol where at any h_t, $P(h_t)$ is the agent, and the set of permissible offers is $\mathcal{O}(h_t) = \mathcal{O} = \{\{i\} | i \in \mathcal{N}\}$, i.e. the singleton subsets of $\mathcal{N}$.

1.4.1.1 Revelation Principle

We consider the class of mechanisms defined in Section 3.1, where the message space is type-dependent, so a type can report only subsets of her available projects. A mechanism q maps any report to a probability of implementation of each project in that report. These are mechanisms with evidence, as each type is only able to report a subset of the projects she has. The message space here satisfies the *normality* condition from Bull and Watson [2007]

We now prove a Revelation Principle for this setting; we show that any social choice function implementable in any protocol is also implementable by a mechanism in this class. We first define a social choice function and an *induced* social choice function.

Definition 3. *A social choice function (SCF) is a function f that maps a set $S \subseteq \mathcal{N}$, to a probability of implementation of each project in S, where $f_S(i)$ denotes the probability of implementing project $i \in S$ from type S.*

An *induced* social choice function is defined to be an SCF that is induced by a strategy of the principal in a protocol and a best response to that strategy. In order to better understand what an induced SCF is, fix a protocol. Consider any strategy of the principal in this protocol and any best response of the agent to this strategy. [13] This pair of strategy and best response induce a probability distribution over outcomes for each type S, where an outcome (i, t) denotes project i being implemented at time period t. For any type S and project $i \in S$, we can condense the discounted probabilities of implementing i at different histories into a single probability, and this probability is denoted by $f^I_S(i)$.[14] The induced SCF is then the function f^I that maps any type S to a probability $f^I_S(i)$ of implementation of each $i \in S$.[15]

Proposition 2. *For any delegation protocol and any SCF f^I induced by a strategy of the principal and a best response of the agent, the static mechanism f^I is incentive compatible.*

Proof. Fix a delegation protocol and an induced SCF f^I in the protocol. Recall, from our description of protocols, that at any history, any action that is available to a type

[13] This pair does not need to constitute an equilibrium, we can think of the principal as being able to commit to a strategy.

[14] For example, if $i \in S$ is implemented with probability $\frac{1}{2}$ at $t = 0$, and with probability $\frac{1}{2}$ at $t = 1$, then $f^I_S(i) = \frac{1}{2} + \delta\frac{1}{2}$.

[15] The details of collapsing the probability of various outcomes involving i into a single probability can be found in the Appendix A.1.3.

S' is also available to S, where $S' \subseteq S$. This is because within the constraints of what the protocol permits, anything that S' can propose or accept, S can as well since S has all the projects S' has. Since the induced SCF comes from a best response of the agent, it must be that the payoff for S from f^I_S is weakly better than the payoff from $f^I_{S'}$, which is what S would get if she imitated the best response of S'. Thus, the incentive compatibility in the mechanism, which requires that no type should find it optimal to report a strict subset, is satisfied. □

This result highlights an important point. In the two-project case, he principal can always attain the commitment payoff in an equilibrium of the sequential delegation game. The above Revelation Principle tells us that this commitment benchmark in fact serves as an upper bound for what the principle can achieve with commitment, and therefore in an equilibrium of a very general class of protocols. So, if the principal lacks commitment, and we have two projects, there is no equilibrium of any other protocol in this class that can do better than the principal-optimal equilibrium of the sequential delegation game.

Note that the standard Revelation Principle from Bull and Watson [2007] does not follow directly here as we do not start out with a fixed evidentiary structure under which we compare various static and dynamic mechanisms. Instead, we start out with type-dependent evidentiary actions, which can be taken at multiple nodes. Moreover, the proposal of the agent also limits what the principal can choose when the agent is the proposer, which is not a feature of standard mechanisms with evidence.

We now argue that as $\delta \to 1$, any SCF that is implementable in a static, stochastic mechanism is implementable in the sequential delegation game if the principal is able to commit to a strategy. It means that our sequential delegation game is an optimal protocol under commitment and sequential rationality is the only restriction it imposes on what is attainable in equilibrium.

Theorem 2. *Fix a social choice function f. There exists a strategy of the principal and a best response of the agent in the sequential delegation game such that the induced SCF from this strategy and best response approaches f as $\delta \to 1$.*

We provide the proof in the Appendix, but the idea is to fix an SCF f and construct a corresponding strategy in the sequential delegation game:

* According to this strategy, the first N time periods, $t \in \{0, 1, \dots N-1\}$, are reserved for *information elicitation* where the agent proposes the available projects in a particular order. Proposals in the first N periods are analogous to a *report* in the

mechanism, and since the agent can only propose projects she has, this captures the fact that a type can only report her subsets in the mechanism.

* In time periods $t > N - 1$, the projects that were proposed in the first $N - 1$ periods are proposed again, and accepted with probabilities such that the agent finds it optimal to report all her projects in the first N periods. As $\delta \to 1$, these probabilities approach the implementation probabilities from the mechanism and the principal's payoff from this strategy approaches his payoff in the mechanism.

This result tells us that the sequential delegation game imposes no constraints as $\delta \to 1$ beyond sequential rationality. So, if the principal can commit to a strategy, but is constrained by the extensive form of the sequential delegation game, i.e. has to operate in a setting where the agent makes proposals, and he can merely respond to them, then this is not really a constraint. Even if he could choose the extensive form from a very general class of extensive forms, and *then* also commit to his strategy in that extensive form, he cannot do any better.

In contrast, there exist delegation protocols where the ability to commit to a strategy may not be enough to attain the commitment payoff from the optimal static mechanism, which is the common commitment benchmark for all the protocols we consider.

Consider the delegation protocol where, in each period, the principal makes proposals by choosing a restriction set, which is any subset O of $\mathcal{N}$ that the agent is allowed to choose from. So, at any history h_t , $P(h_t)$ is the principal, and $\mathcal{O}(h_t) = \mathcal{O} = 2^{\mathcal{N}}$. If the agent's type is S, she can either implement a project $i \in O \cap S$, in which case the game ends, or not implement anything and reject O altogether. In this case the game moves to the next period, and the principal offers another restriction set.

Now consider the example where there are three possible projects, $\mathcal{N} = \{1, 2, 3\}$, and three equally likely types in the support of μ with $\mathcal{S} = \{\{1, 2\}, \{2\}, \{2, 3\}\}$. The payoffs are:

$$\pi_1 = 8, \pi_2 = 3, \pi_3 = 1$$

$$\alpha_1 = 3, \alpha_2 = 8, \alpha_3 = 9$$

The optimal mechanism is as follows:

- From type $\{1, 2\}$, project 1 is implemented with probability one
- From type $\{2\}$, project 2 is implemented with probability $\frac{3}{8}$.
- From type $\{2, 3\}$, project 2 is implemented with probability one.

We can show that there does not exist a commitment strategy for the principal in this protocol that would attain the payoff from the above mechanism.

The details are in the Appendix, but the intuition is as follows: The construction of the commitment strategy in the sequential delegation game (Theorem 2) involves eliciting information from the agent about her type through her proposals and conditioning future responses on these initial proposals. Consider the following strategy of the principal: he rejects any proposal except 1 at $t = 0$. At $t = 1$, he accepts project 2 *only* if project 3 was proposed at $t = 0$, otherwise rejects 2 forever if it is proposed before a certain threshold t^*.[16]

In the alternative delegation protocol, the agent can only accept or reject, and not propose projects herself, and it limits the scope for information elicitation. Without the ability to condition future implementation of projects on the agent's own past proposals, the principal cannot separate type $\{2, 3\}$ from $\{2\}$.

1.4.2 Attaining the Commitment Payoff with N Projects

We now turn our attention back to the sequential delegation game and the possibility of attaining attaining the commitment payoff in equilibrium here. It is natural to ask whether our main result holds beyond two projects and we find that it is not clear that this would always be the case.

While the number of possible projects does not alter the game itself or the commitment benchmark, the problem significantly more complex. Even solving for the optimal mechanism is difficult as it now includes IC conditions for each subset of each type. As a result, the problem loses its tractability. In order to recover some of the lost tractability, focus attention on a restricted class of parameters. More specifically, we consider the model under three assumptions about the payoffs of the projects and the types in the support of μ.

We show that these assumptions are sufficient conditions for existence of an equilibrium of the game that attains the commitment benchmark. We also show through an example that the conditions are not necessary and highlight another signaling opportunity for the agent by proposing redundant projects.

Assumption 1. *(Conflicting preferences) The set of projects $\mathcal{N}$ satisfies*

$$\pi_1 > \pi_2 > \ldots > \pi_{N-1} > \pi_N > 0;$$

[16] The construction is similar to the separating equilibrium from the two-project case.

$$\alpha_N > \alpha_{N-1} > \ldots > \alpha_2 > \alpha_1 > 0.$$

We start by assuming that the set of projects is such that the preferences of the players are diametrically opposed. When there are two projects, the only alternative to opposite preferences is identical preferences in which case the problem would be trivial. Beyond two projects, however, there is an array of possibilities for conflicting preferences. Under Assumption 1, the principal and the agent have exactly opposing preferences.

Assumption 2. *(Linear payoffs) Any two projects $i, j \neq 1$ satisfy*

$$\frac{\pi_1 - \pi_i}{\alpha_i - \alpha_1} = \frac{\pi_1 - \pi_j}{\alpha_j - \alpha_1}.$$

We further simplify the complex incentives beyond two projects by Assumption 2 which requires all possible projects to lie on a line on $\mathbb{R}^2_{++}$ (Figure 1.4).

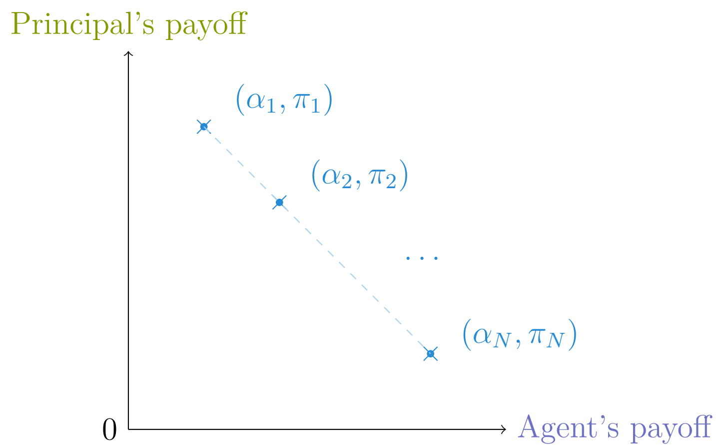

Figure 1.4. The project space with $N > 2$ projects under Assumptions 1 and 2. We see the conflicting preferences of the principal and agent, and the linear payoffs of the projects.

Assumption 3. *(Nested types) The probability distribution over the agent's types μ is such that for any $S, S' \in \mathcal{S}$ with $\mu(S), \mu(S') > 0$, either $S \subseteq S'$ or $S' \subseteq S$.*

Assumption 3 requires the set of possible types to be nested in a way that a type of the agent is either a subset or a superset of any other type. This assumption provides a structure to possible types and simplifies the incentives. Under Assumption 3, there can be at most one type with n projects for each $n \in \{1, 2, \ldots, N\}$.

When the parameters $\mathcal{N}$ and μ satisfy Assumptions 1, 2, and 3, we refer to this restricted type space with the restricted payoff structure as *nested linear type space*. The

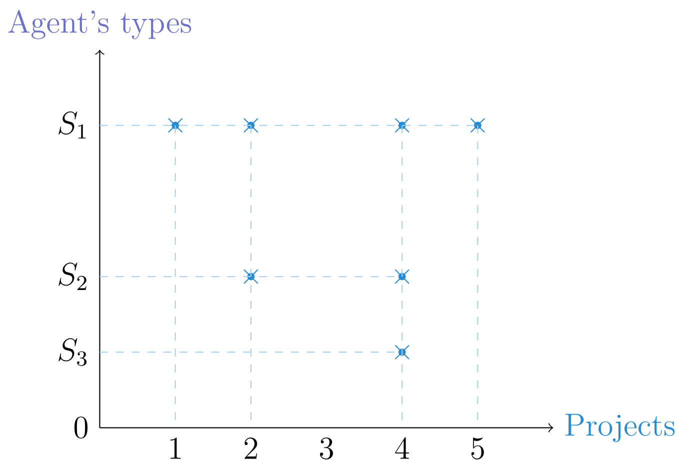

Figure 1.5. The type space $\mathcal{S}$ with $N > 2$ projects under Assumption 3 where project i refers to (α_i, π_i). The type space is nested such that if we take any two types, one would be a subset of the other.

nested linear type space reduces the number of incentive compatibility constraints to at most $(N - 1)$, simplifying the problem significantly (Figure 1.5).

Under these regularity conditions provided by Assumptions 1, 2, and 3, our main result extends to the general model, and the commitment payoff is always attainable in an equilibrium of the sequential delegation game.

Theorem 3. *In the nested linear type space, there exists an equilibrium of the sequential delegation game that attains the principal's commitment payoff as $\delta \to 1$.*

The main idea behind the proof is that we can divide solving for the principal-optimal mechanism into two parts. We first establish that in any optimal mechanism, each type's each possible report generates the same expected payoff v for the agent. Then, we can solve the optimization problem for a fixed value of v for each type. Combined with the fact that the differences between payoffs are linear, the optimal mechanism takes a very clean separating structure. The optimal mechanism can then be replicated in equilibrium with similar strategies as in the separating equilibrium in the two project case.

We can show with an example that our result for the nested linear type space is not tight: there are examples of type spaces outside this class where the commitment payoff can be achieved in equilibrium. Recall the example from Section 1.4.1.1 where there are three possible projects, $\mathcal{N} = \{1, 2, 3\}$, and three equally likely types in the support of μ with $\mathcal{S} = \{\{1, 2\}, \{2\}, \{2, 3\}\}$.

Note that we are outside the linear nested type space introduced in the previous section as the types are not nested and the payoffs are not linear. This type space can be thought of as augmenting our two project case with a type where the bad project is paired with an even worse project. Recall the optimal mechanism:

- From type $\{1, 2\}$, project 1 is implemented with probability one
- From type $\{2\}$, project 2 is implemented with probability $\frac{3}{8}$.
- From type $\{2, 3\}$, project 2 is implemented with probability one.

The structure of an equilibrium that attains the payoff from this mechanism is similar to the separating equilibrium but it exhibits a novel signaling opportunity. The principal always accepts project 1 and never accepts project 3. If project 3 is proposed at $t = 0$, then project 2 is accepted with certainty at $t = 1$. Otherwise, project 2 is only accepted with a delay at $t^* = \min\{t|\delta^t \leqslant \frac{3}{8}\}$. We should highlight that even though project 3 is never implemented, its proposal acts as a screening device and the agent has an opportunity to signal her type by proposing redundant projects.

1.5 Conclusion

In this paper, we study a dynamic principal-agent problem where the agent is privately informed about the feasibility of projects, and the interests of the parties are not aligned. Our main focus is on a dynamic delegation game where the informed agent makes proposals over time and the uninformed principal has the authority to approve without the power to commit to his future responses.

We ask how much of a disadvantage the principal is at here; he lacks proposal power and the ability to commit to his responses to the agent's proposals. Since the principal can only implement the projects that are proposed, we might expect that that the agent can easily *hide* principal-preferred projects by never proposing them. Anticipating that his preferred projects may never be proposed, the principal would in turn capitulate and accept the agent-preferred projects when they are proposed. We show, however, that with two projects, there is always an equilibrium of the game that attains the optimal commitment payoff. We argue that it is in fact the inability to make proposals that enables the principal to *wait*, and for costly delay to emerge as a signalling device in equilibrium. For more than two projects, we identify sufficient conditions on parameters under which the commitment result still holds.

Our setup has natural applications to organizational economics, specifically empire-building by corporate managers. It is well known that managers may not always act in the best interests of the shareholders, but rather act to increase their own influence within the organization. Our analysis has implications for when the manager holds verifiable private information that's relevant to the optimal course of action for the firm, but is motivated by empire-building. We find that the by adopting a bottom-up approach, i.e. eliciting proposals from the manager, the CEO might be able to curb the manager's empire-building plans better than by issuing top-down commands to restrict what a manager can do.

We also define a general class of delegation protocols and show that if the principal is able to commit to his strategy, then the sequential delegation game we consider does as well as any other protocol. This comparison with other protocols highlights an important point. In the two project case, this commitment payoff is achieved in an equilibrium of the sequential delegation game. Therefore there is no equilibrium of any other protocol in our class that results is a strictly higher payoff for the principal. This comparison further reinforces our intuition about the bottom-up approach — even if the principal cannot commit, organizational structures that involve a bottom-up approach might do better than a whole range of other organizational structures.

Chapter 2 How Markets Disrupt Mediated Trade

2.1 Introduction

Adverse selection often causes markets to function inefficiently, and if severe, can result in market breakdowns, as has been known since Akerlof [1970]. The key idea is that asymmetric information about the quality of the asset being traded may prevent mutually beneficial trade, even if it is common knowledge that there are gains from trade.

One way to get *around* this problem of market breakdown is by using an intermediary, namely someone who can help broker trade between sellers (privately informed parties), and buyers (the uninformed party). Examples of such intermediaries include real estate agents, representing someone looking to sell a house, and stockbrokers in financial markets. These intermediaries can create a more efficient channel for sale by reducing the transaction costs associated with the sale, advertising, and even negotiating with potential buyers on behalf of the seller.

But how might such intermediaries solve the problem of adverse selection? The idea is that an intermediary can screen seller's types by offering her a a menu of options. The choices in the menu cause different *types* of the seller to *self-select* into different categories, and this *separation* helps mitigate adverse selection. To understand this, consider the real estate example. Suppose the seller has two options: either she could *only* consider selling at a high price, in which case, there is a chance that the property will remain unsold. Or, she could consider selling at a low price, it which case, it would almost definitely get sold. So the options involve the following trade off: *at a high price,* *the probability of trade is lower.*[1]

Different types of the seller evaluate this trade off differently. A seller with a high quality property might derive a a high benefit from living there herself, if a sale does not happen. So, she would be unwilling to sell at a low price. On the other hand, if the seller knows that there are are issues with her property (that might not be immediately observable to others), then she just wants to ensure a sale, even if its at a low price. So, in equilibrium, a high quality seller would choose the first option and a low quality seller would choose the second. The price at which the property is offered, therefore, acts as an endogenous signal of quality for a buyer who might otherwise not be able to observe all aspects of it perfectly. Since *only* the high quality properties are listed at higher prices, the buyer is willing to buy at these high prices. So, high quality sellers sell with positive probability in equilibrium, and the intermediary is able to mitigate the problem of adverse selection.[2]

So, if an intermediary is operating alone, it can indeed successfully avert a market breakdown. But these intermediaries do **not** typically operate in a vacuum, and the seller often has other ways to sell. For e.g., if the intermediary is unable to broker trade, or the seller is unhappy with the terms offered by the intermediary, she can just sell on her own, on a decentralised market, *without* going through the intermediary.[3]

Given that the seller has the option to sell through an intermediary or directly on the market herself, it's unclear the degree to which intermediation can address adverse selection. This question motivates my analysis:

> *If the seller can also sell **without** the intermediary, are market breakdowns still preventable?*

I find that the answer is **no**. When intermediation is *not* exclusive, there is a

[1] This trade off might arise, for example, because the intermediary has to "search" for a buyer; at a higher price, it might be more *difficult* to find an appropriate buyer, and with some probability, the intermediary might **not** find one. Or, it might be more *costly* to find a high value buyer, and the intermediary might give up at some point.

The interpretation that I go with in this paper however, is that the the intermediary can **commit** to a lower probability of trade for the purpose of screening. The exact interpretation is not important; the idea is that at a higher price, there is some *uncertainty* about whether or not sale will take place, and this uncertainty is what enables screening. I endow the intermediary with commitment power to make it as powerful as possible, so that it can introduce *any* trade off.

[2] Again, one might wonder why the property gets sold with a probability lower than one at the higher price, if the buyer is always willing to buy. As mentioned previously, the interpretation I go with is that intermediary can *commit* to a lower probability of sale, to create a trade off for screening.

[3] Alternatively, the intermediary could represent the "legitimate" channel of sale, and if the seller is unable to sell through the intermediary, she might try to sell on a *black market*.

complete breakdown of trade, where only the lowest quality asset is traded in *any* equilibrium.

I consider a setting with *interdependent values* and severe adverse selection. There is a seller and a buyer. The seller has one unit of an indivisible good for sale. The quality of the good is the seller's *type*, and neither the intermediary, nor the buyer observes it. However, the distribution of types is common knowledge. The seller and the buyer have interdependent values for the good, and the buyer always values the good more than the seller. Even though the seller and buyer commonly know that there are gains from trade, the market suffers from a *lemons problem*: given the prior, the highest price that the buyer is willing to pay for the good is strictly lower that the reservation utility for the highest type of the seller.

The seller has two ways of selling the good: she can either sell through the intermediary, or on a *static competitive market*. The intermediary offers a menu of options, where each option is a tuple (π, p). If the seller chooses (π, p), with probability π, she will get a chance to sell through the intermediary, at price p. An option could, therefore, have some uncertainty associated with it; if the seller chooses it, she may or may not get a chance to sell through the intermediary. For example, in (π, p), if $\pi < 1$, then, with probability $1 - \pi$, the seller will *not* get the chance to sell through the intermediary. On the market, sale happens at a single market clearing price p_M, which is determined in equilibrium.

The **timeline** is as follows: First, the intermediary offers a menu, and the seller chooses an option in that menu. Then, the uncertainty associated with this option is resolved: the seller learns whether or not she has the option to sell through the intermediary. Finally, the seller decides whether to sell through the intermediary (if she has this option), or on the market. In particular, if the seller chooses an option and does *not* get the option to sell through the intermediary, she can, at this point, sell on the market.

The menu offered by the intermediary, therefore, induces a game where the seller chooses an option in the menu, and then decides where to sell. I study the Perfect Bayesian Equilibria of this game. Here, the intermediary has no direct control over the market. In equilibrium however, through the menu that it offers, it influences what types of the seller sell on the market, and therefore the market price. The market price, on the other hand, *endogenously alters* how the seller evaluates different options in the menu, and therefore her choice of allocation.

I find that this equilibrium interaction between the market and the intermediary can completely destroy the intermediary's ability to screen, and therefore its ability to mitigate adverse selection. My main result, stated informally, is as follows:

> **Main Result**. *When the intermediary operates alongside the market, under some condition on parameters, the unique equilibrium outcome is a total breakdown of trade, where only the lowest type trades in equilibrium.*

Hence, the market may completely destroy the efficiency gains that come from intermediated trade. I offer a necessary and sufficient condition for such a breakdown to occur, which I refer to as the *Bottom Lemons Condition (BLC)*. This condition means that there is a *lemons problem* for *every* subset of types at the bottom. So, if there are n possible types, consider any $k \leqslant n$. Conditional on the seller's type being in the set of the k *lowest types*, the buyer's expected value for the good is strictly lower than the reservation utility for the highest type in this set.

The key idea is that when the intermediary is operating *alongside* the market, there is no way to deter the lowest type from "mimicking" the higher types' choice of allocation. And under the BLC, preventing this mimicking is essential for higher types to trade in equilibrium. To see this, suppose that the lowest and the second-lowest type, both sell at the same price in equilibrium. By the BLC, there is a *lemons problem* for the two lowest types. So conditional on this price, the buyer's expected value for the good is strictly lower than the reservation utility of the second-lowest type. So the buyer wouldn't buy.

When there is no market, and the intermediary is the *only* channel of sale, the lowest type can be deterred from mimicking the higher types' choice of allocation by making the probability of trade in these allocations lower. If the lowest type chooses these "high-price" allocations, with some probability, she might not be able to sell *at all*. But when the market is *also* present, the seller can sell on the market, in case sale does not happen through the intermediary. Therefore the lower probability of sale in the high-price allocations is no longer an effective deterrent, and the lowest type finds is optimal to mimic the second lowest type's choice of allocation.

This mimicking, as I argued, implies that neither the lowest, nor the second-lowest type can trade in equilibrium. BLC ensures that this unravelling continues, because *any* subset of types at the bottom suffers from the *lemons problem*. If the third lowest type trades with positive probability, *both* the lowest and second-lowest types would mimic the third lowest type's choice of allocation. But then, buyer's expected value conditional on this allocation's price is strictly lower than the reservation utility of the third lowest price.

Section 2.2 illustrates the main forces at work through a two-type example. Section 2.3

describes the general model, and Section 2.4 contains the main results, including the general result about how the market disrupts trade. Section 2.5 concludes.

2.1.1 Literature Review

The literature on market breakdowns due to adverse selection was initiated by Akerlof [1970]. Akerlof considers a static, competitive market, where trade happens at a single, market clearing price. There is a large literature that takes a mechanism design approach to mitigating the breakdown problem, where there is an intermediary who screens different types of the seller by offering a menu of contracts. Notably, Samuelson [1984] and Myerson [1985] characterise surplus maximising mechanisms in a setting with lemons problem.

My paper *combines* the static competitive setting the mechanism design approach; there is an intermediary, who coordinates the sale of the object between the seller and the buyer, but there is *also* a Walrasian market in the background, and the seller always has the option to sell here. My main finding when the mechanism operates *alongside* the market, rather than replacing it, then the mechanism's ability to screen may be greatly disrupted. Another difference with Samuelson [1984] and Myerson [1985] is that I use a stronger notion of IR for the buyer; there, the buyer's IR is over the entire trading process, and the only requirement is that his ex ante expected payoff from participating in the mechanism needs to be non negative. I require IR to be satisfied at **every** price: for trade to happen at any price, the buyer should not anticipate a loss from trading at that price. This is the same as the *veto incentive compatibility* requirement in Gerardi, Hörner, and Maestri [2014].

The presence of the market alongside the intermediary also connects my paper to the literature on mechanism design with a "competitive fringe", started by Philippon and Skreta [2012], and Tirole [2012]. These papers study optimal government interventions to restore lending and investment in a market with adverse selection, where following government intervention, firms can raise funds in a static, competitive market. Like my paper, the market, and the mechanism offered by the government affect in other in equilibrium. Participation in the government's program signals private information and therefore endogenously affects the market, and the market in turn influences the decision to participate in the government's program. The setting, and the nature of intervention, however, is quite different from mine, and so are some results. In particular, in these papers, the government never benefits from "shutting down" the market, whereas in my setting, under some conditions, the market completely takes away the intermediary's ability to screen, so when these conditions hold, if "shutting down" the market was possible, it would be strictly optimal.

Another strand of literature that this paper is related to is that on *ratifiable mechanisms*, as in Cramton and Palfrey [1995] and Celik and Peters [2011]. In these papers, the outside option to the mechanism takes the form of a *game*. Players can either participate in (or "ratify") the mechanism, or reject it. If any player rejects the mechanism, all players play the game. The similarity with my paper is that the act of rejecting the mechanism conveys information about a player's type, and influences other players' beliefs about him when the game is played. The main difference from my work is that in these papers, the choice between the mechanism and the game is made *ex ante*, and if all agents choose the mechanism, they are *bound* to the mechanism. In my setting, the choice of whether to accept the intermediary's mechanism or not, is made at an *interim* stage, once the seller knows whether the option to sell through the intermediary exists or not. Another difference is that unlike these papers, I consider a setting with adverse selection.

My paper is also related to the literature that combines information design and mechanism design; examples include mechanism design with "aftermarkets", as in Dworczak [2020], conflict resolution as in Balzer and Schneider [2019], and the literature on sequential agency by Calzolari and Pavan [2006] and Calzolari and Pavan [2009]. Like my paper, the design of the mechanism influences what happens outside the mechanism. The difference is that in these papers, the the mechanism designer can choose to reveal some information elicited from the agent, to influence the *post-mechanism* outcome. In my paper, the intermediary cannot directly reveal any information to the market. It can only influence what the market learns about the seller's type *in equilibrium*, through its choice of menu.

2.2 A Two type Example: Complete Market Shutdown

There is a seller and a buyer. The seller has an indivisible good that she'd like to sell. The good's *quality* takes one of two possible values: θ_H, and θ_L, where $\theta_H > \theta_L > 0$, and the probability that the good is of quality θ_H is denoted by $\mu(\theta_H)$. The realisation of θ is the seller's private information, or her *type*. The distribution however, is common knowledge.

If the seller's type is θ, her reservation utility for the good is equal to θ, and the buyers' value for the good is $(1+\alpha)\theta$, where $\alpha \in (0,1)$ reflects the gains from trade. I assume that $(1+\alpha)\mathbb{E}[\theta] < \theta_H$. I refer to this as the *lemons condition*; it means that given the prior, the maximum price that a buyer is willing to pay for the good is strictly lower than the reservation utility of the high type of the seller.

There are two *channels* through which the seller can sell: an *intermediary*, and a *market*. The intermediary offers a menu of options to the seller. Each option in the intermediary's menu is a tuple (π, p), and is characterised by a price (p), and a probability of trade (π). The probability of trade represents the probability with which the seller gets the chance to sell through the intermediary, if she chooses this option.

When the seller chooses (π, p), with probability π, the intermediary offers the seller a chance to sell, and makes a *non binding* recommendation to the seller and the buyer: to the seller, it recommends "*sell at p*", and to the buyer, it recommends "*buy at p*". Then, if the seller decides to sell and the buyer decides to buy, sale happens, through the intermediary, *at price p*. With probability $(1 - \pi)$, the seller **cannot** sell through the intermediary.[4]

The other channel for sale is a *static competitive market*. The market has a single market clearing price, denoted by p_M, which is determined in equilibrium. The market has its own buyers associated with it, whose payoffs are identical to the buyer associated with the intermediary. So the intermediary is mediating trade between the seller and a *particular* buyer, and the market has its own buyers (at least two, so that there's Bertrand competition between them). The buyers cannot move across different channels of sale, and the buyer associated with any channel can only buy through *that* channel.

The **timeline** is as follows:

1. The intermediary commits to a menu.

2. The seller chooses an option from this menu.

3. The seller learns whether or not she has the option to sell through the intermediary. If the seller chose (π', p'):

    - with probability π', she gets the option to sell through the intermediary.

    - with probability $1 - \pi'$, this option does not exist.

4. The seller decides *where* to sell, if at all.

[4] So the intermediary menu is a *stochastic mechanism*. When the seller chooses (π, p), the intermediary *randomises* between offering and not offering the seller the opportunity to sell through the intermediary.

- If the seller *has* the option of selling through the intermediary, she decides between i) selling at p' through the intermediary, ii) selling at p_M on the market, or iii) not selling at all.
- If she does not have the option to through the intermediary, she decides between i) selling at p_M on the market, and ii) not selling.

5. If the seller is selling through the intermediary (resp. the market), the buyer(s) buy as long as *conditional* on sale happening at p' (resp. p_M) the buyer's expected value for the good is at least as much as the price.

Thus, the menu chosen by the intermediary induces a game where the seller chooses an option, and where to sell, and then the buyers choose whether or not to buy. I look at Perfect Bayesian Equilibria (PBE) of this game. In equilibrium, p_M is the market clearing price, given the distribution of seller's types that trade on the market in equilibrium. For example, if only the low type trades on the market in equilibrium, then p_M is $(1 + \alpha)\theta_L$. Finally, observe that I do not specify the objective function of the intermediary; this is because my focus is what is *feasible* in equilibrium for any menu, rather than on which menu is optimal given the intermediary's objective function.

Before getting to the main result of my model, which describes what happens when the intermediary operates alongside the market, it is useful to consider two benchmarks: i) the market is the only channel for sale, or ii) the intermediary is the only channel for sale.

Fact 1. *If the market is the only channel for sale, then the unique equilibrium outcome is that only the low type trades in equilibrium.*

This follows directly from the *lemons condition*. For the high type to be willing to sell, p_M has to be at least θ_H. But then, at such a price, the low type would *also* sell. So, conditional on p_M, the buyers' beliefs equal the prior, and by the *lemons condition*, their expected value for the good is strictly lower than θ_H. So we cannot have $p_M \geqslant \theta_H$ in equilibrium. Thus, the unique equilibrium outcome is that only the low type trades in equilibrium, and $p_M = (1 + \alpha)\theta_L$.

Fact 2. *If the intermediary is the* ***only*** *channel for sale, there exists a* $\pi_H \in (0, 1)$*, and a menu such that when the intermediary offers this menu, there exits an equilibrium where the low type trades with probability one, and the high type sells with probability* $\pi_H \in (0, 1)$*.*

Proof sketch: The main idea here is that the intermediary can screen through a trade off between price and probability of trade. I fix a menu, and argue that if the intermediary is operating in isolation, and offers *this* menu, then there exists an equilibrium in which the high type trades with positive probability.

Consider the menu $\mathcal{M}^* = \{\mathcal{L}, \mathcal{H}\}$, where $\mathcal{L} = (1, (1+\alpha)\theta_L)$, $\mathcal{H} = (\pi_H, \theta_H)$, and π_H satisfies:

$$\pi_H(\theta_H - \theta_L) = (1+\alpha)\theta_L - \theta_L \tag{2.1}$$

Observe that the *lemons condition* implies that $\theta_H > (1+\alpha)\theta_L$. This in turn implies that the price in option $\mathcal{L}$ is strictly lower than the price in option $\mathcal{H}$, and $\pi_H \in (0, 1)$.

So there are two options in the menu, one with lower probability of trade and higher price, and other with probability of trade one, but lower price. When this menu is offered, there exists an equilibrium in which on path:

- The low type chooses $\mathcal{L}$ with probability one.
- The high type chooses allocation $\mathcal{H}$ with probability one.
- If the good is being sold at either price $(1+\alpha)\theta_L$, or price θ_H, the buyer buys with probability one.

Given the buyer's strategy, the low type faces the following trade off: she can choose $\mathcal{L}$, and sell with probability one at price $(1+\alpha)\theta_L$, or choose $\mathcal{H}$, and sell at a strictly higher price θ_H, but with with probability $\pi_H < 1$. The expected payoff from choosing $\mathcal{L}$ is $((1+\alpha)\theta_L - \theta_L)$, and that from choosing $\mathcal{H}$ is $\pi_H(\theta_H - \theta_L)$. By 2.1, the low type is *indifferent* between these two options, so it is indeed optimal for her to choose $\mathcal{L}$. For the high type, since the price in option $\mathcal{L}$ is strictly lower than her reservation utility, it is optimal for her to choose $\mathcal{H}$.

For the buyer, if the good is being sold at price $(1+\alpha)\theta_L$, his equilibrium beliefs are that its the low type with probability one, so it is sequentially rational for the him to buy at this price. Similarly, since *only* the high type is selling at price θ_H, it is optimal for the buyer to buy at this price too. Therefore, in equilibrium,

- The low type trades with probability one, at price $(1+\alpha)\theta_L$.
- The high type trades with probability π_H, at price θ_H.

I now introduce the market, and describe what happens when the intermediary has to operate alongside the market.

Proposition 3. *If the intermediary and market* ***coexist****, then for every menu offered by the intermediary, the unique equilibrium outcome involves the low type trading with probability one and the high type trading with probability zero.*

The above results says that the market completely disrupts the intermediary's functioning. The breakdown occurs because *separation* of the low and high types is no longer possible. When the intermediary is operating alone, separation is achieved through a lower probability of trade for the high type: in $\mathcal{M}^*$, the low type is deterred from choosing $\mathcal{H}$ because $\pi_H < 1$. Therefore, in equilibrium, *only* the high type is selling at the higher price θ_H, so the buyer is willing to buy, and high type trades with positive probability in equilibrium. When the intermediary operates alongside the market, however, the seller can sell on the market, in case she chooses $\mathcal{H}$ and the option to sell through the intermediary does not materialise. Therefore the types can no longer be separated, and θ_H cannot trade in equilibrium.

I now illustrate the core logic of the proof using the menu $\mathcal{M}^*$. Recall that when the intermediary is operating alone, and offers $\mathcal{M}^*$, there there does exist an equilibrium where the high type trades wit positive probability.

Claim: *If the intermediary offers $\mathcal{M}^*$* ***now****, when it is operating alongside the market, there is* ***no*** *equilibrium in which the high type trades with positive probability.*

Proof Sketch: Firstly, observe that in any equilibrium, p_M has to be at least $(1+\alpha)\theta_L$, since the distribution of types conditional on the seller selling on the market cannot be worse than degenerate at θ_L. In fact, *when $\mathcal{M}^*$ is offered, in any equilibrium,* $p_M = (1+\alpha)\theta_L$.

Now suppose, by contradiction, that there *is* an equilibrium with menu $\mathcal{M}^*$ in which the high type trades with positive probability. Fix such an equilibrium. In this equilibrium, it must be that the high type is selling through the intermediary, at price θ_H. This is because the market price p_M, and the other price in $\mathcal{M}^*$, are both equal to $(1+\alpha)\theta_L$, which is strictly lower than θ_H. So, the high type cannot sell on the market or at the other price in the menu.

Therefore, in this equilibrium, the buyer's strategy must be to buy, if the good is being sold at price θ_H through the intermediary. Given the buyer's strategy, the low type *also* strictly prefers option $\mathcal{H}$ to $\mathcal{L}$. This is because by choosing $\mathcal{H}$, with probability $\pi_H > 0$ she can sell at price θ_H, and with probability $1-\pi_H$, when she doesn't have the option to sell through the intermediary, she can sell on the market at $p_M = (1+\alpha)\theta_L$. On the other hand, by choosing $\mathcal{L}$, her only option is to sell at price $(1+\alpha)\theta_L$, either through the intermediary, or on the market. But then, in equilibrium, *both* types will choose $\mathcal{H}$. So, conditional on the price θ_H, the buyer's beliefs equal the prior, and he will not buy. This contradicts that type θ_H is able to sell with positive probability in this equilibrium.

This reasoning extends more generally; there is **no** menu that the intermediary can offer such that there is an equilibrium with that menu where the high type trades with positive probability. As with menu $\mathcal{M}^*$, the idea is that now, in equilibrium, type θ_L is *guaranteed* a price of at least $(1+\alpha)\theta_L$ on the market. This destroys any separation that the intermediary can achieve through allocation probabilities, because now, the lower probability of trade in the allocation *meant* for the high type served is no longer an effective deterrent for the low type to *not* choose this allocation.

2.3 Model

I now develop a model of bilateral trade with an intermediary and a static competitive market.

2.3.1 Setup

I consider a bilateral trade setting with one seller and one buyer. The seller has one unit of an indivisible good for sale. This good has a quality θ associated with it, which is drawn from the set $\Theta = \{\theta_1, \theta_2, \ldots, \theta_n\}$, where $\theta_1 > \theta_2, \ldots > \theta_n > 0$, and θ is drawn according to the distribution $\mu(.)$. The realisation of θ is the seller's private information, but the distribution $\mu(.)$ is common knowledge. I refer to θ as the seller's *type*.

This is a setting with *interdependent values*; for a good of quality θ, the seller's cost of parting with the good is θ, and the buyer's utility from the good is $(1+\alpha)\theta$, where $\alpha \in (0,1)$. So, if the seller sells to the buyer at price p, the seller's payoff is $p - \theta$, and the buyer's payoff is $(1+\alpha)\theta - p$. Since $\theta_n > 0$, this implies that the gains from trade are always strictly positive. The seller and the buyer are risk neutral.

Assumption 4. ***Lemons condition****: The prior $\mu(.)$ satisfies the following condition: $(1+\alpha)\mathbb{E}[\theta] < \theta_1$.*

The above condition says that adverse selection is *severe*; given the prior, the maximum price that the buyer would be willing to pay for the good is strictly less than the cost of providing the good for the highest type.

There are two channels through which the seller can sell her good – an *intermediary* and a *static competitive market* ("market" henceforth). I assume that the buyer is not mobile, and though the seller can sell anywhere, a buyer can only buy through the intermediary. Moreover, I do not model the buyer(s) on the market explicitly. Rather, I assume there is a large number of potential buyers associated with the market, all identical to the buyer associated with the intermediary, and they can only purchase on the market.[5] Henceforth, "buyer" is used to refer to the buyer associated with the intermediary, unless I specify that we are taking about the buyer(s) on the market.

The intermediary offers a finite menu of allocations $\mathcal{M} \subseteq [0,1] \times \mathbb{R}_+$ to the seller. An element of the menu $\mathcal{M}$ is denoted by (π, p); if the seller chooses this allocation, with probability π, she will have the option to sell her good through the intermediary, at price p. Here, p is the price *conditional* on sale, the seller gets it only if a sale actually happens. The market has a single price p_M associated with it, where p_M is determined in equilibrium. Without loss, I assume that the menu offered by the intermediary must always contain the allocation $(0,0)$; this corresponds to the seller deciding to not sell through the intermediary.

2.3.2 Timeline

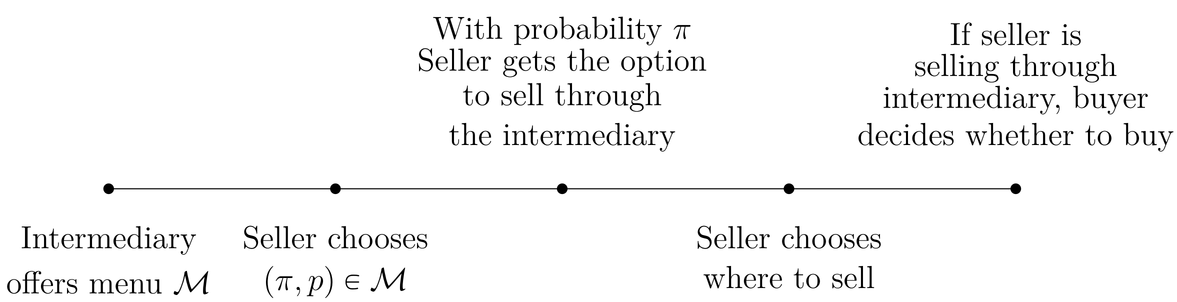

The menu chosen by the intermediary induces a game where the seller first decides which allocation to choose and then where to sell, and then, if the seller is selling through the intermediary, the buyer decides whether or not to buy (if the seller is selling on the market, the good always gets sold). look at the Perfect Bayesian Equilibria (PBE) of

[5] I make this assumption that the buyers are not mobile across channels to abstract away from the buyers' choice of where to trade, as incorporating this aspect would not only make the problem more complicated, but would also distract from my main focus

this game. In equilibrium, the market price p_M is determined by the market clearing condition; it is equal to the expected value of the good for the buyers on the market, where the expectation is taken with respect to the distribution of the seller's types that sell on the market in equilibrium. For example, if *only* type θ sells on the market in equilibrium, then $p_M = (1+\alpha)\theta$. I define the strategies, and equilibrium formally in the next subsection.

2.3.3 Strategies

Fix a menu $\mathcal{M}$ offered by the intermediary. I now define the strategies of the seller and buyer formally. I also define the buyer's beliefs and the beliefs on the market.

Seller's strategy: The seller's strategy has multiple components. It specifies i) the allocation that she chooses, as a function of her type, and ii) where she decides to sell, given her choice of allocation, her type, and whether the option to sell through the intermediary exists. The seller's strategy is given by $\sigma_S = (\sigma(.), \gamma_I(.), \gamma_M(.), \gamma_M'(.))$. I now describe each component individually.

The first component describes the seller's choice of allocation in $\mathcal{M}$, as a function of her type:

$$\sigma : \Theta \to \Delta\mathcal{M}$$

Here, $\sigma((\pi, p)|\theta)$ denotes the probability with which type θ chooses the allocation $(\pi, p) \in \mathcal{M}$. For any choice of allocation, the opportunity to trade through the intermediary may or may not realise. If the seller chooses an allocation, and *gets* the option to sell through the intermediary, she can either choose to sell through the intermediary, or on the market, or not sell at all. This choice is captured by the following two functions:

$$\gamma_I, \gamma_M : \Theta \times \mathcal{M} \to [0, 1]$$

Suppose the seller's type is θ, she chose allocation (π, p) in the menu, *and* has the option to put the good up for sale through the intermediary at price p. Then, $\gamma_I((\pi, p), \theta)$ denotes the probability with which she chooses to sell through the intermediary, at price p, and $\gamma_M((\pi, p), \theta)$ denotes the probability with which she chooses to sell on the market. Therefore, for any θ, and any (π, p), $\gamma_I(\theta, (\pi, p)) + \gamma_M((\pi, p), \theta) \leqslant 1$, where $1 - \gamma_I(\theta, (\pi, p)) - \gamma_M((\pi, p), \theta)$ is the probability with which the seller does not sell.

Lastly, the seller may not get the opportunity to sell through the intermediary. In this case, she can either decide to sell on the market, or not at all. This is the last component of the seller's strategy:

$$\gamma_M^{'} : \Theta \times \mathcal{M} \to [0,1]$$

Here, $\gamma_M^{'}((\pi,p),\theta)$ denotes the probability with which type θ chooses to sell on the market conditional on choosing (π,p) in the menu, and *not* having the option to sell through the intermediary. With probability $1-\gamma_M^{'}((\pi,p),\theta)$, she chooses not to sell at all.

Buyer's strategy: The buyer's strategy describes whether or not the buyer buys at any price p, if the seller is selling through the intermediary at this price. It is a function $\sigma_B : \mathbb{R}_+ \to \{0,1\}$, where $\sigma_B(p) = 1$ denotes that the buyer buys at price p. Observe that I rule out randomisation by the buyer, and restrict the buyer to pure strategies.[6]

I do not specify the strategies of the of the buyers on the market explicitly. There are a large number of potential buyers on the market and in any equilibrium, and in any equilibrium, the market price p_M is determined by the market clearing condition. So, the seller is always able to sell on the market at p_M with probability one, if she decides to do so.

Given the buyer's strategy σ_B, the maximum expected payoff for a seller of type θ, if she chooses an allocation (π,p) in $\mathcal{M}$ is given by:

$$\max_{\gamma_I,\gamma_M,\gamma_M^{'}\in[0,1]} \quad \pi(\gamma_I p \mathbb{1}_{\sigma_B(p)=1} + \gamma_M p_M) + (1-\pi)\gamma_M^{'} p_M \tag{2.2}$$

I denote the expression in 2.2 as $V((\pi,p),\theta|\sigma_B,p_M)$. Here, γ_I, γ_M, and $\gamma_M^{'}$ are shorthand for $\gamma_I((\pi,p),\theta)$, $\gamma_M((\pi,p),\theta)$, and $\gamma_M^{'}((\pi,p),\theta)$ respectively. $\mathbb{1}_{\sigma_B(p)=1}$ is the indicator function that denotes whether the buyer buys or not, if the seller decides to sell at p. Observe that it is without loss that the seller always chooses *some* allocation in the menu, because choosing to sell directly on the market can be captured by choosing the allocation $(0,0)$ in the menu.

Beliefs: I use $\mu_B : \mathbb{R}_+ \to \Delta\Theta$ to denote the buyer's beliefs about the seller's type as a function of the type, and for any p' such that $(\pi',p') \in \mathcal{M}$, $\mathbb{E}_B[\theta|p']$ denotes the expected value of the seller's type, *conditional* on the seller selling at price p', where

[6] Restricting the buyer to pure strategies is actually without loss, as any randomisation that the buyer might do can be built into the allocation probabilities in the menu offered by the intermediary. So, for (π,p) if the buyer randomises between buying and not buying at p, then the same equilibrium outcome is attainable by lowering π, and adjusting the buyer's strategy so that he buys with probability one.

the expectation is taken with respect to $\mu_B(p')$. I denote the beliefs on the market by $\mu_M \in \Delta(\Theta)$, this represents the beliefs of all buyers on the market about the seller's type, conditional on the seller deciding to sell on the market.

Assumption 5. *For any p, $\sigma_B(p) = 1$ if and only if $(1+\alpha)\mathbb{E}_B[\theta|p] \geqslant p$.*

So, I assume that at any price, the buyer buys, as long as given his beliefs about the seller's type conditional on the price, he *does not anticipate a loss* from buying. So he buys when indifferent.[7] I refer to this as the buyer's ***interim IR*** condition.

2.3.4 Equilibrium and Payoffs

The solution concept is Perfect Bayesian Equilibrium. Fix any menu $\mathcal{M}$. An equilibrium of the game induced by the menu is given by $(\sigma_S, \sigma_B, \mu_B, \mu_M, p_M)$, where

- Given σ_B and p_M, for any type θ, the seller's choice of which allocation to choose, and where to sell is sequentially rational. So, $\sigma((\pi, p)|\theta) > 0$ if and only if $V((\pi, p), \theta|\sigma_B, p_M) \geqslant V((\pi', p'), \theta|\sigma_B, p_M)$ for all $(\pi', p') \in \mathcal{M}$. Following any choice of allocation, the choice of if, and where to sell must be sequentially rational, given σ_B, and p_M.[8]

- Given μ_B, the buyer's choice of buying at any price p is sequentially rational; buyer buys if and only if the buyer's interim IR at that price is satisfied, i.e., $(1+\alpha)\mathbb{E}_B[\theta|p] \geqslant p$.

- Given the seller's strategy σ_S, μ_B is derived using Bayes Rule wherever possible.

- $p_M = (1+\alpha)\mathbb{E}_M[\theta]$, where the $\mathbb{E}_M$ denotes the Expectation taken with respect to $\mu_M(.)$, the equilibrium beliefs of the buyers on the market about the types of the seller that sell on the market, where $\mu_M(.)$ is derived from the seller's strategy using Bayes Rule, whenever trade takes place on the market with positive probability in equilibrium.[9] So, when trade takes place with positive probability on the market, $\mu_M(.)$ represents the equilibrium distribution of the types of the seller that sell on the market.

[7]This is again without loss. See footnote 5.

[8]The choice of if, and where to sell must be sequentially rational even if the choice of allocation is *off-path*, i.e., $\sigma((\pi, p)|\theta) = 0$.

[9]When trade takes place on the market with positive probability, all buyers on the market have the same beliefs, but I assume that even when trade takes place on the market with probability zero, all buyers on the market still have the same off path beliefs about the seller's type.

Outcomes and Payoffs: For any menu $\mathcal{M}$, let $P_{\mathcal{M}} = \{p | (\pi, p) \in \mathcal{M}\}$; this is the set of all the prices at which sale can possibly take place through the intermediary. An *outcome* of the game corresponding to $\mathcal{M}$ is a tuple (θ, i, p'), where $i \in \{I, M\}$. The tuple represents the outcome that the good of type θ was sold on i at price p', where i can be either through the intermediary (I), or the market (M). If $i = M$, then p' must be p_M. For a seller of type θ, his payoff from the outcome (θ, i, p') is $p' - \theta$, and the payoff for the buyer who purchased the good is $(1 + \alpha)\theta - p$.

An equilibrium induces, for any type, a probability distribution over outcomes, which is represented by $(\{\pi_p(\theta)\}_{p \in P_{\mathcal{M}}}, \pi_M(\theta))$, where $\pi_p(\theta)$ is the equilibrium probability that the good of type θ is sold through the intermediary at price p, and $\pi_M(\theta))$ is the probability that this good is sold on the market at p_M. For any θ, It must be that $\sum\{\pi_p(\theta)\}_{p \in P_{\mathcal{M}}} + \pi_M(\theta)) \leqslant 1$, where, if this sum is strictly less than 1, then this means that with some probability, type θ does not sell in this equilibrium. The seller and the buyers are risk neutral; the seller's expected payoff from $(\{\pi_p(\theta)\}_{p \in P_{\mathcal{M}}}, \pi_M(\theta))$ is $\sum_{p \in P_{\mathcal{M}}} \pi_p(\theta)(p - \theta) + \pi_M(p_M - \theta)$. For the buyer, at the time of buying, this expectation is taken with respect to μ_B, her equilibrium beliefs about the seller's type. On the market, by definition of p_M, any buyer who buys gets zero payoff.

Menus vs Direct Mechanisms: In my model, the intermediary can commit to menus, which are indirect mechanisms. I do not consider direct mechanisms that map reports in Θ to tuples (π, p), because direct mechanisms are *not* without loss here. This is because I am implicitly restricting attention to a special class of *deterministic* mechanisms: although the menus are stochastic in the sense that at any price, trade may or may not happen, each allocation in the menu consists of a *single* price, so the intermediary is not offering a menu of *lotteries* over prices.[10] As Strausz [2003] shows, when restricting attention to deterministic mechanisms, direct IC mechanisms may *not* be without loss. The main idea is that in the game *induced* by the mechanism, there are outcomes which are only attainable when the seller plays a mixed strategy. A mechanism that offers lotteries over prices can randomise *for* the seller, but my class of mechanisms cannot, so the standard Revelation Principle does not hold.[11] Section 2.4.5 contains a more in-depth discussion of what happens when the intermediary can offer lotteries over prices.

[10] I restrict attention to this class of mechanisms for tractability, but as I discuss in Section 2.4.5, the breakdown result with two types still goes through, even if I allow for lotteries over prices.

[11] One might wonder why this randomisation over prices can make a difference, since the seller and buyer are risk neutral. The reason is that in my model, the buyer's IR must hold for *every* price, and therefore randomisation over prices can help by allowing more flexibility in varying the information contained in any *given* price.

2.4 Main Results

With two-types, the market completely disrupts the operation of the intermediary. In this section, I study the market's effect on the intermediary's functioning more generally. But, before getting to the main results, I first describe how the intermediary screens when there is no market, and what changes when the intermediary is operating alongside the market. This will be useful for understanding the challenges to screening when the intermediary has to operate alongside the market.

2.4.1 Screening When There Is No Market

When the intermediary is operating in isolation, it screens through a trade-off between prices and probability of trade. In the intermediary's menu, allocations that have higher price have lower probabilities of trade.

Why does this trade off cause separation of higher and lower types in equilibrium? The key idea is that higher types have a higher cost of parting with the good. So, while comparing two allocations, they might find allocation with a lower probability of trade more attractive, because it involves a lower expected cost of parting with the good. In equilibrium, this results in higher types choosing allocations with higher prices and lower probability of trade.

To see this more clearly, observe that for any allocation (π, p), and any price θ, the expected payoff from choosing this allocation is $\pi(p - \theta) = \pi p - \pi\theta$, where πp is the *expected price from sale* in this allocation, and $\pi\theta$ is the *expected cost of parting with the good*. Now, consider allocations (π, p) and (π', p'), where $\pi < \pi'$, and $p > p'$. So, allocation (π', p') has a higher probability of trade, and a lower price. If type θ is comparing the two allocations, then:

$$\pi'(p' - \theta) \geqslant \pi(p - \theta) \Longleftrightarrow \pi' p' - \pi p \geqslant (\pi' - \pi)\theta \tag{2.3}$$

Therefore, the comparison between the allocations boils down to a comparison between $\pi' p' - \pi p$, and $(\pi' - \pi)\theta$. Here, $\pi' p' - \pi p$ is the difference between the expected prices in the two allocations, and $(\pi' - \pi)\theta$ is the difference between the expected cost of parting with the good. Suppose $\pi' p' - \pi p > 0$, so (π', p'), the allocation with the higher probability of sale and lower price, has higher *expected* price from sale than (π, p). But since $(\pi' - \pi)\theta > 0$, (π', p'), *also* involves higher expected cost of parting with the good.

From 2.3, we can see that for lower values of θ, the higher expected price dominates the higher expected cost, and they prefer (π', p') to (π, p). On the other hand, for higher values of θ, the effect of higher expected cost of parting with the good dominates, and they prefer (π, p). I now sum up this discussion in the following lemma:

Lemma 1. *If type θ is indifferent between allocations (π, p) and (π', p'), where $\pi < \pi'$, and $p > p'$, then any $\theta^{'} < \theta$ strictly prefers (π', p') to (π, p), and any $\theta^{''} > \theta$ strictly prefers (π, p) to (π', p').*

Proof. Follows directly from 2.3 □

2.4.2 Screening In The Presence of the Market

In this section, I describe how the market impacts the way the intermediary can screen. The trade off that enables screening remains the same: in equilibrium, higher types trade with lower probabilities, and at higher prices. However, the presence of the market implies certain constraints for the prices at which trade can happen through the intermediary in equilibrium. It also endogenously alters the reservation utility for certain types in equilibrium, thereby changing the way these types evaluate allocations in the intermediary's menu. I now state some lemmas about these equilibrium constraints that will be useful in understanding subsequent results.

Lemma 2. *In equilibrium, if the market price is p_M, then any trade that takes place through the intermediary must be at a price (weakly) greater than p_M. Moreover, in equilibrium, if **any** trade takes place through the intermediary, at a price strictly greater than p_M, then **all** trade through the intermediary must take place at a place at a price strictly greater than p_M.*

Proof. The first part is straightforward. Since, the seller can always sell on the market at p_M, in equilibrium, no type of the seller would sell through the intermediary at $p < p_M$.

For the second part, since there exists a p at which trade happens with positive probability in equilibrium, there must be an allocation (π, p), such that $p > p_M$, and $\sigma_B(p) = 1$, i.e., the buyer's strategy is to buy at p. So, by choosing (π, p), with probability π, the seller can sell at $p > p_M$ (and with $(1 - \pi)$, sell at p_M on the market). So, no type of the seller would choose an allocation (π', p') with $p' = p_M$. □

Lemma 2 is at the root of the breakdown result in the next section. The key idea is that now, types greater than p_M trading in equilibrium has an additional "cost" that was not present when there was no market: it means that all types less than p_M must *also* trade at prices strictly greater than p_M. This, combined with the fact that at any price, the buyer's interim IR constraint must also be satisfied, makes harder for higher types to trade in equilibrium.

I now argue that for types $\theta \leqslant p_M$, the presence of the market alters their reservation utility, while evaluating allocations in the menu.

Lemma 3. *Suppose the equilibrium market price is p_M. Then any two types, (weakly) lower than p_M, have the same "effective" reservation utility, and therefore have the same ranking over any two allocations.*

Proof. Consider allocations (π, p) and (π', p') in $\mathcal{M}$, such that $\pi < \pi'$, $p > p' > p_M$, and in equilibrium, the buyer' strategy is to buy at both prices p and p'. Then, for any $\theta \leqslant p_M$, the payoff from choosing (π, p) is

$$\pi(p - \theta) + (1 - \pi)(p_M - \theta) = \pi(p - p_M) + p_M - \theta$$

Similarly, the payoff from choosing (π', p') is $\pi'(p' - p_M) + p_M - \theta$. Therefore, for θ, the comparison between allocations (π, p) and (π', p') boils down to the comparison between $\pi(p - p_M)$ and $\pi'(p' - p_M)$. Exactly the same thing is true for any $\theta' \leqslant p_M$. So, all types lower than p_M have the same ranking over any two allocations.

□

Therefore, all $\theta \leqslant p_M$ evaluate choices in the menu *as if* their type is p_M, and the intermediary is operating in isolation. I refer to p_M as the *effective type* for all types $\theta \leqslant p_M$. I now use this fact to prove the following lemma:

Lemma 4. *Suppose, in equilibrium, an allocation (π, p) is chosen by some $\theta \leqslant p_M$, and by some $\theta' > p_M$. Then, in equilibrium, is is also chosen by $\{\theta | p_M < \theta < \theta'\}$.*

Proof. This follows directly from Lemma 1 and the notion of *effective type*. Fix any θ such that $p_M < \theta < \theta'$ (if such a θ exists). Then, θ strictly prefers (π, p) to all allocations (π', p') such that $\pi' > \pi$, and $p' < p$. This is because *effective type* p_M chooses, and therefore weakly prefers (π, p) to (π, p). Therefore, by Lemma 1, since $\theta > p_M$, θ strictly prefers (π, p) to (π, p). Similarly, we can argue that θ will not choose any allocation (π'', p'') such that $\pi'' < \pi$ and $p'' > p$, by using θ'. □

2.4.3 Main Result: Market Breakdown

With two-types, the presence of a market leads to a trading impasse. This leads to the question:

With more than two types, is there ever a complete breakdown of trade? If so, when?

The answer to the first part of the above question is *yes*; the impasse result holds more generally. I provide a ***necessary and sufficient condition*** on the prior for finitely many types under which, *only the lowest type trades in any equilibrium.* Recall that the set of types is given by $\Theta = \{\theta_1, \theta_2, \ldots \theta_n\}$, where $\theta_1 > \theta_2 \ldots > \theta_n$, and the probability of type θ is $\mu(\theta)$.

Definition 4. ***Bottom Lemons Condition***: *The prior $\mu(.)$ satisfies the Bottom Lemons Condition (BLC) if for any $k \in \{1, 2 \ldots n-1\}$, we have that $(1+\alpha)\mathbb{E}[\theta'|\theta' \leqslant \theta_k] < \theta_k$.*

Observe that when there are two types, BLC is *equivalent* to the *lemons condition.* With more than two types, this condition says that for *any* subset of types at the bottom, there is a *lemons problem.* Before stating the main result, I provide two examples to help understand the BLC better.

Suppose there are three possible types: $\theta_1 = 3$, $\theta_2 = 2$, and $\theta_3 = 1$, and $\alpha = 0.2$. So the buyer's payoff from a good of type θ is 1.2θ. Keeping α and the set of possible types same, I vary the prior to provide two examples: one where the prior satisfies the BLC, and another where it does not.

Example 1. *Prior satisfies BLC: $\mu(1) = \mu(2) = \mu(3) = \frac{1}{3}$.* First consider the two lower types: conditional on $\theta \in \{\theta_2, \theta_3\}$, the buyer's expected value for the good is $(1+\alpha)\mathbb{E}[\theta|\theta \leqslant \theta_2] = (1.2)(\frac{2+1}{2}) = 1.8$ which is strictly lower than θ_2. Now, consider the entire set of types. The buyer's expected value is $(1+\alpha)\mathbb{E}[\theta] = 2.4$, which is strictly lower than θ_1.

Example 2. *Prior does **not** satisfy BLC: $\mu(2) = \frac{3}{5}$, and $\mu(1) = \mu(3) = \frac{1}{5}$.* Now, $(1+\alpha)\mathbb{E}[\theta|\theta \leqslant \theta_2] = 2.8 > \theta_2$, so the BLC fails, since conditional on the lower two types, the buyer's expected value is strictly higher than θ_2. The *lemons condition* still holds though, as $(1+\alpha)\mathbb{E}[\theta] = 2.4 < \theta_1$. So, the prior violates the BLC, but still satisfies the overall *lemons condition.*

I now state the main result, which says that BLC characterises conditions under which total breakdown of trade is the *unique* equilibrium outcome.

Theorem 4. *The unique equilibrium outcome involves type θ_n trading with probability one, and all $\theta > \theta_n$ trading with probability zero if and only if the BLC is satisfied.*

Before getting to the sketch of the proof, I make the following observation: *BLC implies that types at the bottom cannot be "pooled".* To see this, suppose there is such pooling in equilibrium, i.e., there exists a $\overline{\theta} > \theta_n$, such that all $\theta \leqslant \overline{\theta}$ trade at the same price p' in equilibrium. Then, conditional on p', the buyer's expected value for the good is $(1+\alpha)\mathbb{E}[\theta|\theta \leqslant \overline{\theta}]$, which, by BLC, is *strictly* lower than $\overline{\theta}$. But since $\overline{\theta}$ is selling at p' in equilibrium, we must have $p' \geqslant \overline{\theta}$. This is a contradiction.

I will argue that any equilibrium where types $\theta > \theta_n$ trade, *must* have such pooling, thereby resulting in a contradiction. The key force behind the inevitability of pooling is that when the market is present, the lowest type can never be deterred from mimicking some higher type's choice of allocation, and everything unravels from here. I now go through the sketch of the proof in a series steps. The full proof is in the Appendix.

Proof Sketch:

Step 1: In any equilibrium, the price on the market is $p_M = (1+\alpha)\theta_n$, and ***only*** *θ_n can trade on the market in equilibrium.*

I do not provide a proof of why $p_M = (1+\alpha)\theta_n$ here, but I prove the second part taking this as given. BLC implies that $(1+\alpha)\theta_n$ is strictly lower than θ_{n-1}, the second-lowest type. Therefore, $(1+\alpha)\theta_n$ is strictly lower than any type greater than θ_n. Since $p_M = (1+\alpha)\theta_n$ in any equilibrium, only θ_n can trade on the market.

So, for a seller of type $\theta > \theta_n$, since the market price is lower than her reservation utility, if she trades in equilibrium, it must be through the intermediary. *Now suppose there is an equilibrium in which types greater than θ_n trade with positive probability.* Fix such an equilibrium. Then, the following is true:

Step 2: In equilibrium, θ_n mimics the choice of ***some*** *higher type.*

Suppose not, i.e., in equilibrium, θ_n chooses an option that's *not* chosen by *any* $\theta > \theta_n$. Let this option be (π', p'). Since, in equilibrium, *only* θ_n is choosing (π', p'), therefore the buyer only finds it optimal to buy if $p' \leqslant (1+\alpha)\theta_n$.

But it cannot be optimal for θ_n to choose an option with $p' \leqslant (1+\alpha)\theta_n$ in equilibrium. Recall that $\theta > \theta_n$ trade with positive probability in equilibrium, and these types can only trade through the intermediary. So, if they trade with positive probability in equilibrium, there exists a (π, p), where $p \geqslant \theta_{n-1} > (1+\alpha)\theta_n$, such that by choosing this option, the seller can sell at p with probability $\pi > 0$. This is a contradiction. Therefore, in equilibrium, θ_n must choose an option that's *also* chosen by some $\theta > \theta_n$.

Step 3: Step 2 implies that there must be ***pooling at the bottom****; there exists a type* $\overline{\theta} > \theta_n$*, such that* ***all*** $\theta \leqslant \overline{\theta}$ *choose the* ***same option*** *in equilibrium.*

Let the option chosen by θ_n in equilibrium be (π^*, p^*). By *Step 2*, there exists a $\theta' > \theta_n$, such that in equilibrium, θ' chooses (π^*, p^*) as well. Let $\overline{\theta}$ be the highest θ that chooses (π^*, p^*) in equilibrium. Since both $\theta_n < p_M$, and $\overline{\theta} > p_M$ choose (π^*, p^*) in equilibrium, by Lemma 4, we have that all types in the set $\{\theta | p_M < \theta < \overline{\theta}\}$ (if any), must *also* choose (π^*, p^*) in equilibrium. Therefore, the set of types that chooses (π^*, p^*) in equilibrium, is given by $\{\theta | \theta \leqslant \overline{\theta}\}$.

Step 4: Step 3 results in a ***contradiction****.*

By *Step 3*, there exists a $\overline{\theta} > \theta_n$, such that all $\theta \leqslant \overline{\theta}$ choose the same allocation in equilibrium. Let this allocation be (π^*, p^*), so the buyer's expected value for the good, conditional on price p^*, is $(1+\alpha)\mathbb{E}[\theta | \theta \leqslant \overline{\theta}]$. But by BLC, $(1+\alpha)\mathbb{E}[\theta | \theta \leqslant \overline{\theta}] < \overline{\theta}$, so the IR of $\overline{\theta}$ is violated at p^*, which contradicts that $\overline{\theta}$ chooses (π^*, p^*) in equilibrium.

This completes the sketch of the proof. If we start with an equilibrium where types greater than θ_n trade with positive probability, we reach a contradiction, so there can be no such equilibrium.

I only sketched the proof for sufficiency: the result in Theorem 4 is an *if and only if*, so if the BLC is *not* satisfied, then there always exists an equilibrium in which types greater than θ_n trade with positive probability. To see this, observe that when BLC fails, a complete breakdown of trade can be avoided even if there is *just* a market, and no intermediary. If BLC fails, there exists a $\theta' > \theta_n$ such that $(1+\alpha)\mathbb{E}[\theta | \theta \leqslant \theta'] \geqslant \theta$. Let $\overline{\theta}$ be the highest such θ. Then, if there is just a market, there exists an equilibrium, where the market price is $(1+\alpha)\mathbb{E}[\theta | \theta \leqslant \overline{\theta}]$ and all $\theta \leqslant \overline{\theta}$ trade with probability one. Of course, if there is an intermediary, this is trivially still an equilibrium as the intermediary can offer a menu $\mathcal{M} = \{(0,0)\}$.

BLC represents ***severe adverse selection***: it characterises the condition on the prior under which, if there is just a market, then trade breaks down to the lowest type. So the result in Theorem 4 essentially says the following:

If and only if adverse selection is so severe that trade breaks down to the lowest type in the absence of the intermediary, can the intermediary do nothing to avert this breakdown either, when it has to operate alongside the market.

Of course, when BLC is satisfied, we can do *strictly* better in terms of surplus from trade when the intermediary is operating in isolation. One trivial construction that allows this is the exactly menu that I constructed for the two type case; the menu offers exactly two allocations and only the lowest two types, θ_n and θ_{n-1} trade with positive probability in equilibrium. θ_n trades with probability one at price $(1+\alpha)\theta_n$, and θ_{n-1} with probability $\pi_{n-1} < 1$ at price θ_{n-1}. Of course, this might not be the surplus maximising menu, and we can have menus such that $\theta > \theta_{n-1}$ also trade in equilibrium. The main point is if there is no market, we can *always* do better than *just* θ_n trading with probability one. To sum up, see Table 2.1.

2.4.4 What if the BLC fails?

We saw in the last section that if adverse selection is "severe", as implied by BLC, then a complete breakdown of trade is unavoidable, with or without the intermediary. *But what happens when adverse selection is less severe, i.e. the BLC fails?*

As I argued earlier, in this case, we can always avoid a complete breakdown of trade. There are two possibilities: The first is that the presence of the market reduces surplus from trade, compared to when there is *just* an intermediary. So, although a complete breakdown of trade can be avoided, the market might still result in some inefficiency. The second possibility is that the presence of the market can sometimes help *improve* efficiency.

2.4.4.1 Market causes inefficiency

I provide a sufficient condition on parameters under which the market reduces the surplus from trade attainable in equilibrium, compared to when the intermediary is operating alone. But before that, I state the result informally, and illustrate the main logic through a three type example.

Only Market	Only Intermediary	Intermediary + Market
Only lowest type trades	Higher types *can* trade in equilibrium	Only lowest type trades

Table 2.1. BLC (severe adverse selection)

> **Theorem**: *If a subset of types is "concentrated" at the bottom, such that the highest type in this subset is "sufficiently" lower than all types **not** in the subset, then the presence of the market results in loss of efficiency.*

Before stating the formal result, I establish some notation.

Definition 5. ***Pooling Type***: *A type θ is said to be a Pooling Type if $(1+\alpha)\mathbb{E}[\theta'|\theta' \leqslant \theta] \geqslant \theta$. Let Θ_{Pool} denote the set of all such types.*

A type θ is *Pooling Type* is such that if the intermediary is operating in isolation, we can find a price p such that in equilibrium, if all types $\theta' \leqslant \theta$ are selling at p with probability one, then both both the buyer's interim IR, and the seller's IR are satisfied. This is because $(1+\alpha)\mathbb{E}[\theta'|\theta' \leqslant \theta] \geqslant \theta$, so we can choose $p \geqslant \theta$, and $\leqslant (1+\alpha)\mathbb{E}[\theta'|\theta' \leqslant \theta]$.

Let $\tilde{\theta} = \max\{\theta|\theta \in \Theta_{Pool}\}$. This denotes the *Highest Pooling Type*.

Theorem 5. *For every Θ, $\mu(.)$, and α, there exists an $\epsilon(\alpha, \mu(.)) > 0$, such that if $\tilde{\theta} > \theta_n$, $\tilde{\theta} < (1+\alpha)\theta_n$, and $\tilde{\theta} - \theta_n < \epsilon(\alpha, \mu(.))$, then, in the surplus maximising equilibrium with a market, i) there is no breakdown, and ii) the expected surplus from trade is strictly lower than the expected surplus in the surplus maximising equilibrium when there is no market.*

By definition of $\tilde{\theta}$, if $\tilde{\theta} > \theta_n$, the the BLC is *not* satisfied. The condition in Theorem 5 says that if $\tilde{\theta} > \theta_n$ *and* sufficiently close to θ_n, i.e. if types at the bottom are *sufficiently* concentrated, then the market leads to loss of efficiency. As in the breakdown case, the main idea behind the inefficiency here is that the market creates an endogenous outside option, and therefore an endogenous IR constraint for the lower types, and increases the minimum price that they would accept in equilibrium. In any equilibrium, $p_M \geqslant (1+\alpha)\theta_n$. Therefore, in any equilibrium, for any $\theta \leqslant \tilde{\theta}$, the *effective type* type is p_M. Since $\tilde{\theta} < (1+\alpha)\theta_n$, we have that for all $\theta \leqslant \tilde{\theta}$, their new , "equilibrium" reservation utility is strictly higher than their original reservation utility. The lower types are precisely the types responsible for the *lemons* problem, so their *effective type* increasing means the gap between the seller's IR and the buyer's interim IR widens, resulting in loss of efficiency.

2.4.4.2 Market improves efficiency

The final question one might have is that when the BLC fails, is *can the presence of the market can ever **improve** efficiency, compared to when the intermediary is operating in* *isolation?*[12] The answer to this question is yes as well! This may seem counter-intuitive at first; in this setting the screening device is the trade off between price and probability of trade, and it seems natural that the presence of a market where the seller can *always* sell should reduce the intermediary's ability to screen. Indeed, this is what we have seen till now.

The reason for this is that the intermediary has limited commitment power in the sense that it cannot offer a menu of *lotteries* over prices; each allocation in the intermediary's menu has exactly *one* price associated with it. In Section 2.4.5, I discuss why this is a limitation, and why randomisation over prices might help improve surplus in equilibrium, when the intermediary is operating in isolation. The idea behind the market improving surplus is exactly this: it helps with such randomisation, when the intermediary *itself* cannot randomise. I now illustrate why the market can help improve surplus through a three type example. I will focus on the core logic here, and skip many details, which are contained in the Appendix.

Example 1. *Consider a setting with three types $\{\theta_1, \theta_2, \theta_3\}$, where $\theta_1 > \theta_2 > \theta_3$, and the probability of type θ is denoted by $\mu(\theta)$. Suppose the prior $\mu(.)$ satisfies the following:*

- *The lemons condition:* $(1+\alpha)\mathbb{E}[\theta] < \theta_1$
- $(1+\alpha)\mathbb{E}[\theta|\theta \in \{\theta_2, \theta_3\}] > \theta_2$*, and* $\theta_2 = (1+\alpha)\theta_3$[13]
- $(1+\alpha)\mathbb{E}[\theta|\theta \in \{\theta_1, \theta_2\}] < \theta_1$

Then, there exists an equilibrium in the presence of the market that results in strictly higher surplus than any equilibrium when the intermediary is operating alone.

I first describe what happens if there is *just* an intermediary and what is the surplus maximising outcome in this case. Then, I argue how the presence of the market can help improve this. If there is just an intermediary,then the surplus maximising equilibrium is one in which types θ_2 and θ_3 trade with probability one, and θ_1 trades with probability π^*, where $\pi^* \in (0,1)$, and is such that:

$$\pi^*(\theta_1 - \theta_2) = (1+\alpha)\mathbb{E}[\theta|\theta \in \{\theta_2, \theta_3\}] - \theta_2$$

[12] When the BLC *is* satisfied, as I showed in Section 2.4.3, the market always results in a strict reduction in surplus.

[13] We don't need $\theta_2 = (1+\alpha)\theta_3$; the fact that the market can improve surplus works even without it. This is just for simplicity of exposition, and there is a more general version of this example in the Appendix.

This equilibrium outcome is achieved when the intermediary offers a menu $\mathcal{M}^* = \{\mathcal{L}^*, \mathcal{H}^*\}$, where $\mathcal{L}^* = (1, (1+\alpha)\mathbb{E}[\theta|\theta \in \{\theta_2, \theta_3\}])$, and $\mathcal{H}^* = (\pi^*, \theta_1)$.[14] When $\mathcal{M}^*$ is offered, given the condition on π^*, θ_2 is *indifferent* between $\mathcal{L}^*$ and $\mathcal{H}^*$. So, by Lemma 1, θ_1 and θ_3 *strictly* prefer $\mathcal{H}^*$ and $\mathcal{L}^*$ respectively. So, it is indeed optimal for θ_2 and θ_3 to choose $\mathcal{L}^*$, and for θ_1 to choose $\mathcal{H}^*$. Therefore, in equilibrium, θ_2 and θ_3 trade with probability one, and θ_1 trades with probability $\pi < 1$.

When there is just an intermediary, θ_1 cannot trade with a probability greater than π^* in equilibrium. The idea is that given $\mathcal{M}^*$, θ_2 is indifferent between the two options. So in order to increase the probability with which θ_1 trades, $\mathcal{H}^*$ has to be made more unattractive for θ_2, relative to $\mathcal{L}^*$. This can be done by either reducing the price in $\mathcal{H}^*$, or increase the price in $\mathcal{L}^*$, or both. Reducing the price in $\mathcal{H}^*$ would make it lower than θ_1, so this doesn't work.

Can we increase the price in $\mathcal{L}^*$ so that its greater than $(1+\alpha)\mathbb{E}[\theta|\theta \in \{\theta_2, \theta_3\}]$? In the equilibrium that I described above, *only* θ_1 is choosing $\mathcal{H}^*$, so at price θ_1, the buyer's value for the good is $(1+\alpha)\theta$. There is therefore, some "room" here: if θ_3 randomises between the two options and chooses $\mathcal{H}^*$ with positive probability, as long as this probability is small enough, the buyer's beliefs at price θ_1 will be such that the buyer is still willing to buy at this price. Such randomisation by θ_3 would also enable us to increase the price in option $\mathcal{L}^*$; since θ_3 is now choosing $\mathcal{L}^*$ with a probability strictly lower than one, the conditional expected value (for the buyer), at the price in $\mathcal{L}^*$ is strictly greater than $(1+\alpha)\mathbb{E}[\theta|\theta \in \{\theta_2, \theta_3\}]$.

But this doesn't work, because given Lemma 1, if θ_3 is indifferent between the two options, θ_2 strictly prefers $\mathcal{H}^*$. But then, if θ_1 and θ_2 are both choosing $\mathcal{H}^*$ with probability one (and θ_3 with a probability lower than one), the buyer's conditional expected value at price θ_1 falls below θ_1, because the prior satisfies $(1+\alpha)\mathbb{E}[\theta|\theta \in \{\theta_1, \theta_2\}] < \theta_1$. In short, there is no way to make *just* θ_3 randomise between the two options. The market solves this problem.

> *When the intermediary is operating alongside the market, there exists a menu, and an equilibrium with this menu, where θ_2 and θ_3 trade with probability one, and θ_1 trades with a probability $\pi' > \pi^*$.*

So how does the market help? The main idea is that the market endogenously alters the reservation utilities of types. So, if $p_M \geqslant \theta_2$, then both θ_2 and θ_3 have the same *effective*

[14] $\pi^* = \frac{(1+\alpha)\mathbb{E}[\theta|\theta \in \{\theta_2, \theta_3\}] - \theta_2}{\theta_1 - \theta_2}$ which is strictly lower than one, because by the *lemons condition*, $(1+\alpha)\mathbb{E}[\theta|\theta \in \{\theta_2, \theta_3\}] < \theta_1$.

type of p_M, and have the same ranking over options in the intermediary's menu. So, if θ_2 is indifferent between two options, θ_3 is too. Consider the menu $\mathcal{M}^{**} = \{\mathcal{L}^{**}, \mathcal{H}^{**}\}$ where $\mathcal{L}^{**} = (1, p')$, with $p' > (1+\alpha)\mathbb{E}[\theta|\theta \in \{\theta_2, \theta_3\}]$ and $\mathcal{H}^{**} = (\pi', \theta_1)$, where $\pi' > \pi^*$ and satisfies:

$$\pi'(\theta_1 - \theta_2) = p' - \theta_2$$

When the intermediary offers this menu, there is an equilibrium where $p_M = (1 + \alpha)\theta_3 = \theta_2$, and:

- θ_2 and θ_3 trade with probability one. θ_2 chooses $\mathcal{L}^{**}$ and trade with probability one at price p'. θ_3 *randomises* between choosing $\mathcal{L}^{**}$ and $\mathcal{H}^{**}$. If she chooses $\mathcal{L}^{**}$, then she sells at price p', and if she chooses $\mathcal{H}^{**}$, then with probability π', she sells at price θ_1, and with probability $(1 - \pi')$, she sells on the market at $p_M = \theta_2$.
- θ_1 chooses $\mathcal{H}^{**}$, and trades with probability π' at price θ_1.
- Buyer's strategy is to buy at both prices p' and θ_1

If $p_M = \theta_2$, both θ_2 and θ_3 have the same *effective type* of θ_2. Given the condition on π', and the buyer's strategy, both are indifferent between $\mathcal{L}^{**}$ and $\mathcal{H}^{**}$, and therefore the strategy of both θ_2 and θ_3 is sequentially rational. Since only θ_3 is selling on the market in equilibrium (if she chooses $\mathcal{H}^{*}$, and cannot sell through the intermediary), $p_M = (1 + \alpha)\theta_3 = \theta_2$.

The only thing left to argue is that given the seller's strategy, the buyer's strategy is optimal. As with menu $\mathcal{M}^{*}$, only θ_2 and θ_3 are choosing option $\mathcal{L}^{**}$, so why is the buyer willing to buy at a price $p' > (1+\alpha)\mathbb{E}[\theta|\theta \in \{\theta_2, \theta_3\}]$? The reason is that θ_3 is choosing this option with a probability strictly lower than one, so the expected value of the good, conditional on price p', is indeed greater than $(1+\alpha)\mathbb{E}[\theta|\theta \in \{\theta_2, \theta_3\}]$. And the randomisation by θ_3 is such that the buyer's expected value conditional on price θ_1 is still (weakly) greater than θ_1.

To sum up, when the BLC is not satisfied, two things can happen. Either the presence of the market reduces surplus, compared to when the intermediary operates alone; Theorem 5 provides a sufficient condition for this to happen. The second possibility is that the presence of the market improves surplus, compared to when the intermediary is operating alone. Section B.0.7 of the Appendix contains an example of this.[15]

[15] It can never be the case that the intermediary *strictly* reduces surplus, compared to when there is just a market. This is because the intermediary can always offer a menu with just allocation $(0, 0)$, so in equilibrium, it is as if there is no intermediary.

2.4.5 What if the Intermediary Could Offer a Lottery Over Prices?

I specified the mechanism offered by the intermediary as a menu, where each allocation in the menu consists of a *single* price. In this section, I first point out that this specification of the intermediary's mechanism is *with* loss, and then talk about how much of my analysis still holds, and which results survive, if I consider more general mechanisms.

Recall that in my model, the buyer's IR must be satisfied at *every* price at which trade happens through the intermediary. As Gerardi, Hörner, and Maestri [2014] show, when the buyer's IR must be satisfied at every price, it is *with* loss to consider allocations with one price. So, there are outcomes attainable when the intermediary offers a menu of *lotteries over prices*, that are not attainable when each allocation in the menu can contain only one price. Formally, Gerardi, Hörner, and Maestri [2014] show that in this setting, it is without loss to focus on the following direct mechanisms: the intermediary maps each report θ to f_θ, a probability distribution over $\{0,1\} \times \mathbb{R}_+$. Here, $f_\theta(p)$ is the probability of trade happening at price p, if θ is reported, and $f_\theta(0)$ is used to denote the probability of no trade. In the Appendix, I describe the class of mechanisms in detail.

But why does offering lotteries over prices expand the set of attainable outcomes? The idea is that in equilibrium, prices contain information about the seller's type. When each allocation only contains a single price, the only information contained in this price is what types of the seller chose the allocation with this price in equilibrium. When the intermediary can map reports to lotteries over prices, it allows the intermediary greater flexibility in *how* to communicate the information elicited from the seller, to the buyer.

The above discussion might lead us to believe that it is this limited commitment on part of the intermediary that allows the market to disrupt its operation. However, with two types, if the prior satisfies the *lemons condition*, the presence of the market would *still* result in a breakdown:

Proposition 4. *Suppose there are two possible types, the lemons condition is satisfied, and the intermediary can offer a mechanism that maps reports to lotteries over prices. Then, the unique equilibrium outcome still involves only the lower type trading.*

The proof of the above proposition is in the Appendix. With more than two types, the analysis with lotteries over prices becomes quite complex, and therefore for tractability, I restrict attention to the case where any option in the menu has a single price. I should point out however, that with more that two types, it is still possible to argue that if types are *sufficiently* far apart, the presence of the market leads to a breakdown. However, getting a closed-form condition analogous to the BLC is difficult.

2.5 Conclusion

In this paper, I highlight the extent to which the presence of outside trading opportunities can disrupt intermediated trade. A seller who decides to trade through an intermediary, usually *also* has the option to sell her good without the intermediary. Selling without the intermediary can take several forms. For example, the seller can negotiate with a potential buyer directly. Or the intermediary could represent the legitimate channel of sale; if seller is unable to sell through this channel, she can sell on a "black market".

I model this outside selling opportunity as a static competitive market, where trade takes place at a single price. My main result is that under some conditions, the presence of the market completely destroys any efficiency gains from intermediated trade: in the unique equilibrium outcome, *only* the lowest type trades. The market "infects" the intermediary; in equilibrium, it is as if there is *just* a static, competitive market plagued with severe adverse selection, and no intermediary. I also provide conditions under which the the market results in inefficiency, but does not cause a total breakdown of trade.

In this paper, I go with a particular interpretation of how the intermediary operates. I model this as a bilateral trade setting, where the intermediary is brokering trade between a seller and a particular buyer, and this buyer *will* buy as long as he anticipates no loss from buying. The intermediary in my model can *commit* to randomising between offering and not offering the seller the seller the opportunity to trade. While this interpretation might make sense for some settings, in others, it might be unrealistic that the intermediary can commit to randomise.

However, even if the the intermediary cannot commit to randomise, a common feature of many settings is that there is some uncertainty associated with trade at higher prices. This uncertainty might arise because the intermediary has to "search" for potential buyers, and at a high price, it might not be able to find a buyer. This could be the case when potential buyers have identical payoff functions with respect to the good, but are heterogeneous in terms of wealth. For example, it could be that all buyers *want* to but at a high price, if *only* high quality goods are selling at higher prices in equilibrium, but most buyers are cash constrained. Therefore, at a high price, with some probability, the intermediary may not be able to find an appropriate buyer.

But even with this alternative interpretation, the key idea behind separation of high and low types remains the same: at higher prices, the probability of trade is lower. This suggests that my analysis of the disruption caused by the market goes though with alternative interpretations as well. I choose not to model the "search" for the buyer, or any other features of the setting that may give rise to the probability of trade-price trade off. I assume that the intermediary can commit to any menu. I show that *even* with this commitment power that allows the intermediary to create arbitrary trade offs between price and probability of trade, breakdown cannot be avoided under certain conditions. I leave the more general analysis of intermediation with outside selling opportunities for future work.

Appendix A
Appendix for Chapter 1

A.1 Proofs

A.1.1 Proof of Proposition 1

Proof. Given that $q_{Gg} = 1$ and $q_{Mg} + q_{Mb} = 1$, our maximisation problem reduces to:

$$\max_{q_{Mg}, q_{Bb} \in [0,1]} \quad \mu_G \pi_g + \mu_B q_{Bb} \pi_b + \mu_M q_{Mg} \pi_g + \mu_M (1 - q_{Mg}) \pi_b$$

$$\text{subject to} \quad q_{Mg} \alpha_g + (1 - q_{Mg}) \alpha_b \geqslant q_{bB} \alpha_b$$

In an optimal mechanism, we must also have that $q_{Mg}\alpha_g + (1 - q_{Mg})\alpha_b = q_{bB}\alpha_b$. To see this, suppose the inequality is strict. Let ${q_{Mg}}' = q_{Mg} + \varepsilon$, and ${q_{Mb}}' = 1 - {q_{Mg}}' = q_{Mb} - \varepsilon$, where $\varepsilon > 0$. For ε small enough, IC_{MB} is still satisfied and the principal's expected payoff increases by $\varepsilon \mu_M (\pi_g - \pi_b)$. We can therefore substitute $q_{Mg}\alpha_g + (1 - q_{Mg})\alpha_b = q_{bB}\alpha_b$ into our objective function, and we get

$$\mu_G \pi_g + \mu_B \pi_b (1 - q_{Mg}(1 - \frac{\alpha_g}{\alpha_b})) + \mu_M q_{Mg} \pi_g + \mu_M (1 - q_{Mg}) \pi_b$$

$$= \mu_G \pi_g + \mu_M \pi_b + q_{Mg} \{ \mu_M (\pi_g - \pi_b) - \mu_B \pi_b (1 - \frac{\alpha_g}{\alpha_b}) \}$$

The only choice variable is q_{Mg} now, and whether the above expression is increasing or decreasing in q_{Mg} depends on the sign of it's coefficient, $\{\mu_M(\pi_g - \pi_b) - \mu_B \pi_b (1 - \frac{\alpha_g}{\alpha_b})\}$. If the coefficient is strictly positive, then the optimal mechanism has $q_{Mg} = 1$. Also, because IC_{MB} holds with equality, we have $q_{Bb} = \frac{\alpha_g}{\alpha_b}$ in the optimal mechanism. A bit of rearranging gives us that $\{\mu_M(\pi_g - \pi_b) - \mu_B \pi_b (1 - \frac{\alpha_g}{\alpha_b})\} > 0$ is equivalent to $\lambda > \lambda^*$.

Similarly, if $\lambda < \lambda^*$, the optimal mechanism has $q_{Mg} = 0$, and therefore $q_{Bb} = 1$. If $\lambda = \lambda^*$, principal's expected payoff is constant in q_{Mg} and therefore any $q_M g \in [0, 1]$ is optimal, with q_{Bb} again being determined by the equality of IC_{MB}.

□

A.1.2 Proof of Theorem 1

We first establish some notation. Recall that for any time period t, the set of all possible period t histories is denoted by $\mathcal{H}_t$, where $\mathcal{H}_t = (\mathcal{N} \cup \varnothing)^t$. The representative period t history is denoted by $h_t \in \mathcal{H}_t$. The action (or proposal) space of an agent of type S at any history is given by $A_S(h_t) = A_S = \{S \cup \varnothing\}$, where $\varnothing$ represents remaining silent. An element of $A_S(h_t)$ is given by a_S^t. Given a history h_t, and a proposal in $(\mathcal{N} \cup \varnothing)$ by the agent, the principal can either *accept*, or *reject* this proposal.

A behaviour strategy maps histories and types into action spaces. For the agent of type S, $\sigma_S(a_S^t|h_t)$ denotes the probability of choosing a_S^t at history h_t. For the principal, $\sigma_P(h_t, i)$ denotes the probability of accepting proposal $i \in \mathcal{N}$ at history h_t. If at h_t, agent is silent, then $\sigma_P(h_t, \varnothing) = 0$, as there is no project to accept. At any history h_t, we denote the probability the principal attaches to type S by $\mu_S(h_t)$. If, following h_t, i is proposed, the updated beliefs are given by $\mu_S(h_t, i)$ for each S, and by $\mu_S(h_t, \varnothing)$ if the agent is silent at h_t. For any t and any $t' < t$, let $h_t(t') \in (\mathcal{N} \cup \varnothing)$ be the proposal at period t', along this history h_t. We denote by $h_{t-1}(h_t)$ the period $t-1$ history obtained by removing proposal $h_t(t-1)$ from h_t, so that $h_t = (h_{t-1}(h_t), h_t(t-1))$.

Our solution concept is Perfect Bayesian Equilibrium, as defined in Fudenberg and Tirole (1991). We want to highlight that here, something stronger than Bayes' Rule is used to update beliefs following any proposal at any history. To understand this, fix any history h_t. Even if this is a history that arises with probability zero along the equilibrium path, beliefs following a proposal i at this history are updated using Bayes' Rule if $\exists$ a type S such that $\mu_S(h_t) > 0$ and $\sigma_S(i|h_t) > 0$. Beliefs are allowed to be *completely arbitrary* only if given h_t, the proposal made by the agent had probability zero, according to the agent's strategy. However, even when following a proposal i at history h_t, beliefs are allowed to be completely arbitrary, they must still have support in $\{S \subseteq \mathcal{N} | i \in S\}$, i.e. the set of types that have i, because only these types can possibly propose i. This is in contrast to models without hard evidence.

A.1.2.1 Pooling equilibrium:

Strategies: The principal's strategy σ_P is: for any history h_t, and proposal i, $\sigma_P(h_t, i) = 1$, i.e. at each history, the principal accepts any proposal with probability one. If the agent is of the empty type, then for any h_t, $\sigma_E(\varnothing|h_t) = 1$. For an agent of type $S \neq E$, let $i^* = \min\{i|i \in S\}$, i.e. the most preferred project of the agent in S. At any history h_t, $\sigma_S(i^*|h_t) = 1$, so at any history, the agent always proposes her favorite available project with probability one.
Beliefs: Fix any h_t (if $t = 0$, this is the null history). (i) If the agent is silent at h_t, the beliefs are $\mu_E(h_t, \varnothing) = 1$ and $\mu_S(h_t, \varnothing) = 0$ for all $S \neq E$. (ii) If the agent proposes g at h_t, the beliefs are $\mu_G(h_t, g) = 1$ and $\mu_S(h_t, g) = 0$ for all $S \neq E$. (iii) If the agent proposes b at h_t, the beliefs are $\mu_B(h_t, b) = \frac{\mu_B}{\mu_B + \mu_M}$, and $\mu_M(h_t, b) = \frac{\mu_M}{\mu_B + \mu_M}$.

We now argue that the beliefs satisfy the requirements for PBE. At any h_t, if the agent takes a probability zero action, PBE imposes no restrictions on beliefs. So, we only need to worry about the case where the agent's proposal at h_t has positive probability given this history. In this case, PBE requires the beliefs to be determined by Bayes' Rule. At h_t, silence is a positive probability action only if $\mu_E(h_t) > 0$, since only type E has $\sigma_E(\varnothing|h_t) > 0$. In this case, the belief $\mu_E(h_t, \varnothing) = 1$ is precisely the one determined by Bayes' Rule, since $\mu_E(h_t) > 0$, $\sigma_E(\varnothing|h_t) = 1 > 0$, and $\sigma_S(\varnothing|h_t) = 0$ for all $S \neq E$. Similarly, g is a positive probability action only if $\mu_G(h_t) > 0$ and b is a positive probability action only if $\mu_B(h_t) > 0$. By identical reasoning as before, we can argue that in both these cases, the specified beliefs are the ones determined by Bayes' Rule.

Lemma 5. *The strategies and beliefs described above constitute a Perfect Bayesian Equilibrium.*

Proof. We have already argued that the beliefs that we described satisfy the requirements for being part of a PBE. Now, we only have to argue that in the continuation game starting at any history h_t, the strategies of the principal and agent constitute a Bayes Nash Equilibrium (BNE), given the principal's beliefs $\mu(h_t)$ at that history.

Fix any history h_t. If the agent is silent, there is no action for the principal to take. If the agent proposes g, then irrespective of $\mu(h_t)$, it is sequentially rational to accept g, since this is the highest payoff the principal can get. If the agent proposes b, again, the exact beliefs of the principal do not matter. Whatever the beliefs are, they have support in $\{S \subseteq \mathcal{N}|b \in S\} = \{B, M\}$. So, the principal must believe with probability one that it is a type that *has* b, and that therefore will propose b in the next period, if the principal rejects this proposal (in fact, in every future period). So, it is sequentially rational for the principal to accept this proposal too.

For the agent, if it is of type E, it can only stay silent. If it is of type G, it again has no profitable deviation to proposing g, as the principal will accept it if proposed. If it is of type B or M, again, the principal will accept b if proposed, no there is no profitable deviation to proposing b. □

Lemma 6. *The principal's payoff from the pooling equilibrium is the same as his payoff from the pooling mechanism.*

Proof. Along the equilibrium path, from type G, g is proposed and accepted at $t = 0$, and from types B and M, b is proposed and accepted at $t = 0$. This exactly replicates the implementation probabilities of $q_{Gg} = 1$, and $Q_{Bb} = q_{Mb} = 1$ from the pooling mechanism. □

A.1.2.2 Separating equilibrium:

Before describing the strategies and beliefs, we define two classes of histories.

Definition 6. *A **type 1** history h_t is one where there is **no** $t^{'} < t^*(\delta)$ such that $h_t(t^{'}) = b$. In other words, there is no $t^{'} < t^*(\delta)$ such that along h_t, b was proposed at $t^{'}$. The null history is a **type 1** history.*

Definition 7. *A **type 2** history h_t is one where $t \neq 0$ and there **is** a $t^{'} < t^*(\delta)$ such that $h_t(t^{'}) = b$. In other words, there is a $t^{'} < t^*(\delta)$ such that along h_t, b was proposed at $t^{'}$.*

Strategies:

- The principal's strategy σ_P is: At any history h_t, $\sigma_P(h_t, g) = 1$. So, g is accepted if proposed at any history. If h_t is such that $t < t^*(\delta)$, $\sigma_P(h_t, b) = 0$. If $t \geqslant t^*(\delta)$, and the history is of type 1, then $\sigma_P(h_t, b) = 1$. If $t \geqslant t^*(\delta)$, and the history is of type 2, then $\sigma_P(h_t, b) = 0$.

- If the agent's type is E, for any history h_t, $\sigma_E(\varnothing|h_t) = 1$.

- If the agent's type is G, for any history h_t, $\sigma_G(g|h_t) = 1$.

- If the agent's type is B, and h_t is such that $t < t^*(\delta)$, then $\sigma_B(\varnothing|h_t) = 1$. If $t \geqslant t^*(\delta)$ and h_t is of type 1, then $\sigma_B(b|h_t) = 1$. If $t \geqslant t^*(\delta)$ and h_t is of type 2, then $\sigma_B(\varnothing|h_t) = 1$.

- If the agent's type is M, then at h_0, $\sigma(g|h_0) = 1$. Now consider h_t with $t > 0$. If h_t is of type 2, then $\sigma_M(g|h_t) = 1$. If h_t is of type 1, and $t < t^*(\delta)$, then $\sigma_M(\varnothing|h_t) = 1$. If h_t is of type 1, and $t \geqslant t^*(\delta)$, then $\sigma_M(b|h_t) = 1$.

Beliefs:

At the null history h_0:

- If the agent is silent, the beliefs are $\mu_E(h_0, \varnothing) = \frac{\mu_E}{\mu_E+\mu_B}$ and $\mu_B(h_0, \varnothing) = \frac{\mu_B}{\mu_E+\mu_B}$. If the agent proposes g, the beliefs are $\mu_G(h_0, g) = \frac{\mu_G}{\mu_G+\mu_M}$, $\mu_M(h_0, g) = \frac{\mu_M}{\mu_G+\mu_M}$. If the agent proposes b, the beliefs are $\mu_M(h_0, b) = 1$.

- These beliefs all satisfy the conditions for PBE. The agent staying silent and proposing g both occur with positive probability at h_0, as $\sigma_E(\varnothing|h_0) = \sigma_B(\varnothing|h_0) = 1$ and $\sigma_G(g|h_0) = \sigma_M(g|h_0) = 1$, and these two cases, beliefs are the ones determined by Bayes' Rule. Proposing b is a probability zero action here, thus, beliefs can be arbitrary.

At a type 1 history h_t, where $t > 0$:

- If the agent is silent, beliefs are $\mu_B(h_t, \varnothing) = 1$. If the agent proposes g, beliefs are $\mu_G(h_t, g) = 1$. If the agent proposes b, and $t < t^*(\delta)$ beliefs are $\mu_M(h_t, b) = 1$. If $t \geqslant t^*(\delta)$, beliefs are $\mu_B(h_t, b) = 1$.

- The beliefs are consistent with PBE. Silence is a positive probability action at h_t only if $t < t^*(\delta)$ and $\mu_B(h_t) > 0$ or $\mu_M(h_t) > 0$. This is the case only if $h_t(t-1) = \varnothing$, and in this case beliefs are the same as the ones determined by Bayes' Rule. Proposing g is never a positive probability action at h_t so beliefs can be arbitrary in a PBE and we're done.

- Proposing b is a positive probability action only if $t \geqslant t^*(\delta)$. (1)If $t = t^*(\delta)$, this is the case only if $h_t(t-1) = \varnothing$. In this case, $\mu_B(h_t) = 1$, and $\sigma_B(b|h_t) = 1$, so beliefs are precisely the ones determined by Bayes' Rule. (2) If $t > t^*(\delta)$, then proposing b is a positive probability event only if $h_t(t-1) = \varnothing$ or $h_t(t-1) = b$. In this case, $\mu_B(h_t) = 1$ and $\sigma_B(b|h_t) = 1$, so again, the beliefs we specified are the ones determined by Bayes' Rule.

At a type 2 history:

- If the agent is silent, the beliefs are $\mu_E(h_t, \varnothing) = \frac{\mu_E}{\mu_E+\mu_B}$ and $\mu_B(h_t, \varnothing) = \frac{\mu_B}{\mu_E+\mu_B}$. If the agent proposes g, beliefs are $\mu_M(h_t, g) = 1$. If the agent proposes b, beliefs are $\mu_M(h_t, b) = 1$.

- The beliefs are consistent with PBE. If the agent is silent, there are two possibilities. Either, $h_t(t-1) = \varnothing$, in which case $\mu_E(h_t) = \frac{\mu_E}{\mu_E+\mu_B}$ and $\mu_E(h_t) = \frac{\mu_E}{\mu_E+\mu_B}$, so the beliefs that we mentioned are precisely the ones determined by Bayes' Rule. If $h_t(t-1) = b$, or $h_t(t-1) = g$, silence is a probability zero action at h_t, and PBE allows beliefs to be arbitrary.

- If the agent proposed b at h_t, irrespective of what was proposed at $t-1$, b is a probability zero action. This is because $\sigma_M(b|h_t) = \sigma_B(b|h_t) = 0$ at h_t. So again, PBE allows beliefs to be arbitrary. If agent proposed g at h_t, g is a positive probability action here only if $h_t(t-1) = b$, or $h_t(t-1) = g$. In this case since $\mu_M(h_t) = 1$ and $\sigma_M(g|h_t) = 1$, beliefs are precisely the ones determined by Bayes' Rule. If $h_t(t-1) = \varnothing$, beliefs can be arbitrary.

Lemma 7. *The strategies and beliefs described above constitute a Perfect Bayesian Equilibrium.*

Proof. We have already argued that the beliefs that we described satisfy the requirements for being part of a PBE. We must now argue that in the continuation game starting at any history h_t, the strategies of the principal and agent constitute a Bayes Nash Equilibrium (BNE), given the principal's beliefs $\mu(h_t)$ at that history.

For the principal, fix any history h_t. If the agent is silent, there is no action for the principal to take. If the agent proposes g, then irrespective of $\mu(h_t)$, it is sequentially rational to accept g, since this is the highest payoff the principal can get. If the agent proposes b, we need to consider three cases. (i) $t < t^*(\delta)$: The principal's strategy is to reject b at such a history. His beliefs following a proposal of b at h_t are $\mu_M(h_t, b) = 1$. If the principal accepts, he gets α_b, and if he rejects, he expects the agent to propose g in the next period. So, rejection is indeed sequentially rational if the principal is sufficiently patient. (ii) $t \geqslant t^*(\delta)$ and the history is of type 1. In this case, if b is proposed, the principal's beliefs are $\mu_B(h_t, b) = 1$ and if he rejects b, he expects b to be proposed again in the next period. So, accepting b is sequentially rational. (iii) $t \geqslant t^*(\delta)$ and the history is of type 2. In this case, the principal's strategy is to reject b. His beliefs following a proposal of b are $\mu_M(h_t, b) = 1$, so the same reasoning as case (i) follows, and rejection is sequentially rational.

For the agent, fix a history h_t. If her type is E, she can only be silent at any history. If her type is G, her strategy is to propose g, which is optimal, since the principal would accept it. If her type is B, and (i) $t < t^*(\delta)$, her strategy is to stay silent. Consider a one-shot deviation where she deviates by proposing b. This proposal would be rejected and since her strategy involves staying silent at any future period following this rejection, therefore getting a payoff of zero, she cannot be better off by this deviation. (ii) If $t \geqslant t^*(\delta)$ and the history is of type 1, proposing b is optimal since it would be accepted and α_b is the highest payoff the agent can get. (iii)If $t \geqslant t^*(\delta)$ and the history is of type 2, silence is optimal. Consider a one-shot deviation where the agent proposes b instead. It would be rejected, and the agent's strategy is to stay silent in each period that follows. So, this is not a profitable deviation.

If the agent's type is M, at h_0, her strategy is to propose g. Consider the one-shot deviation where she proposes b instead. It will be rejected and she will propose g at $t = 1$, which will get accepted. Clearly, this is not profitable, as she can propose g at $t = 0$ and it will get accepted. If the one-shot deviation involves silence at $t = 0$, her strategy then is to stay silent till $t^*(\delta)$, at which point she proposes b can it is accepted. So her payoff is $\delta^{t^*(\delta)}\alpha_b$ which is $\leqslant \alpha_g$. So, this deviation is not profitable either. We can similarly rule out deviations at other histories.

□

A.1.3 Induced Social Choice Function

We provide the details of collapsing the probability of various outcomes involving i into a single probability here.

Fix a protocol, a strategy the principal has committed to, and a best response of the agent. Let (i, t) denote the outcome that project i is implemented (proposed and accepted) at t. No project ever being implemented is also a possible outcome. The proof proceeds in two steps. We first show that the strategy and the best response induce, for any type S, a probability distribution over outcomes. We then condense these probabilities to arrive at the *induced* SCF.

Fix a type S and a project $i \in S$. Recall that at any history h_t, the proposer is $P(h_t)$. Let $x(i|h_t)$ be the probability with which, according to $P(h_t)$'s strategy, the proposer $P(h_t)$ offers $O(h_t) \in \mathcal{O}(h_t)$ such that $i \in O(h_t)$ at h_t. Let $y(h_t, i)$ be the probability that the other party, who is not the proposer, *accepts* i in the offer $O(h_t)$. We define the probability of any history inductively. We denote period 0 history by h_0, so at $t = 1$, for any $h_1 = (h_0, i)$, where $i \in S \cup \{\varnothing\}$. We define $\nu(h_1) = x(i|h_0)(1 - y(h_0, i))$, which is just the probability that i was proposed at $t = 0$ but not accepted, and thus the probability of history h_1 at $t = 1$. This is clearly a number in $[0, 1]$. Given that we have defined $\nu(h_t)\ \forall\ h_{t'},\ t' \leqslant t$, and any $h_{t+1} = (h_t, i)$ for some $i \in S \cup \{\varnothing\}$, we have that $\nu(h_{t+1}) = \nu(h_t)x(i|h_t)(1 - y(h_t, i))$.

We define the probability of outcome (i, t) as

$$p_S(i,t) := \sum_{h_t \in \mathcal{H}_t} \nu(h_t) x(i|h_t) y(h_t, i)$$

for all $i \in S$. It can be verified that the sum of probabilities for all outcomes in which a project *is* implemented,

$$\sum_{t=0}^{\infty} \sum_{i \in S} p_S(i,t) \leqslant 1,$$

where the probability of the outcome that *no* project is ever implemented is

$$1 - \sum_{t=0}^{\infty} \sum_{i \in S} p_S(i,t).$$

Thus, for any type S, the strategy and best response induce a probability distribution over outcomes.

We now construct the corresponding Induced SCF. For any S', probability of implementation of $i \in S'$ is

$$f^I_{S'}(i) = \sum_{t=0}^{\infty} \delta^t p_{S'}(i,t).$$

A.1.4 Proof of Theorem 2

Proof. Fix a static IC mechanism. In this mechanism, any report $S = \{i_1, i_2, \ldots, i_m\}$, is mapped to implementation probability q_{Sk} for project i_k, and $\alpha_{i_1} < \alpha_{i_2} \ldots < \alpha_{i_m}$. We first define, for any S,

$$y_S(\delta) := \min\{\frac{1}{q_{S1} + \frac{q_{S2}}{\delta} + \ldots \frac{q_{Sm}}{\delta^{m-1}}}, 1\}$$

and let $y(\delta) = \min y_S(\delta)$.

We now construct the corresponding strategy in the sequential delegation game. According to this strategy:

- At any $t \in \{0, 1, \ldots, N-1\}$, the principal accepts *no* proposal irrespective of history.
- At every history where at any $t \in \{0, 1, \ldots, N-1\}$ the agent proposed anything other than project $t+1$ or $\varnothing$, the principal rejects any proposal.
- Fix a history where the set of projects proposed from $t = 0$ until $t = N-1$ is $S' = \{i_1, i_2, \ldots, i'_m\}$, and each project i was proposed at $t = i-1$. We call this

history $h_N^{S'}$. Let $q_{S'1},\ q_{S'2},\ \dots\ q_{S'm'}$ be the probabilities of implementation of each project in S', when S' is reported in the mechanism we have fixed.

- At h_N^S, at $t = N$, if i_1 is proposed, the principal accepts with probability $y(\delta)q_{S1}$.
- If agent does *not* propose i_1 at h_N^S, the principal rejects any proposal at any $t > N$.
- At the history $(h_N^S, i_1, i_2 \dots, i_{k-1})$ if the agent proposes i_k, it is accepted with probability $\frac{y(\delta)q_{Sk}}{\delta^{k-1}(q_{S1}+\frac{q_{S2}}{\delta}+\dots+\frac{q_{S(k-1)}}{\delta^{k-2}})}$. If agent does *not* propose i_k at $(h_N^S, i_1, i_2, \dots, i_{k-1})$, the principal rejects the current proposal and any proposal at any future period t.
- Period $N + (m' - 1)$ onward (given that history until period N is $h_N^{S'}$), no project is accepted, irrespective of history.

Note that since the set of projects proposed in the first N periods is S', it is optimal for the agent to report projects in decreasing order of the principal's preference in the next m' periods. If the agent stays silent at any of the m' periods that follow, or recommends a project *out of turn*, the principal never accepts any project again. Let the expected payoff from reporting S' in the mechanism be $E_{S'}$. It is easy to see that if the agent proposes all projects she has in the first N periods, she gets an expected payoff of $\delta^N y(\delta)E_{S'}$. Thus, since the mechanism was IC, it is indeed a best response for the agent to propose *all* projects she has in the first N periods. The principal's payoff from this strategy and best response is therefore the product of $\delta^N y(\delta)$ and the expected payoff from the mechanism. It can be easily verified that $y(\delta) \to 1$ as $\delta \to 1$. So, the principal's payoff from this strategy approaches the payoff from the mechanism as $\delta \to 1$.

□

We now provide an example that shows that with certain delegation protocols, the ability to commit to a strategy may not be enough to attain the commitment payoff from the optimal static mechanism. Recall the delegation protocol and example in Section 1.4.2 where there are three possible projects, $\mathcal{N} = \{1, 2, 3\}$, and three equally likely types in the support of μ with $\mathcal{S} = \{\{1, 2\}, \{2\}, \{2, 3\}\}$.

Proof. Suppose there is a commitment strategy of the principal in this alternative game, and a best response of the agent that attains the payoff from the optimal mechanism. This commitment strategy and best response give rise to an *induced* SCF. Recall from the construction of the induced SCF, that for any type S, and any $i \in S$, we have $f_{S'}^I(i) = \sum_{t=0}^{\infty} \delta^t p_{S'}(i,t)$. We therefore must have that $f_{2,3}^I(2) \to 1$ and $f_{1,2}^I(2) \to \frac{3}{8}$ as

$\delta \to 1$. Recall that a history here is simply a sequence of restriction sets that have been rejected by the agent. Now, for type $\{2,3\}$, consider any t, and history h_t at which 2 is implemented with positive probability. At this history, it must be that $2 \in O(h_t)$ and this is accepted by $\{2,3\}$ with positive probability. But, even $\{1,2\}$ can accept 2 at this history. Moreover, since the history involves a sequence of rejections, and the ability to reject is not type-dependent, *any* such history can also be reached when the type is $\{2,3\}$. This contradicts the fact that $f^I_{1,2}(2) \to \frac{3}{8}$. □

A.1.5 Proof of Theorem 3

Proof. The proof proceeds in three steps. Broadly, we first show that in solving for the optimal mechanism, the optimization problem can be divided into two parts. Then we show that if the solution to the second part corresponds to a value in a certain set, the payoff from the optimal mechanism can be replicated in equilibrium. Lastly, we show that the solution to the second part must always correspond to a value in this set.

Let the set of types be $\{S_1, S_2 \dots S_M\}$ where for any $i < i'$, we have that $S_{i'} \subset S_i$. For any type S_i, let $\mu(S_i) = \mu_i$ and let $q_{i,j}$ be the probability of implementation of $j \in S_i$ in the mechanism when the report is S_i. Let the expected value to the agent, corresponding to any report S_i, be denoted by E_i, where $E_i := \sum_{j \in S_i} q_{i,j}\alpha_j$. Observe that due to type-dependent message spaces and the support of $\mu(.)$, the IC constraints here boil down to $(M-1)$ inequalities:

$$E_1 \geqslant E_2, \dots \geqslant E_M,$$

and we refer to the inequalities $E_{i+1} \dots \geqslant E_M$ as the IC constraints *below* i and the inequalities $E_1 \geqslant E_2 \dots \geqslant E_i$ as the IC constraints *above* i.

Lemma 8. *In any optimal mechanism, any report must generate the same expected payoff for the agent. Formally, for any two reports S_i and $S_{i'}$, it must be that $E_i = E_{i'}$.*

Proof. We prove this by contradiction. Suppose, in an optimal mechanism there are i, i' such that $E_i \neq E_{i'}$. Without loss, let $i < i'$. This implies, given the nature of the support of $\mu(.)$, that $S_{i'} \subset S_i$. Since the optimal mechanism is IC, it must be that $E_i \geqslant E_{i'}$, and since $E_i \neq E_{i'}$, we have that $E_i > E_{i'}$. This in turn implies that we can find *consecutive types* $k, k+1$ such that $i \leqslant k < k+1 \leqslant i'$ and $E_k > E_{k+1}$. So without loss, let $i' = i+1$. Let i^* be the lowest indexed project in S_i.

- **Case 1**: *Corresponding to report S_i, $q_{i,j} > 0$, for some $j > i^*$.*

In this case, consider the following perturbation: Let $q'_{i,j} = q_{i,j} - \varepsilon$ and $q'_{i,i*} = q_{i,i*} + \varepsilon$. Now, $E'_i = E_i - \varepsilon(\alpha_j - \alpha_i) < E_i$. The ε in the perturbation is small enough that $E'_i > E_{i+1}$. Other than this change, all allocation probabilities corresponding to all other reports are unchanged, relative to the original mechanism. This new mechanism is IC, because since the original mechanism was IC, and we have reduced E_i, all IC constrains *above* i still hold. The inequality between E'_i and E_{i+1} is preserved, so this IC still holds. All IC constraints *below* i still hold, clearly. Thus we have constructed another IC mechanism in $(*)$ that gives a strictly higher expected payoff to the principal, as his payoff from type S_i increases by $\varepsilon(\pi_{i*} - \pi_j)$. Thus, the mechanism we started out with cannot be optimal.

- **Case 2**: *Corresponding to report S_i, $q_{i,j} = 0$, for every $j > i^*$.*

  We can again construct an IC mechanism in $(*)$ that gives strictly higher expected payoff to the principal. Since $q_{i,j} = 0$, for every $j > i^*$, we have that $E_i = q_{i,i*}\alpha_{i*} \leqslant \alpha_{i*}$, as only i^* might have positive allocation probability in S_i. Also, $E_i > E_{i+1}$, so it must be that $\sum_j q_{i+1,j} < 1$. This is because all projects in S_{i+1} have weakly higher payoff for the agent than α_{i*}, so if their allocation probabilities sum up to one, we would have $E_{i+1} \geqslant \alpha_{i*} \geqslant E_i$, which cannot be. So, since $\sum_j q_{i+1,j} < 1$, in particular, $q_{i+1,(i+1)*} < 1$. Consider the following perturbation: let $q'_{i+1,(i+1)*} = q_{i+1,(i+1)*} + \varepsilon$ where ε is small enough that $E'_{i+1} = E_{i+1} + \varepsilon\alpha_{(i+1)*} < E_i$, so the IC constraint between $i, i+1$ is preserved. Everything else is unchanged with respect to the original mechanism. Clearly, all IC constraints *above* i hold, and all *below* i hold as well as we increased E_{i+1}. This mechanism gives the principal a higher expected payoff since the payoff from type S_{i+1} has increased. So, the mechanism we started out with cannot be optimal.

This completes the proof of our claim that in any optimal mechanism, *any* report must generate the same expected payoff for the agent.

□

Now that we have shown this, the principal's optimization problem (finding the payoff-maximizing mechanism among all IC mechanisms in $(*)$) can be divided into two parts. First, for any expected value v, find the optimal mechanism corresponding to *this* v; the mechanism that maximizes the principal's payoff when each report generates an expected payoff of v for the agent. Then, maximize the principal's payoff over the possible values of v, i.e. find the values of v the optimal mechanism corresponding to which generates the highest expected payoff for the principal. Our aim is not to solve for the optimal mechanism, but rather show it is always the case that the principal's expected payoff from the optimal mechanism can be attained in equilibrium. We do this in the steps that follow.

Lemma 9. *In any optimal mechanism, it cannot be that $v > \alpha_{M*}$, where M^* is the lowest indexed project in S_M.*

Proof. Suppose in the optimal mechanism, $v > \alpha_{M*}$. Then it must be that there is a project $j \in S_M$ such that $j > M^*$, since the expected value that type S_M gets is $> \alpha_{M*}$. So, we can perturb this mechanism as follows:$q'_{M,j} = q_{M,j} - \varepsilon$, and $q'_{M,M*} = q_{M,M*} + \varepsilon$. Clearly, the new mechanism is IC as there is no IC below M. And it results in higher expected payoff for the principal. □

Lemma 10. *Let the set of all projects in types $\{S_1, S_2 \dots S_M\}$ be $\{\alpha_1, \alpha_2, \dots, \alpha_N\}$, where $\alpha_1 < \alpha_2 \dots \alpha_N$. For any $v \in \{\alpha_1, \alpha_2, \dots, \alpha_N\}$, such that $v \leqslant \alpha_{M*}$, there exists an equilibrium of the sequential delegation game in which the principal attains the payoff from the optimal mechanism corresponding to v, as $\delta \to 1$.*

Proof. Recall that for any project i, $\frac{\pi_1 - \pi_i}{\alpha_i - \alpha_1} = K$ where K is a constant. Let $v = \alpha_k < \alpha_{M*}$. For each S_i, we solve the following problem:

$$\begin{aligned} \max_{\{q_{i,j} | j \in S_i\}} \quad & \sum_{j \in S_i} q_{i,j} \pi_j \\ \text{subject to} \quad & \sum_{j \in S_i} q_{i,j} \alpha_j = \alpha_k \end{aligned} \tag{A.1}$$

There are two possibilities: Either $i^* < k$ or $i^* \geqslant k$. If $i \geqslant k$, the solution is $q^*_{i,i*} = \frac{\alpha_k}{\alpha_i}$, because we cannot do better than assigning positive probability to *only* the principal-favorite project in S_i, which is i^*, and since $i^* \geqslant k \implies \alpha_i \geqslant \alpha_k$, we can do so.

Now, let us consider the case where $i < k$. Here, $\alpha_{i*} < \alpha_k$, so we can no longer assign positive probability only to i^* in S_i. In this case, any solution to the above optimization problem must satisfy $\sum_j \{q^*_{i,j} | j \in S_i\} = 1$. If $\sum \{q^*_{i,j} | j \in S_i\} < 1$, then in particular $q^*_{i,i*} < 1$, and $q^*_{i,j} > 0$ for *some* $j' > i^*$, because we must have $\sum_{j \in S_i} q^*_{i,j} \alpha_j = \alpha_k$. Fix any such $j' > i^*$, such that $q^*_{i,j'} > 0$. We can now perturb the allocation probabilities as follows: Let $q^{**}_{i,j'} = q^*_{i,j'} - \varepsilon$, $q^{**}_{i,i*} = q^*_{i,i*} + \varepsilon \frac{\alpha_j}{\alpha^*_i}$, and $q^{**}_{i,j''} = q^*_{i,j''}$ $\forall$ $j'' \neq \{i^*, j\}$. It is straightforward to check that $\sum_{j \in S_i} q^{**}_{i,j}\alpha_j = v$, the principal's expected payoff is strictly higher, and for ε small enough, $\sum_{j \in S_i} q^{**}_{i,j}\alpha_j \leqslant 1$. So, if $i < k$, we must have $\sum\{q^*_{i,j} | j \in S_i\} = 1$. The constraint in the optimization problem can be thus be rewritten substituting $q_{i,i^{**}} = 1 - \sum_{\{j \in S_i | j < i^{**}\}} q_{i,j}$, where i^{**} is the highest indexed project in S_i.

$$\sum_{\{j \in S_i | j < i^{**}\}} q_{i,j}(\alpha_{i^{**}} - \alpha_j) = \alpha_{i^{**}} - \alpha_k \tag{A.2}$$

We can also, after the same substitutions, rewrite the objective function and get:

$$\pi_{i^{**}} + \sum_{\{j \in S_i | j < i^{**}\}} q_{i,j}(\pi_j - \pi_{i^{**}}),$$

which, after substituting (A.2), is just equal to

$$\pi_{i^{**}} + \sum_{\{j \in S_i | j < i^{**}\}} q_{i,j} K(\alpha_{i^{**}} - \alpha_j) = \pi_{i^{**}} + (\alpha_{i^{**}} - \alpha_k)K$$

Observe that the last expression is a constant independent of allocation probabilities. So, *any* allocation probabilities that satisfy $\sum\{q_{i,j} | j \in S_i\} = 1$ and $\sum_{j \in S_i} q_{i,j}\alpha_j = \alpha_k$, solves the optimization problem. In particular, $q_{i,k} = 1$ solves (A.1) when $i < k$.

To sum up, for any $v \in \{\alpha_1, \alpha_2, \ldots, \alpha_N\}$, *an* optimal mechanism *corresponding to* v is as follows:

$$q^*_{i,i^*} = \frac{\alpha_k}{\alpha^*_i}, q^*_{i,j} = 0 \; \forall \; j > i^*, \; if \; i^* > k,$$

and

$$q^*_{i,k} = 1, q^*_{i,j} = 0 \; \forall \; j \neq k, \; if \; i^* \leqslant k$$

We now construct an equilibrium that replicates the payoff from the above mechanism. The construction is very similar to our *separating* equilibrium from the two-project case. Fix $v = \alpha_k$. Consider the equilibrium where on path, at $t = 0$, all types that have project k report it, and this proposal is accepted right away. For every S_i such that $i^* > k$, there exists a threshold $t^*_{i^*}(\delta)$, such that type S_i, which does not have project k, proposes i at t^*_i, which is then accepted by the principal. We define this threshold inductively:

$$t^*_{k+1}(\delta) := \min\{t : \alpha_k \geqslant \delta^t \alpha_{k+1}\},$$

and, given that we have defined t^*_{k+j}, we define t^*_{k+j+1} as follows:

$$t^*_{k+j+1}(\delta) := \min\{t : \delta^{t^*_{k+j}(\delta)} \alpha_{k+j} \geqslant \delta^t \alpha_{k+j+1}\}$$

We omit the details of the strategies, as they are very similar to the *separating* equilibrium. But intuitively, this on path behavior can be supported in equilibrium as if the principal sees a proposal $i^* > k$ before $t^*_{i^*}$, his off path belief is that it is type S_1 with probability one, and if this proposal is rejected, S_1 will propose 1 in the next period. The thresholds are such that any type that does not have k will find it optimal to propose the principal's favorite project that it has, at the appropriate threshold.

Note that in this proof we have implicitly assumed that all types where $i^* < k$ have project k. In case they do not, this construction does not work. However, the next Lemma will show that we do not have to worry about these cases; if such an $\alpha_k = v$ in the optimal mechanism, we can find another v' that attains same or strictly higher payoff for the principal, such that the optimal mechanism corresponding to this v' is implementable in equilibrium.

□

We have thus shown that every optimal mechanism must have *some* v which is the expected payoff to each type of the agent, and if the optimal mechanism has $v^* \in \{\alpha_1, \alpha_2, \ldots, \alpha_N\}$, there always exists an equilibrium where the principal attains the payoff from the optimal mechanism as $\delta \to 1$. We now show, in the next lemma, that there is always a $v \in \{\alpha_1, \alpha_2, \ldots, \alpha_N\}$ such that the optimal mechanism *corresponding to* v is indeed optimal. This would complete our proof that commitment payoff can be attained for this case of nested types.

Lemma 11. *For any v such that $v \in (\alpha_k, \alpha_{k+1})$ for some $k \in \{1, 2 \ldots N-1\}$, either v cannot be part of the optimal mechanism, or there exists $v' \in \{\alpha_1, \alpha_2 \ldots \alpha_N\}$ such that the principal's payoff from the optimal mechanism corresponding to v' is the same as his payoff from the optimal mechanism corresponding to v.*[1]

Proof. Let $v \in (\alpha_k, \alpha_{k+1})$ for some $k \in \{1, \ldots, N\}$. The principal's objective is to maximize his expected payoff by choosing implementation probabilities for each type:

$$\max \sum_{j \geqslant i^*} q_{i,j} \pi_j \quad \text{subject to} \quad \sum_{j \geqslant i^*} q_{i,j} \alpha_j = v, \quad \forall i \in \{1, \ldots, M\}$$

For all types S_i with $i^* \geqslant k+1$, the optimal mechanism assigns $q_{i,i^*} = \frac{v}{\alpha_{i^*}} < 1$ and $q_{i,j} = 0$ for all other $j > i^*$.

[1] The perturbations that we construct here will also work if $v = \alpha_k$ sor some k butt all types where $i^* < k$ do not have k.

For the rest of the types S_i with $i \leqslant k+1$, as we argued in Lemma 10, we have $\sum_j q_{i,j} = 1$. Since given the constraint that the agent's expected payoff equals v, any randomization is optimal, we consider one particular randomization as part of the optimal mechanism.

An optimal mechanism that corresponds to the expected value $v \in (\alpha_k, \alpha_{k+1})$ is as follows:

- Consider $i = \min\{i' | i'^* < k+1\}$. In this case, $(i+1)^*$ is the lowest indexed project of the type S_{i+1}, so it must be that $(i+1)^* \geqslant k+1$. Also all types above $(i+1)$ will have both i^* and i^*+1.
- for all types above i, only projects i^* and i^*+1 are implemented with positive probability, with the appropriate mixture to provide the expected payoff of v, Let these probabilities be q_{i^*} and $q_{(i+1)^*}$.
- for all types below i, only the project with the lowest index is implemented with positive probability.

We now argue that there exists some v' such that there is an IC mechanism in which each type of the agent gets v' and the principal gets a strictly higher payoff than the mechanism we describe above. Consider the following perturbation:

- for all types such that lowest indexed project is $< k+1$, implement $(i+1)^*$ with probability $q_{(i+1)^*} - \varepsilon$ and i^* with probability $q_{i^*} + \varepsilon$;
- for all types such that lowest indexed project is $\geqslant k+1$, implement this lowest indexed project i' with probability $q_{i,i'} - \frac{\varepsilon(\alpha_{(i+1)^*} - \alpha_{i^*})}{\alpha_{i'}}$.

Now the gain for the principal is

$$\sum_{i' \leqslant i} \mu_{i'} \varepsilon(\pi_{i^*} - \pi_{(i+1)^*})$$

and the loss is

$$\sum_{i' > i} \mu_i \frac{\varepsilon(\alpha_{(i+1)^*} - \alpha_{i^*})}{\alpha_{i'}} \pi_{i'}.$$

We can see that ϵ gets canceled out, and the comparison only depends on the parameters, probabilities of the types and the payoffs.

Either the gain is greater or the loss, or they are exactly equal. If the gain is greater than the loss, then the perturbed mechanism is an IC mechanism where the principal is strictly better off, and the original mechanism cannot be optimal. If the loss is greater than the gain, then we can reverse the signs of the perturbation and achieve an IC mechanism where the principal is strictly better off again, making the previous mechanism not optimal. Finally, if the gain and the loss are exactly the same, then any perturbation would result in the same expected payoff for the principal. In this case, we can perturb the mechanism such that the expected payoff for all types is α_{k+1} and this would be an optimal mechanism as well. In addition, this optimal mechanism can be implemented in an equilibrium of the game as established in Lemma (10).

□

□

Appendix B
Appendix for Chapter 2

B.0.1 Notation

I first establish some notation. For any menu $\mathcal{M}$ offered by the intermediary, and any equilibrium of the game induced by this menu, let $\Theta_{(\pi,p),\mathcal{M}} = \{(\theta| \ \sigma((\pi, p)|\theta) > 0\}$, so $\Theta_{(\pi,p),\mathcal{M}}$ is the set of all types that in equilibrium, choose allocation $(\pi, p) \in \mathcal{M}$ with positive probability. Recall that $P_{\mathcal{M}} = \{p| \ \exists \ (\pi', p') \in \mathcal{M} \ with \ p' = p\}$: this is the set of all possible prices in the menu offered by the intermediary.

It is without loss to consider menus such that for any price p, there is at most one allocation in the menu with this price. Let $P_\theta = \{p| \ \sigma((\pi, p)|\theta) > 0\}$; this is the set of all prices such that on path, type θ chooses an allocation with this price with positive probability. For any on path price $p \in \bigcup_\theta P_\theta$, let $\mathbb{E}[\theta|p]$ denote the expected value of types that choose to sell at this price in equilibrium. Since this is an on path price, this Expectation is derived from the seller's strategy using Bayes Rule. Recall that the seller's strategy is given by $\sigma_S = (\sigma(.), \gamma_I(.), \gamma_M(.), \gamma_M'(.))$, where $\gamma_I : \Theta \times \mathcal{M} \to [0, 1]$, and $\gamma_I((\pi, p), \theta)$ is the probability with which a seller of type θ chooses to sell through the intermediary, *conditional* on choosing (π, p) in the menu, *and* having the option to sell through the intermediary. So,

$$\mathbb{E}[\theta|p] = \frac{\sum_{\{\theta|p \in P_\theta\}} \mu(\theta)\sigma((\pi, p)|\theta)\gamma_I((\pi, p), \theta)(1 + \alpha)\theta}{\sum_{\{\theta|p \in P_\theta\}} \mu(\theta)\sigma((\pi, p)|\theta)\gamma_I((\pi, p), \theta)}$$

For any $p \in P_{\mathcal{M}}$, $\pi_p(\theta)$ denotes the probability with which type θ sells at price p in equilibrium, and $\pi_M(\theta))$ denotes the probability with which θ sells on the market. So, for any p, if $(\pi, p) \in \mathcal{M}$ is the allocation in the menu with this price, then $\pi_p(\theta) = \pi\sigma((\pi, p)|\theta)\gamma_I((\pi, p), \theta)\sigma_B(p)$. Similarly, $\pi_M(\theta) = \sum_{\{(\pi,p) \in \mathcal{M}\}} \sigma((\pi, p)|\theta)\{\pi\gamma_M((\pi, p), \theta) + (1 - \pi)\gamma_M'((\pi, p), \theta)\}$. Let Θ_+ be the set of all types that trade with positive probability in equilibrium. So, since $\sum_{P_{\mathcal{M}}} \pi_p(\theta) + \pi_M(\theta)$ is the total probability with which type θ trades in equilibrium, $\Theta_+ = \{\theta | \sum_{P_{\mathcal{M}}} \pi_p(\theta) + \pi_M(\theta) > 0\}$. Also, let $P_{(\mathcal{M},+)}$ be the set of all prices in the menu such that in equilibrium, trade happens at these prices with positive probability, through the intermediary. So, $P_{(\mathcal{M},+)} = \{p \in \mathcal{M} | \sum_{\theta} \pi_p(\theta) > 0\}$.

B.0.2 Two Useful Results

Before proving the Theorems, I state two results that will be useful for proving the Theorems. I provide the proof for these results at the end, after the proofs of the Theorems.

We begin with a useful simplification. We can restrict attention to equilibria where the seller's strategy is such that for any θ, and any (π, p) such that $\sigma((\pi, p)|\theta) > 0$, $\gamma_I((\pi, p), \theta)$=1. So, it is without loss to restrict attention to equilibria in which if the seller chooses an allocation (π, p) with positive probability in equilibrium, then given the opportunity to sell through the intermediary at price p, the seller will do so with probability one. To state the result formally, let us define when two equilibria are outcome equivalent.

Fix any two menus $\mathcal{M}$, and $\mathcal{M}'$, and an equilibrium of the game induced by each of these menus. Let $P_{\mathcal{M},\mathcal{M}'} = P_{\mathcal{M}} \bigcap P_{\mathcal{M}'}$, the prices that are part of *both* menus. For the menu induced by $\mathcal{M}$, let the market price be given by p_M, and $\pi_p(\theta)$ denotes the probability with which type θ sells at price p in equilibrium, and $\pi_M(\theta))$ denotes the probability with which θ sells on the market. Similarly, for the menu induced by $\mathcal{M}'$, let the market price be given by p'_M, and $\pi'_p(\theta)$ denotes the probability with which type θ sells at price p in equilibrium, and $\pi'_M(\theta))$ denotes the probability with which θ sells on the market

Definition 8. *The two equilibria are outcome equivalent if i)* $p_M = p'_M$*, ii)in each equilibria, trade happens with positive probability at the same set of prices, that are in both menus, i.e.* $P_{(\mathcal{M},+)} = P_{(\mathcal{M}',+)} \subseteq P_{\mathcal{M},\mathcal{M}'}$*, and iii) for any* θ*, and any* $p \in P_{\mathcal{M},\mathcal{M}'}$*,* $\pi_p(\theta) = \pi'_p(\theta)$*, and* $\pi_M(\theta)) = \pi'_M(\theta))$*.*

Proposition 5. *Fix a menu* $\mathcal{M}$*, and an equilibrium of the game induced by this menu. Suppose there exists a* θ*, and a* $(\pi, p) \in \mathcal{M}$ *such that* $\sigma((\pi, p)|\theta) > 0$*,* $\gamma_I((\pi, p), \theta) < 1$*. Then we can construct another equilibrium, that is outcome equivalent to this equilibrium, where* $\gamma'_I((\pi, p), \theta) = 1$

Proof. If there exists a θ, and a $(\pi, p) \in \mathcal{M}$ such that $\sigma((\pi, p)|\theta) > 0$, $\gamma_I((\pi, p), \theta) < 1$, then it must be that $p = \theta$, otherwise it cannot be sequentially rational for θ to choose

$\gamma_I((\pi,p),\theta) < 1$. This is because by Lemma 2, $p \geqslant p_M$, so $\gamma_I((\pi,p),\theta)$ cannot be lower than one because $p < p_M$. So, $\theta = p \geqslant p_M$. Consider the following modification to strategy of the seller of type θ: $\sigma'((\pi,p)|\theta) = \sigma((\pi,p)|\theta)\gamma_I((\pi,p),\theta)$, and $\gamma_I'((\pi,p),\theta) = 1$.

So, I modify type θ's strategy such that she chooses the allocation (π,p) with a strictly lower probability, and *conditional* on choosing (π,p) and having the opportunity to sell through the intermediary, it does so with probability one. The second modification in the strategy of type θ is that $\sigma'((0,0)|\theta) = \sigma((0,0)|\theta) + \sigma((\pi,p)|\theta)(1-\gamma_I((\pi,p),\theta))$. Since I reduced the probability of θ choosing (π,p), the probability of θ choosing *some* other allocation in the menu must increase. I add this residual probability, $\sigma((\pi,p)|\theta)(1-\gamma_I((\pi,p),\theta))$, to $\sigma((0,0)|\theta)$, the allocation that represents *not* participating in the intermediary's trading process.

Once these modifications are made, we can find the appropriate $\gamma_M((\pi,p),\theta)$ and $\gamma_M'((\pi,p),\theta)$ such that the *equilibrium outcome* remains same. This is because if $\theta = p$, then when it doesn't sell at p, it either sells on the market, or it does not sell at all. By increasing $\sigma((0,0)|\theta)$ to $\sigma'((0,0)|\theta)$, and then making appropriate modifications to $\gamma_M((\pi,p),\theta)$ and $\gamma_M'((\pi,p),\theta)$, we can make sure that the probability with which θ sells on the market is the same as the original equilibrium. Observe that under the modified strategies, the probability with which θ trades at p, through the intermediary, is also the same as in the original equilibrium. So, the equilibrium outcome remains the same as before. □

I now state another result that will be useful for proving Theorem 4.

Proposition 6. *Fix a menu $\mathcal{M}$ and an equilibrium of the game induced by this menu, such that in this equilibrium, trade happens both through the intermediary and on the market with positive probability. Then, it must be that $(1+\alpha)\mathbb{E}[\theta'|\theta' \leqslant p_m] > p_M$.*

B.0.3 Proof of Theorem 4

Proof. I first consider equilibria in trade trade takes place with positive probability, *both* through the intermediary, and on the market. Therefore, if the intermediary, offers the menu $\mathcal{M}$, we have that $P_{(\mathcal{M},+)} \neq \varnothing$, *and* $\sum_\theta \pi_M(\theta) > 0$. I first show that p_M is uniquely determined in any such equilibrium.

Lemma 12. *When BLC is satisfied, then in any equilibrium where trade takes place with positive probability, both through the intermediary, and on the market, $p_M = (1+\alpha)\theta_n$.*

Proof. In any equilibrium, we must have $p_M \geqslant (1+\alpha)\theta_n$ in any equilibrium, as θ_n is the lowest type. Suppose $p_M > (1+\alpha)\theta_n$. Then it must be that in equilibrium, *some* $\theta > \theta_n$ trades on the market with positive probability. So, there exists a $k \geqslant (n-1)$ such that $k = \min\{k' \geqslant (n-1) | p_M \geqslant \theta_{k'}\}$. So, θ_k denotes the highest type (recall that types with lower indices are higher), such that $p_M \geqslant \theta_k$. Therefore, $\mathbb{E}[\theta' | \theta' \leqslant p_M] = \mathbb{E}[\theta' | \theta' \leqslant \theta_k]$. From BLC, we know that $\mathbb{E}[\theta' | \theta' \leqslant \theta_k] < \theta_k$, since $k \geqslant (n-1)$. So, $\mathbb{E}[\theta' | \theta' \leqslant \theta_k] < p_M$ as well, since $p_M \geqslant \theta_k$. This is a contradiction to Proposition 6, since now, we have that $\mathbb{E}[\theta' | \theta' \leqslant p_M] < p_M$. Therefore, in equilibrium, we cannot have $p_M > (1+\alpha)\theta_n$, and p_M must be $(1+\alpha)\theta_n$ in any equilibrium.

□

Proposition 7. *When BLC is satisfied, there can be no equilibrium where trade takes place with positive probability, both through the intermediary, and on the market, and a type $\theta > \theta_n$ trades with positive probability.*

Proof. Suppose there is such an equilibrium. By Lemma 12, we have $p_M = (1+\alpha)\theta_n$. I first argue that $p_M < \theta$ for *any* $\theta > \theta_n$. To see this, consider the lowest two types, θ_{n-1}, and θ_n. By BLC, we have that $(1+\alpha)\mathbb{E}[\theta' | \theta' \in \{\theta_{n-1}, \theta_n\}] < \theta_{n-1}$, which implies that $(1+\alpha)\theta_n < \theta_{n-1}$. Since $p_M < \theta_{n-1}$, it is also strictly lower than any other $\theta > \theta_n$. I divide the proof in steps:

Step 1: *Any type $\theta > \theta_n$ that trades with positive probability in equilibrium, must trade through the intermediary, at a price strictly greater than $p_M = (1+\alpha)\theta_n$.*

This follows directly from the fact that p_M is strictly lower than any $\theta > \theta_n$.

Step 2: *For any $p' \in P_{(\mathcal{M},+)}$, $p' > p_M$.*

By **Step 1**, $\theta > \theta_n$ can *only* trade through the intermediary, at a price strictly greater than p_M, so there exists an allocation $(\pi, p) \in \mathcal{M}$ such that $p > p_M$, and $p \in P_{(\mathcal{M},+)}$, i.e., there is some price that's strictly greater than p_M, at which trade takes place with positive probability. The claim in **Step 2** then follows from Lemma 2.

Step 3: *In equilibrium, there exists an allocation $(\pi', p') \in \mathcal{M}$, such that for θ_n, $\sigma((\pi', p') | \theta_n) = 1$.*

Suppose θ_n chooses more than one allocation in $\mathcal{M}$ with positive probability in equilibrium.

Then, θ_n must be indifferent between all these allocations. By Lemma 3, that the *effective type* of θ_n is $p_M = (1+\alpha)\theta_n$. So, if θ_n is indifferent between allocations (π', p'), and (π'', p''), such that $p' < p''$, and $\pi' > \pi''$, then it is *as if* the intermediary is operating in isolation, and there is a hypothetical type p_M, which is indifferent between these allocations:

$$\pi'(p' - p_M) = \pi''(p'' - p_M)$$

This implies that that for any $\theta > p_M$, (π'', p'') is strictly preferred to (π', p'), by Lemma 1. Since any $\theta > \theta_n$ is also strictly greater than p_M, therefore, in equilibrium, *no* $\theta > \theta_n$ will choose (π', p') with positive probability. But if only θ_n is selling at p', then the buyer wouldn't buy, as $p' > (1+\alpha)\theta_n$. This is a contradiction.

Step 4: *Let* $(\pi', p') \in \mathcal{M}$ *be such that for* θ_n, $\sigma((\pi', p')|\theta_n) = 1$. *Then,* $\{\theta > \theta_n | \sigma((\pi', p')|\theta) > 0\} \neq \emptyset$. *Also, if* $\theta^{'} > \theta > \theta_n$, *and* $\sigma((\pi', p')|\theta^{'}), \sigma((\pi', p')|\theta) > 0$, *then it must be the case that* $\sigma((\pi', p')|\theta) = 1$.

Since $p' > (1+\alpha)\theta_n$, therefore, for the buyer to buy at this price, some $\theta > \theta_n$ must choose the allocation (π', p') with positive probability in equilibrium. Suppose $\theta^{'} > \theta > \theta_n$ and $\sigma((\pi', p')|\theta^{'}) > 0$, and $\sigma((\pi', p')|\theta) < 1$. Then, there must be another allocation, $(\pi^{''}, p^{''})$, such that $\sigma((\pi^{''}, p^{''})|\theta) > 0$. Therefore, θ is indifferent between (π', p') and $(\pi^{''}, p^{''})$. Then, there can be two cases.

Either, $\pi' > \pi^{''}$, and $p' < p^{''}$. In this case, θ is indifferent, so $\theta^{'}$ would strictly prefer $(\pi^{''}, p^{''})$ to (π', p'), which is a contradiction, as $\sigma((\pi', p')|\theta^{'}) > 0$. The second case is that $\pi' < \pi^{''}$, and $p' > p^{''}$. In this case, θ_n strictly prefers $(\pi^{''}, p^{''})$, which is again a contradiction. So, we must have $\sigma((\pi', p')|\theta) = 1$.

From **Step 4**, we derive the desired contradiction to the fact that there exists an equilibrium, where trade takes place both through the intermediary and on the market, *and* types higher than θ_n trade with positive probability. By **Step 3**, in this equilibrium, there is an allocation (π', p') such that θ_n chooses this allocation with probability one. **Step 4** implies that there exists a highest type $\theta^* > \theta_n$, such that in equilibrium, $\sigma((\pi', p')|\theta^*) > 0$, and for all $\theta < \theta^*$, $\sigma((\pi', p')|\theta) = 1$. So, in equilibrium, conditional on price p', the maximum value of buyer's expected value for the good is $\mathbb{E}[\theta|\theta \leqslant \theta^*]$.[1] But, by BLC, $\mathbb{E}[\theta|\theta \leqslant \theta^*] < \theta^*$, which violates the IR for type θ^*. This is the desired

[1]It need not be equal to this, as θ^* can randomise.

contradiction, and completes the proof of Proposition 7.

□

Now suppose we look at equilibria where trade takes place only through the intermediary.

Proposition 8. *When BLC is satisfied, there is no equilibrium in which trade takes place only through the intermediary, and types $\theta > \theta_n$ trades with positive probability.*

Proof. Suppose by contradiction, that there exists an equilibrium where trade takes place only through the intermediary, and types $\theta > \theta_n$ trades with positive probability. Since no trade takes place on the market, p_M is determined by off path beliefs and does not have to be equal to $(1+\alpha)\theta_n$. But it must be that $p_M \geqslant (1+\alpha)\theta_n$. So, $\theta_n < p_M$.

In this equilibrium, it must be that θ_n must choose an allocation (π', p') with probability one, such that $\pi' = 1$, and $p' > (1+\alpha)\theta_n$. π' must be one, because if $\pi' < 1$, then since $p_M > \theta_n$, θ_n *will* sell on the market if trade does not happen through the intermediary. But in equilibrium, no trade takes place on the market. Secondly, it must be that $p' > (1+\alpha)\theta_n$. This is because $\theta > \theta_n$ trade with positive probability in equilibrium, so there exists an allocation (π, p) such that $p \geqslant \theta_{n-1} > (1+\alpha)\theta_n$, and $\sigma_B(p) = 1$. Therefore, θ_n would never choose (π', p') if $p' = (1+\alpha)\theta_n$.

So, suppose θ_n chooses (π', p') with probability one in equilibrium. Since $p' > (1 = \alpha)\theta_n$, then, we can show, by using an argument like *Step 4* of Proposition 7, that there exits type θ^*, such that in equilibrium, θ^* is the highest type that chooses (π', p') with positive probability, and all $\theta < \theta^*$ choose (π', p') with probability one. Thus we have the desired contradiction: conditional on p', buyer's expected value for the good is at most $\mathbb{E}[\theta|\theta \leqslant \theta^*]$. But, by BLC, $\mathbb{E}[\theta|\theta \leqslant \theta^*] < \theta^*$, which violates the IR for type θ^*. This is the desired contradiction.

□

Propositions 7 and 8 complete the proof of sufficiency in Theorem 4. There can be two possible kinds of equilibria: where trade takes place both through intermediary and market, and where trade takes place only through the intermediary.[2] I show, by Propositions 7 and 8 respectively, that there can be no equilibrium of either kind where $\theta > \theta_n$ trade with positive probability, if the BLC is satisfied.

I now provide the proof of necessity: If the BLC is *not* satisfied, then there is always an equilibrium where $\theta > \theta_n$ trade with positive probability.

[2]Technically, there can also be equilibria where trade takes place *only* on the market. But when BLC is satisfied, in such equilibria, $\theta > \theta_n$ can never trade.

Proposition 9. *If the prior does not satisfy the BLC, there is always an equilibrium where types greater than θ_n trade with positive probability.*

Proof. I now construct such an equilibrium. Since the BLC is not satisfied, there exists a $\theta' > \theta_n$, such that $(1+\alpha)\mathbb{E}[\theta|\theta \leqslant \theta'] \geqslant \theta'$. Let $\theta^* = \max\{\theta'|(1+\alpha)\mathbb{E}[\theta|\theta \leqslant \theta'] \geqslant \theta'\}$. Let $\theta^{**} = \min\{\theta|\theta > \theta^*\}$, so θ^{**} is the next higher type after θ^{**}. Observe that such a type always exists, since the prior satisfies the *lemons condition*, so θ^* cannot be θ_1. Now consider the following menu: $\mathcal{M} = \{(1, (1+\alpha)\mathbb{E}[\theta|\theta \leqslant \theta^*])\}$.[3]

If the intermediary offers this menu, there exists an equilibrium where 1) No trade takes place on the market, 2) All $\theta \leqslant \theta^*$ trade with probability one and 3) $p_M = (1+\alpha)\theta_n$.

To see this, suppose $p_M = (1+\alpha)\theta_n$, and the buyer's strategy is $\sigma_B((1+\alpha)\mathbb{E}[\theta|\theta \leqslant \theta^*]) = 1$. Clearly, $(1+\alpha)\mathbb{E}[\theta|\theta \leqslant \theta^*] > p_M$, so any type would prefer to sell at $(1+\alpha)\mathbb{E}[\theta|\theta \leqslant \theta^*]$, as opposed to p_M. Now observe that by definition of θ^*, it must be that $\theta^{**} > (1+\alpha)\mathbb{E}[\theta|\theta \leqslant \theta^*]$. This is because $(1+\alpha)\mathbb{E}[\theta|\theta \leqslant \theta^{**}]$ is a convex combination of $(1+\alpha)\mathbb{E}[\theta|\theta \leqslant \theta^*]$ and $(1+\alpha)\theta^{**}$. So, if $\theta^{**} \leqslant (1+\alpha)\mathbb{E}[\theta|\theta \leqslant \theta^*]$, then $(1+\alpha)\mathbb{E}[\theta|\theta \leqslant \theta^{**}] > \theta^{**}$, and this contradicts the fact that $\theta^* = \max\{\theta'|(1+\alpha)\mathbb{E}[\theta|\theta \leqslant \theta'] \geqslant \theta'\}$. So, the set of types that sell at $(1+\alpha)\mathbb{E}[\theta|\theta \leqslant \theta^*]$ is $\{\theta \leqslant \theta^*\}$. Therefore, given p_M, and the buyer's strategy, types $\theta \leqslant \theta^*$ would choose to sell at price $(1+\alpha)\mathbb{E}[\theta|\theta \leqslant \theta^*]$. Since they are able to sell at this price with probability one, no trade takes place through the market. For the buyer, his strategy of buying at $(1+\alpha)\mathbb{E}[\theta|\theta \leqslant \theta^*]$ is obviously optimal, given the seller's strategy. Lastly, $p_M = (1+\alpha)$ is determined by the off-path belief that if the seller is selling on the market, her type must be θ_n.

□

This completes the proof of Theorem 4.

□

B.0.4 Proof of Proposition 6

Proof. Suppose BLC is satisfied, and fix an equilibrium trade occurs with positive probability both through the intermediary and the market. There are two possible cases: $\Theta_+ = \{\theta'|\theta' \leqslant p_M\}$, i.e., there is *no* θ' strictly greater than p_M that trades with positive probability in equilibrium, or there exists a $\theta \in \Theta_+$ such that $\theta > p_M$. I consider these two cases separately, and show that in each case, we must have $\mathbb{E}[\theta'|\theta' \leqslant p_M] \geqslant p_M$.

Lemma 13. *Suppose $\Theta_+ = \{\theta'|\theta' \leqslant p_M\}$. Then, it must be that $\mathbb{E}[\theta'|\theta' \leqslant p_M] \geqslant p_M$.*

[3] Technically, $(0, 0)$ is always in the menu, but I don't write it explicitly here.

Proof. Suppose, by contradiction, that in equilibrium, $\mathbb{E}[\theta'|\theta' \leqslant p_M] < p_M$. Starting with with assumption, we will reach a contradiction. The proof proceeds in the following steps:

Step 1: *$P_{(\mathcal{M},+)}$ cannot contain more than one price.*

Suppose not, i.e., $P_{(\mathcal{M},+)}$ contains more than one price. By Lemma 2, $p' \geqslant p_M$ for all $p' \in P_{(\mathcal{M},+)}$. Let p, p'' be two prices in $P_{(\mathcal{M},+)}$, and without loss, let $p'' > p' \geqslant p_M$. By Lemma 2, this implies that $p' > p_M$ as well. Therefore, $p' > p_M$ for all $p' \in P_{(\mathcal{M},+)}$.

Since at any $p' \in P_{(\mathcal{M},+)}$, trade is taking place with positive probability (by definition of $P_{(\mathcal{M},+)}$), it must be that that for any such p', $\sigma_B(p') = 1$, i.e the buyer's strategy must be to buy at all these prices. Therefore, at any $p' \in P_{(\mathcal{M},+)}$, the buyer's interim IR must be satisfied. So, $(1+\alpha)\mathbb{E}[\theta'|p'] \geqslant p'$ for any $p' \in P_{(\mathcal{M},+)}$. Since $p' > p_M$ for any $p' \in P_{(\mathcal{M},+)}$, we have that $(1+\alpha)\mathbb{E}[\theta'|p'] > p_M$ for all $p' \in P_{(\mathcal{M},+)}$.

Since $p' > p_M$ for all $p' \in P_{(\mathcal{M},+)}$, any $\theta \in \Theta_+$ will sell on the market, only if she chooses an allocation in the intermediary's menu and does not get the option to trade through the intermediary. Therefore no type would choose to trade directly on the market, i.e. for all $\theta \in \Theta_+$, $\sigma((0,0)|\theta) = 0$. So, all types in the set $\{\theta'|\theta' \leqslant p_M\}$ choose *some* $p' \in P_{(\mathcal{M},+)}$ in equilibrium, i.e. $\sum_{p' \in P_{(\mathcal{M},+)}} \sigma((\pi', p')|\theta) = 1$. So, by law of total expectation, we have that

$$(1+\alpha)\mathbb{E}[\theta'|\theta' \leqslant p_M] = \sum_{p' \in P_{(\mathcal{M},+)}} \sum_{\theta' \leqslant p_M} \mu(\theta')\sigma((\pi', p')|\theta')(1+\alpha)\mathbb{E}[\theta'|p']$$

But as we argued earlier, by the buyer's interim IR, $(1+\alpha)\mathbb{E}[\theta'|p'] > p_M$ for all $p' \in P_{(\mathcal{M},+)}$, so this implies that $(1+\alpha)\mathbb{E}[\theta'|\theta' \leqslant p_M] > p_M$, which is a contradiction.

Step 2: *If $P_{(\mathcal{M},+)}$ is a singleton, $\{p'\}$, then $p' = p_M$.*

If $p'' > p_M$, by the buyer's interim IR, it must be that $(1+\alpha)\mathbb{E}[\theta'|p'] \geqslant p'' > p_M$. But then, as before, all types in $\Theta_+ = \{\theta'|\theta' \leqslant p_M\}$ will choose the allocation (π', p') with probability one. But this implies that $(1+\alpha)\mathbb{E}[\theta'|\theta' \leqslant p_M] > p_M$, which is a contradiction. Now, p' cannot be strictly lower than p_M, so it must be that $p' = p_M$.

I now show that **Step 2** leads to a contradiction. The only possibility is that $P_{(\mathcal{M},+)} = \{p_M\}$. So, every $\theta \leqslant p_M$ is indifferent between trading directly on the market, i.e.

choosing $(0,0)$ in $\mathcal{M}$, or choosing (π', p'), and selling on the market if trade does not happen through the intermediary. For any $\theta \leqslant p_M$, $\sigma((\pi', p'|\theta) + \sigma((0,0)|\theta) = 1$.

Observe that since $p' = p_M$, therefore, by the buyer's interim IR, $(1+\alpha)\mathbb{E}[\theta|p'] \geqslant p_M$. So, since $(1+\alpha)\mathbb{E}[\theta|\theta \leqslant p_M] < p_M$, therefore, by Law of Total Expectation, we have that $(1+\alpha)\mathbb{E}[\theta|(0,0)] < p_M$, where $\mathbb{E}[\theta|(0,0)]$ is the expected value of the seller's type, conditional on choosing $(0,0)$ in the menu.

The price p_M on the market, is determined by the market clearing condition. Observe that for any $\theta < p_M$, $\gamma_M^{'}((\pi', p'), \theta) = \gamma_M^{'}((0,0), \theta) = 1$. If there is a $\theta = p_M$, then this type may randomise between selling and not selling on the market, but we assume that it always sells (nothing will change if we don't assume this, its just to simplify notation).

$$p_M = \frac{(1-\pi')\sum_{\theta \leqslant p_M} \mu(\theta)\sigma((\pi', p')|\theta)(1+\alpha)\theta + \sum_{\theta \leqslant p_M} \mu(\theta)\sigma((0,0)|\theta)(1+\alpha)\theta}{(1-\pi')\sum_{\theta \leqslant p_M} \mu(\theta)\sigma((\pi', p')|\theta) + \sum_{\theta \leqslant p_M} \mu(\theta)\sigma((0,0)|\theta)}$$

$$\implies \sum_{\theta \leqslant p_M} \mu(\theta)\sigma((0,0)|\theta)\{p_M - (1+\alpha)\theta)\} = (1-\pi') \sum_{\theta \leqslant p_M} \mu(\theta)\sigma((\pi', p')|\theta)\{(1+\alpha)\theta - p_M\}$$

$$\implies \sum_{\theta \leqslant p_M} \mu(\theta)(1-\sigma((\pi', p')|\theta))\{p_M - (1+\alpha)\theta)\} = (1-\pi') \sum_{\theta \leqslant p_M} \mu(\theta)\sigma((\pi', p')|\theta)\{(1+\alpha)\theta - p_M\}$$

In the above equation, the RHS is equal to $(1-\pi')(\sum_{\theta \leqslant p_M} \mu(\theta)\sigma((\pi', p')|\theta))\{\mathbb{E}[\theta|p'] - p_M\}$, which is non negative, as $\mathbb{E}[\theta|p'] \geqslant p_M$. Now, consider the LHS:

$$\sum_{\theta \leqslant p_M} \mu(\theta)(1-\sigma((\pi', p')|\theta))\{p_M - (1+\alpha)\theta)\}$$

$$= \sum_{\theta \leqslant p_M} \mu(\theta)\{p_M - (1+\alpha)\theta)\} - \sum_{\theta \leqslant p_M} \mu(\theta)\sigma((\pi', p')\{p_M - (1+\alpha)\theta)\}$$

$$= \{p_M - (1+\alpha)\mathbb{E}[\theta|\theta \leqslant p_M]\}(\sum_{\theta \leqslant p_M} \mu(\theta)) + \sum_{\theta \leqslant p_M} \mu(\theta)\sigma((\pi', p')\{(1+\alpha)\theta - p_M\}$$

$$> \sum_{\theta \leqslant p_M} \mu(\theta)\sigma((\pi', p')\{(1+\alpha)\theta) - p_M\}$$

$$\geqslant (1-\pi') \sum_{\theta \leqslant p_M} \mu(\theta)\sigma((\pi', p')\{(1+\alpha)\theta) - p_M\}$$

This is because $p_M > (1+\alpha)\mathbb{E}[\theta|\theta \leqslant p_M]$, and $\sum_{\theta \leqslant p_M} \mu(\theta)\sigma((\pi', p')\{(1+\alpha)\theta - p_M\} \geqslant 0$. Observe that the last expression is the RHS, so LHS is strictly greater than RHS, but this is a contradiction, as we started with LHS=RHS. This concludes the proof of Lemma 13. □

I now consider the case where there exists a $\theta > p_M$ in Θ_+.

Lemma 14. *Suppose there exists a $\theta \in \Theta_+$ such that $\theta > p_M$. Then, it must be that $\mathbb{E}[\theta'|\theta' \leqslant p_M] \geqslant p_M$.*

Proof. Suppose, by contradiction, that in equilibrium, $(1+\alpha)\mathbb{E}[\theta'|\theta' \leqslant p_M] < p_M$. Starting with with assumption, we will reach a contradiction. The proof proceeds in the following steps:

Step 1: *All allocations (π'', p'') that are chosen with positive probability in equilibrium must have $p'' > p_M$.*

Since $\theta > p_M$ trade with positive probability in equilibrium, there exists an allocation (π, p) such that $p > p_M$, and at p, $\sigma_B(p) = 1$. Therefore, by choosing (π, p), the seller can sell at $p > p_M$ with positive probability. Since such an allocation exists, therefore in equilibrium, *all* types must choose (π'', p'') with $p'' > p_M$.

Step 2: *In equilibrium, there must exists at least one allocation (π', p') which is chosen with positive probability by both $\theta \leqslant p_M$, and by $\theta > p_M$.*

Now, suppose there is no allocation that's chosen with positive probability by both $\theta \leqslant p_M$, and by $\theta > p_M$. So, any allocation that's chosen with positive probability in equilibrium, is either *only* chosen by $\theta \leqslant p_M$, or *only* chosen by $\theta > p_M$. Let $supp(\sigma)_{\theta \leqslant p_M} = \{(\pi, p) \in \mathcal{M} | \sigma((\pi, p)|\theta) > 0 \ for \ some \ \theta \leqslant p_M\}$. This is the set of all allocations chosen with positive probability by $\theta \leqslant p_M$. By **Step 1**, for any $(\pi, p) \in supp(\sigma)_{\theta \leqslant p_M}$, $p > p_M$. Therefore, we must have $(1+\alpha)\mathbb{E}[\theta|p] > p_M$, to satisfy the buyer's interim IR. So, $\sum_{(\pi,p) \in supp(\sigma)_{\theta \leqslant p_M}} \sigma((\pi, p)|\theta) = 1$ for every $\theta \leqslant p_M$, and for every $(\pi, p) \in supp(\sigma)_{\theta \leqslant p_M}$, we have $(1+\alpha)\mathbb{E}[\theta|p] > p_M$. Therefore, by Law of Total Expectation, we have $(1+\alpha)\mathbb{E}[\theta|\theta \leqslant p_M] > p_M$. But this is a contradiction, since since

$(1+\alpha)\mathbb{E}[\theta|\theta \leqslant p_M]$.

Step 3: *There exists exactly one allocation* (π', p') *which is chosen with positive probability by both* $\theta \leqslant p_M$*, and by* $\theta > p_M$*.*

Suppose there is more than one such allocation. Recall that the *effective type* of all $\theta \leqslant p_M$ is p_M, so "type" p_M must be indifferent between all such allocations. But then, for any two allocations, if p_M is indifferent, then by Lemma 1 any $\theta > p_M$ must strictly prefer the allocation with the lower probability of trade and higher price. This contradicts the fact that *both* allocations are chosen with positive probability by both $\theta \leqslant p_M$, and by $\theta > p_M$.

Step 4: *Let* (π', p') *denote the allocation that's chosen with positive probability, both by types* $\theta < p_M$*, and by types* $\theta > p_M$*. Then* $(1+\alpha)\mathbb{E}[\theta \leqslant p_M|p'] < p_M$*, where* $\mathbb{E}[\theta \leqslant p_M|p']$ *denotes the expected value of the seller's type, conditional on being weakly lower than* p_M*, and choosing* (π', p')*.*

This follows from Law of Total Expectation. Let $supp(\sigma)_{(\theta \leqslant p_M)}$ denote the set of allocation that's chosen by *only* types $\theta \leqslant p_M$ in equilibrium. So, the sell of *all* allocations chosen with positive probability by $\theta \leqslant p_M$ is given by $supp(\sigma)_{(\theta \leqslant p_M)} \bigcup (\pi', p')$.

Now, by **Step 1**, for any $(\pi, p) \in supp(\sigma)_{(\theta \leqslant p_M)}$, $p > p_M$. So, to satisfy the buyer's interim IR at p, we must have $(1+\alpha)\mathbb{E}[\theta|p] > p_M$. Recall that for any $\theta \leqslant p_M$ $\sum_{(\pi,p)\in supp(\sigma)_{(\theta \leqslant p_M)}} \sigma((\pi,p)|\theta) + \sigma((\pi',p')|\theta) = 1$. Now, the claim in **Step 4** follows from the Law of Total Expectation. Since $(1+\alpha)\mathbb{E}[\theta|\theta \leqslant p_M] < p_M$, and $(1+\alpha)\mathbb{E}[\theta|p] > p_M$ for all $(\pi, p) \in supp(\sigma)_{(\theta \leqslant p_M)}$, we have that $(1+\alpha)\mathbb{E}[\theta \leqslant p_M|p'] < p_M$.

Now I show that **Step 4** results in a contradiction. To see this, recall that the market price p_M is determined by the market clearing condition in equilibrium. Observe that for any $\theta < p_M$, $\gamma'_M((\pi', p'), \theta) = \gamma'_M((\pi, p), \theta) = 1$, for any $(\pi, p) \in supp(\sigma)_{(\theta \leqslant p_M)}$. If there is a $\theta = p_M$, then this type may randomise between selling and not selling on the market, but we assume that it always sells (nothing will change if we don't assume this, its just to simplify notation). I now denote $supp(\sigma)_{(\theta \leqslant p_M)}$ by the shorthand notation $S_{(\leqslant)}$. So,

$$p_M = \frac{\sum_{S_{(\leqslant)}} (1-\pi) \sum_{\theta \leqslant p_M} \mu(\theta)\sigma(\pi,p)|\theta)(1+\alpha)\theta + (1-\pi') \sum_{\theta \leqslant p_M} \mu(\theta)\sigma((\pi',p')|\theta)(1+\alpha)\theta}{\sum_{S_{(\leqslant)}} (1-\pi) \sum_{\theta \leqslant p_M} \mu(\theta)\sigma(\pi,p)|\theta) + (1-\pi') \sum_{\theta \leqslant p_M} \mu(\theta)\sigma((\pi',p')|\theta)}$$

$$\implies (1-\pi') \sum_{\theta \leqslant p_M} \mu(\theta)\sigma((\pi', p')|\theta)(p_M - (1+\alpha)\theta) = \sum_{S_{(\leqslant)}} (1-\pi) \sum_{\theta \leqslant p_M} \mu(\theta)\sigma(\pi, p)|\theta)((1+\alpha)\theta - p)$$

The RHS is equal to $\sum_{S_{(\leqslant)}}(1-\pi)(\sum_{\theta \leqslant p_M} \mu(\theta)\sigma((\pi, p)|\theta))\{\mathbb{E}[\theta|p] - p_M\}$, which strictly positive, as $\mathbb{E}[\theta|p] > p_M$ for every $(\pi, p) \in S_{(\leqslant)}$. From here, we can reach a contradiction in exactly the same manner as we do at the end of Lemma 13. To see this, first observe that $\pi^{'} < \pi$, and $p^{'} > p$ for all $(\pi, p) \in S_{(\leqslant)}$. This follows directly from Lemma 1, since in equilibrium, $(\pi^{'}, p^{'})$ is chosen by $\theta > p_M$ with positive probability, and any $(\pi, p) \in S_{(\leqslant)}$ is chosen by *only* $\theta \leqslant p_M$. This, in combination with the fact that $p_M > (1+\alpha)\mathbb{E}[\theta \leqslant p_M|p']$, will lead to LHS > RHS, which is the desired contradiction. This completes the proof.

□

B.0.5 Proof of Theorem 5

The proof broadly proceeds in the following steps. I first establish some properties of equilibria when the intermediary is operating alongside the market. Then, I will show that for any such equilibrium, there exists an ϵ, such that of $\tilde{\theta} - \theta < \epsilon$, then we can construct an equilibrium for when the intermediary is operating alone, that generates strictly higher surplus that this equilibrium. Finally, I will argue that since we can do this for any equilibrium, this must also be true for the surplus maximising equilibrium when the intermediary is operating alongside the market.

I first show that when the intermediary is operating alongside the market, all equilibria must have a particular form. Let $\theta^* = \min\{\theta | \theta > \tilde{\theta}\}$.

Lemma 15. *When the intermediary is operating alongside the market, in any equilibrium, $p_M \in [(1+\alpha)\theta_n, \theta^*)$. Therefore, in any equilibrium, only $\theta \leqslant \tilde{\theta}$ can trade on the market.*

Proof. This follows from Proposition 6. Recall that $\tilde{\theta}$ is the highest type θ' such that $(1+\alpha)\mathbb{E}[\theta|\theta \leqslant \theta^{'}] \geqslant \theta^{'}$. Since θ^* is the type immediately higher that $\tilde{\theta}$, by Proposition 6, it cannot be that $p_M \geqslant \theta^*$. Because if this is the case, then let $\theta^{**} \geqslant \theta^*$ be the maximum type that's (weakly) lower than p_M. Then $(1+\alpha)\mathbb{E}[\theta|\theta \leqslant p_M] = (1+\alpha)\mathbb{E}[\theta|\theta \leqslant \theta^{**}]$, but $(1+\alpha)\mathbb{E}[\theta|\theta \leqslant \theta^{**}] < \theta^{**}$ by definition of $\tilde{\theta}$. Also, $\theta^{**} \leqslant p_M$, so we have that $(1+\alpha)\mathbb{E}[\theta|\theta \leqslant \theta^{**}] < p_M$, which contradicts Proposition 6.

Therefore $p_M < \theta^*$. In any equilibrium, we must have $p_M \geqslant (1+\alpha)\theta_n$. So, $p_M \in [(1+\alpha)\theta_n, \theta^*)$. Since $\tilde{\theta} < (1+\alpha)\theta_n$, this implies that all types $\theta \leqslant \tilde{\theta}$ can trade on the market in equilibrium. And for any type $\theta' > \theta$, as I argued, it cannot be that $p_M \geqslant \theta'$. So, no such type can trade on the market in equilibrium. $\square$

Therefore,in any equilibrium, all $\theta \leqslant \tilde{\theta}$ must trade with probability one. For any equilibrium, let $\Theta_{(+,>)}$ denote the set of types (if any) that are strictly greater than $\tilde{\theta}$, and trade with positive probability in equilibrium. If $\Theta_{(+,>)} \neq \varnothing$, let $\overline{\theta} = \max\{\theta | \theta \in \Theta_{(+,>)}\}$, i.e., $\overline{\theta}$ is the highest type that trades with positive probability in equilibrium. Also, recall that $\theta^* = \min\{\theta | \theta > \tilde{\theta}\}$, so θ^* is the lowest type in $\Theta_{(+,>)}$.

I now show that any equilibrium induces a *segmentation* of types in $\Theta_{(+,>)}$.

Lemma 16. *Suppose in equilibrium, $\Theta_{(+,>)} \neq \varnothing$. Then, there exists a partition of $\Theta_{(+,>)}$, denoted by $\{\theta^1, \theta^2, \dots \theta^m\}$, where $\theta^* \leqslant \theta^1 < \theta^2 \dots < \theta^m \leqslant \overline{\theta}$, where the ith segment is given by $\Theta_i = \{\theta | \theta^{i-1} < \theta \leqslant \theta^i\}$[4] In equilibrium, all types in the same segment choose the same allocation, and if $i < i'$, then the ith segment trades with higher probability and at a lower price in equilibrium than the i'th segment.*

Proof. This follows directly from Lemma 1. First, observe that since any $\theta \in \Theta_{(+,>)}$ is strictly greater than p_M, these types can only trade through the intermediary.

In equilibrium, if a type $\theta \in \Theta_{(+,>)}$ chooses an allocation (π, p), then it weakly prefers (π, p) to all other allocations that are chosen with positive probability in equilibrium. So for any allocation (π', p') such that $\pi' < \pi$, and $p' > p$, by Lemma 1, if θ weakly prefers (π, p) to (π', p'), then all $\theta' < \theta$ strictly prefer (π, p) to (π', p'), and therefore no θ' will choose (π', p') in equilibrium. Similarly, no $\theta' > \theta$ will choose (π', p') with $\pi' > \pi$, and $p' < p$.

So, let (π^1, p^1) be the allocation chosen by θ^*, the lowest type in $\Theta_{(+,>)}$. Then no $\theta \in \Theta_{(+,>)}$ will choose an allocation with a lower price an higher allocation probability. Let θ'' be the highest type that chooses (π^1, p^1) in equilibrium. The fact that all $\theta^* < \theta < \theta''$ also choose (π^1, p^1) in equilibrium follows from Lemma 1, since both $\theta^* < \theta$, and $\theta'' > \theta$ weakly prefer (π^1, p^1) to all other allocations. Therefore, $\theta'' = \theta^1$, and $\{\theta | \theta^* \leqslant \theta''\}$ constitute the lowest segment in $\Theta_{(+,>)}$. Similarly, we can argue for the higher segments. $\square$

Now, fix an equilibrium when the intermediary is operating alongside the market. I will consider two cases and show that in each case, I can construct an equilibrium for when the intermediary is operating is isolation, which has strictly higher expected surplus from trade.

[4] if $i = 1$, then $\theta^{i-1} = \theta^*$.

Lemma 17. *Suppose $\Theta_{(+,>)}$ is empty, so only $\theta \leqslant \tilde{\theta}$ trade in equilibrium. Then, there exists an equilibrium with strictly higher surplus from trade than this one when there is no market.*

Proof. Suppose the intermediary is operating in isolation. Consider the menu $\mathcal{M} = \{(1,(1+\alpha)\theta),(\pi_H,\theta^*)$, where $\pi_H \in (0,1)$. Without going into the argument in detail (such arguments appear elsewhere in the paper), I claim that if π_H is low enough, there exists an equilibrium where $\theta \leqslant \tilde{\theta}$ trade with probability one, and type θ^* trades with probability π_H. In this equilibrium, θ^* chooses allocation (π_H,θ^*) and all $\theta \leqslant \tilde{\theta}$ choose $(1,(1+\alpha)\theta)$. Obviously, such an equilibrium results in strictly higher surplus than the original equilibrium, with the market because now, θ^* is also trading with positive probability. □

Lemma 18. *Now suppose $\Theta_{(+,>)} \neq \varnothing$. Here too, there exists an equilibrium with strictly higher surplus from trade than this one when there is no market.*

Proof. As I showed in Lemma 16, such an equilibrium consists of a partition of $\Theta_{(+,>)}$. Let the segmentation of $\Theta_{(+,>)}$ in this equilibrium be given by $\{\theta^1,\theta^2\ldots,\theta^m\}$, and let allocation chosen by segment Θ_i be denoted by (π^i,p^i). First, I fix the segmentation $\{\theta^1,\theta^2\ldots,\theta^m\}$, and provide an upper bound for the expected surplus from trade in equilibrium with this $\Theta_{(+,>)}$ and segmentation. After establishing this upper bound, I argue that this upper bound, and therefore the surplus in the original equilibrium, can be improved upon when there is no market.

In the original equilibrium, types $\theta \leqslant \tilde{\theta}$, all have an effective type that's $p_M \geqslant (1+\alpha)\theta_n$. By **Step 3** of Lemma 14, there is an allocation (π,p) that's chosen by only types $\theta \leqslant \tilde{\theta}$ in equilibrium. So, it must be that effective type p_M weakly prefers (π,p) to (π^1,p^1), the allocation chosen by Θ_1, the lowest segment of $\Theta_{(+,>)}$. This puts an upper bound on π^1, and therefore an upper bound on the probabilities of trade of all subsequent segments, since the highest type in any segment Θ_i weakly prefers (π^i,p^i) to (π^{i+1},p^{i+1}).

The upper bound on π^1 is given by $\frac{(1+\alpha)\tilde{\theta}-(1+\alpha)\theta_n}{p^1-(1+\alpha)\theta_n}$. This is because π^1 is highest when effective type p_M is indifferent between (π,p) and (π^1,p^1). Since (π,p) is chosen by only types $\theta \leqslant \tilde{\theta}$, the maximum value of p is $(1+\alpha)\tilde{\theta}$. Also, the lowest value of p_M is $(1+\alpha)\theta_n$. So, the highest value that effective type p_M can get in equilibrium, is $(1+\alpha)\tilde{\theta}-(1+\alpha)\theta_n$. The better off type p_M is, the higher we can make π^1, if we keep p^1 fixed. After this, we can inductively modify the probability of trade of each segment accordingly, so that the highest θ in any segment Θ_i is indifferent between choosing (π^i,p^i), with the modified π^i, or between (π^{i+1},p^{i+1}).

Now I argue that we can construct an equilibrium for the case where the intermediary operates in isolation, and the same segmentation of $\Theta_{(+,>)}$. To see this, observe that we can keep the segmentation fixed, and modify the probabilities of trade such that types $\theta \leqslant \tilde{\theta}$ trade with probability one at price $(1+\alpha)\mathbb{E}[\theta|\theta \leqslant \tilde{\theta}]$. Unlike the case when the market is present, now, we need to make $\tilde{\theta}$ indifferent between $(1, (1+\alpha)\mathbb{E}[\theta|\theta \leqslant \tilde{\theta}])$, and (π^1, p^1), so if we keep p^1 the same as the original equilibrium, $\pi^1 = \frac{(1+\alpha)\mathbb{E}[\theta|\theta \leqslant \tilde{\theta}] - \tilde{\theta}}{p^1 - \tilde{\theta}}$. It is easy to see that if $\tilde{\theta}$ is close enough to θ_n, then, $\frac{(1+\alpha)\mathbb{E}[\theta|\theta \leqslant \tilde{\theta}] - \tilde{\theta}}{p^1 - \tilde{\theta}} > \frac{(1+\alpha)\tilde{\theta} - (1+\alpha)\theta_n}{p^1 - (1+\alpha)\theta_n}$, the upper bound derived for the equilibrium with the market. If π^1 is strictly greater, the probability of trade of all subsequent segments can be made strictly higher. □

This completes the proof. For any equilibrium in the presence of a market, we can construct an improvement. Therefore, whatever the surplus maximising equilibrium is, we can construct an improvement over that as well.

□

B.0.6 Proof of Proposition 4

I first describe the setup and the class of mechanisms I consider formally. As in the two type example, there are two types, θ_L and θ_H, where $\theta_L < \theta_H$, and the *lemons* condition holds: $(1+\alpha)\mathbb{E}[\theta] < \theta_H$. A mechanism is denoted by (M, f, P), where M is the set of messages, P is the set of "possible" prices at which trade can take place through the intermediary, and $f : M \to \Delta(P)$, so each report is mapped to a *lottery* over prices. I use $f_m(p)$ to denote the probability with which the seller is offered the chance to sell at price p if she reports m, and for any $m \in M$, $\sum_{p \in P} f_m(p) \leqslant 1$, where the weak inequality captures the fact that with some probability, the seller may not be offered the opportunity to sell through the intermediary. I assume that P is finite, but relaxing this assumption will not change the result in Proposition 4.

The timeline for the game is the same as before: 1) the intermediary offers a mechanism, 2) the seller makes a report to the mechanism, 3) the uncertainty associated with the mechanism is resolved; the seller learns whether or not she can trade through the intermediary, and if so, at what price, and 4) the seller decides where to sell, given her choices. In this setting, when the intermediary operates alongside the market, one can show, using standard arguments that the Revelation Principle holds, and it is without loss to restrict attention to *direct mechanisms*, so $M = \{\theta_H, \theta_L\}$. A mechanism is *IC* if, in the game induced by the mechanism, each type of the seller finds it optimal to report truthfully. As before, the buyer's interim IR must be satisfied, i.e., for trade to take place at any price in equilibrium, the buyer must find it optimal to buy at that price, given her beliefs.

I now show that there is *no* direct, *IC* mechanism, such that θ_H trades with positive probability in the game induced by that mechanism. To this end, I first argue that $p_M = (1+\alpha)\theta_L$ in any equilibrium.

Lemma 19. *For any mechanism, and any equilibrium of the game induced by this mechanism,* $p_M = (1+\alpha)\theta_n$.

Proof. Fix a mechanism, and an equilibrium induced by the mechanism. For any type θ, let $\pi_p(\theta)$ denote the probability with which type θ sells through the intermediary, at price p, and $\pi_M(\theta)$ denote the probability with which θ sells on the market, in equilibrium. Therefore:

$$p_M = \frac{\mu(\theta_H)\pi_M(\theta_H)(1+\alpha)\theta_H + \mu(\theta_L)\pi_M(\theta_L)(1+\alpha)\theta_L}{\mu(\theta_H)\pi_M(\theta_H) + \mu(\theta_L)\pi_M(\theta_L)}$$

Suppose $p_M > (1+\alpha)\theta_L$. Then, it must be that $\pi_M(\theta_H) > 0$. This in turn implies that $p_M \geqslant \theta_H$, because p_M must satisfy θ_H's IR. By the *lemons condition*, $p_M \geqslant \theta_H$ implies that $\pi_M(\theta_H) > \pi_M(\theta_L)$. I now make two observations. Firstly, for the seller of type θ_L, the option to sell on the market at $p_M \geqslant \theta_H$ exists, so she will never sell through the intermediary at at $p < \theta_H$. For trade to take place at p, the buyer's interim IR must be satisfied, so if $p \geqslant \theta_H$, then $\pi_p(\theta_H) > \pi_p(\theta_L)$.

The second observation is that because $p_M > \theta_L$, so since the seller can always sell on the market, θ_L will sell with probability one *overall*, across the mechanism and the market. So, if P is the set of all prices in the mechanism, we have:

$$\sum_{p\in P} \pi_p(\theta_L) + \pi_M(\theta_L) = 1 \tag{B.1}$$

Combined with the first observation, this leads to a contradiction. This is because $\pi_p(\theta_H) > \pi_p(\theta_L)$ for every p at which θ_L trades though the intermediary, and $\pi_M(\theta_H) > \pi_M(\theta_L)$, so B.1 implies that $\sum_{p\in P} \pi_p(\theta_H) + \pi_M(\theta_H) > 1$. This cannot be the case because $\sum_{p\in P} \pi_p(\theta_H) + \pi_M(\theta_H)$ is the total probability of trade of type θ_H across the intermediary and the mechanism, and cannot exceed one.

□

So, I have shown that in any equilibrium, $p_M = (1+\alpha)\theta_L$. I will now argue that there exits no *IC* mechanism that induces an equilibrium where θ_H trades with positive probability.

Lemma 20. *There is no equilibrium where θ_H trades with positive probability.*

Proof. Fix a mechanism offered by the intermediary, and an equilibrium induced by this mechanism. Suppose, by contradiction, that θ_H trades with positive probability in equilibrium.

Firstly, observe that θ_H can only trade through the intermediary, because the *lemons condition* implies that $\theta_H > p_M$. For any $p \in P$, such that $f_{\theta_H}(p) > 0$, it must be that $p \geqslant \theta_H$, to satisfy the high type's IR. For sale to happen at any such p, it must also be the case that the buyer's interim IR at p is satisfied. So, if $p \geqslant \theta_H$, and the buyer's strategy is to buy at this price in equilibrium, then we must have:

$$\frac{\mu(\theta_L) f_{\theta_L}(p)(1+\alpha)\theta_L + \mu(\theta_H) f_{\theta_H}(p)(1+\alpha)\theta_H}{\mu(\theta_L) f_{\theta_L}(p) + \mu(\theta_H) f_{\theta_H}(p)} \geqslant p \geqslant \theta_H \tag{B.2}$$

B.2 says that the expected value of the good for the buyer, conditional on price p, must be at least θ_H. Let $P_H \subseteq P$ be the set of all prices p such that $p \geqslant \theta_H$, and in equilibrium, trades takes place with positive probability at p. So, for any $p \in P_H$, B.2 is satisfied. Because of the lemons condition, it must therefore be that $f_{\theta_L}(p) < f_{\theta_H}(p)$ for any $p \in P_H$. So, if there is a price $p \geqslant \theta_H$, then report θ_H is mapped to that price with strictly higher probability that report θ_L. Also, by the buyer's interim IR, any price not in P_H must be weakly lower than $(1+\alpha)\theta_L$, since *only* θ_L is being mapped to that price. This contradicts the fact that the mechanism is *IC*. Because type θ_L, in expectation, would get a strictly higher payoff if she reports θ_H.

□

B.0.7 Market Can Improve Surplus

I now provide a three-type example of how the the presence of the market can sometimes *obfuscate* information contained in prices in a way that improves surplus.

Proposition 10. *Consider a setting with three types $\{\theta_1, \theta_2, \theta_3\}$, where $\theta_1 > \theta_2 > \theta_3$, and the probability of type θ is denoted by $\mu(\theta)$. Suppose the prior $\mu(.)$ satisfies the following:*

- *The lemons condition:* $(1+\alpha)\mathbb{E}[\theta] < \theta_1$
- $(1+\alpha)\mathbb{E}[\theta|\theta \in \{\theta_2, \theta_3\}] > \theta_2$, *and* $\theta_2 \geqslant (1+\alpha)\theta_3$
- $(1+\alpha)\mathbb{E}[\theta|\theta \in \{\theta_1, \theta_2\}] < \theta_1$

Then, there exists an equilibrium in the presence of the market that results in strictly higher surplus than any equilibrium when the intermediary is operating alone.

Proof. Under the conditions on parameters, when the intermediary is operating in isolation, the surplus maximising equilibrium pools θ_2 and θ_3, and separates θ_1. I omit the proof for this here, but the idea is that the only other options are separating θ_3, and pooling θ_1 and θ_2, and separating all three types. The first out of these is not feasible, because of the the condition $(1+\alpha)\mathbb{E}[\theta|\theta \in \{\theta_1, \theta_2\}] < \theta_1$, and some straightforward algebra shows that the second is never optimal.

The surplus maximising menu consists of two allocations; it is given by $\mathcal{M}^{**} = \{(\pi_1, p_1), (\pi_{\{2,3\}}, p_{\{2,3\}})\}$, where $\pi_{\{2,3\}} = 1, p_{\{2,3\}} = \mathbb{E}[\theta|\theta \in \{\theta_2, \theta_3\}],\ p_1 = \theta_1$, and $\pi_H = \frac{p_{\{2,3\}} - \theta_2}{p_1 - \theta_2}$. If the intermediary offers this menu, then there exists an equilibrium in which:

- both θ_2 and θ_3 choose allocation $(\pi_{\{2,3\}}, p_{\{2,3\}})$, and trade with probability one, at price $p_{\{2,3\}}$.
- θ_1 chooses allocation (π_1, p_1), and trades with probability $\pi_1 \in (0,1)$, at price p_1.

In this equilibrium, π_1, the probability of trade of θ_1, is such that type θ_2 is indifferent between the two allocations. Also, observe that $p_{\{2,3\}}$, the price that θ_2 and θ_3 get, is is equal to $(1+\alpha)\mathbb{E}[\theta|\theta \in \{\theta_2, \theta_3\}]$; this is the *maximum* possible price these two types can get if they are selling at one price, otherwise the buyer's interim *IR* will not be satisfied.

If somehow, we could increase the price in the allocation that θ_2 and θ_3 are choosing, to $p'_{\{2,3\}} > p_{\{2,3\}}$, then to make θ_2 indifferent between the two allocations, now, ${\pi_1}' = \frac{p'_{\{2,3\}} - \theta_2}{p_1 - \theta_2}$, which is strictly greater than π_1. Thus, we would be able to increase surplus, because θ_2 and θ_3 are still trading with probability one, and θ_1 is trading at a strictly higher probability. The idea is that if θ_2 and θ_3 are able to trade at a strictly higher price, their payoff from choosing the other allocation decreases, and we can increase the probability of trade for that allocation till θ_2 is indifferent again.

This is exactly what the market helps with. The intuition is the following: with the menu $\mathcal{M}^{**}$, only θ_1 chooses (π_1, p_1), and since $p_1 = \theta_1$, the expected value for the good for the buyer, conditional on p_1, is $(1+\alpha)\theta_1 > p_1$. So, there is some "room" here, in the sense that if θ_2 and θ_3 were *also* choosing this allocation with a (low enough) strictly positive probability, then the buyer's interim IR would still be satisfied at p_1.

he presence of the market helps exploit this room. Consider the following menu: $\mathcal{M}' = \{(1, p''), (\pi'', p''_1)\}$, where $p''_1 = \theta_1$, $\pi'' = \frac{p'' - \theta_2}{\theta_1 - \theta_2}$, and $p'' = \frac{(1+\alpha)[\mu(\theta_2)(1-\sigma_2)\theta_2 + \mu(\theta_3)(1-\sigma_3)\theta_3]}{\mu(\theta_2)(1-\sigma_2) + \mu(\theta_3)(1-\sigma_3)}$, where $\sigma_2,\ \sigma_3$ are such that $\frac{(1+\alpha)[\mu(\theta_2)\sigma_2\theta_2 + \mu(\theta_3)\sigma_3\theta_3]}{\mu(\theta_2)\sigma_2 + \mu(\theta_3)\sigma_3} = \theta_2$, and $(1+\alpha)\mathbb{E}[\theta|p''_1] \geqslant \theta_1$. Then, the following is an equilibrium:

- $p_M = \frac{(1+\alpha)[\mu(\theta_2)\sigma_2\theta_2 + \mu(\theta_3)\sigma_3\theta_3]}{\mu(\theta_2)\sigma_2 + \mu(\theta_3)\sigma_3} = \theta_2$

- Given p_B, θ_2, and θ_3 are indifferent between the two allocations, and randomise between them as part of their equilibrium strategy.
- Strategy for θ_2: $\sigma((\pi'', \theta_1)|\theta_2) = \sigma_2$, and $\sigma((1, p'')|\theta_2) = 1 - \sigma_2$.
- Strategy for θ_3: $\sigma((\pi'', \theta_1)|\theta_3) = \sigma_3$, and $\sigma((1, p'')|\theta_3) = 1 - \sigma_3$.
- θ_1 chooses (π'', θ_1) with probability one
- For θ_2 and θ_3, if they choose (π'', θ_1), and the opportunity to trade through intermediary *not* realise, they sell on the market.

Observe that $p'' > (1 + \alpha)\mathbb{E}[\theta|\theta \in \{\theta_2, \theta_3\}]$ by Law of Iterated Expectations, since since $(1 + \alpha)\mathbb{E}[\theta|\theta \in \{\theta_2, \theta_3\}] > \theta_2$, and $\frac{(1+\alpha)[\mu(\theta_2)\sigma_2\theta_2+\mu(\theta_3)\sigma_3\theta_3]}{\mu(\theta_2)\sigma_2+\mu(\theta_3)\sigma_3} = \theta_2$. So, $p'' > p_{\{2,3\}}$. It is easy to see that given the strategies of θ_2 and θ_3, the buyer's interim IR is satisfied at p''. Also, in equilibrium, θ_2 and θ_3 trade on the market only if they choose allocation (π'', θ_1), and can't trade through the intermediary. So, p_M is indeed $\frac{(1+\alpha)[\mu(\theta_2)\sigma_2\theta_2+\mu(\theta_3)\sigma_3\theta_3]}{\mu(\theta_2)\sigma_2+\mu(\theta_3)\sigma_3} = \theta_2$. Given p_B, the strategies are also sequentially rational. So, this is an equilibrium.

Since $p'' > p_{\{2,3\}}$, we have that $\pi'' > \pi_{\{2,3\}}$. So, in this equilibrium, both θ_2 and θ_3 still trade with probability one overall, and θ_1 trades with strictly higher probability than before. So, the expected surplus from trade is strictly higher than in the optimal equilibrium when the intermediary was operating in isolation.

So how exactly does the market help? With the market, in the equilibrium we constructed with menu $\mathcal{M}'$, the buyer's interim IR is satisfied at $p'' > (1 + \alpha)\mathbb{E}[\theta|\theta \in \{\theta_2, \theta_3\}] > \theta_2$ because of the manner in which θ_2 and θ_3 randomise between the two allocations. It is important to note that *both* θ_2 and θ_3 cannot be indifferent between the two allocations in $\mathcal{M}'$. This is because in absence of the market, for any two allocations, if θ_2 is indifferent between them, then θ_3 strictly prefers the one with the higher allocation probability. So, without the market, it cannot be that in equilibrium, *both* these types randomise between the *same* two allocations. So, with menu $\mathcal{M}^{**}$, θ_2 is indifferent and can randomise, but if only θ_2 randomises, the expected value conditional on $p_{\{2,3\}}$ becomes strictly lower than $(1 + \alpha)\mathbb{E}[\theta|\theta \in \{\theta_2, \theta_3\}]$, which defeats the purpose of trying to increase $p_{\{2,3\}}$ while satisfying the buyer's interim IR.

Now come back to the case where the market is present and intermediary offers menu $\mathcal{M}'$. Consider the equilibrium I constructed. Since $p_M = \theta_2$, the *effective* type of both θ_2 and θ_3 is θ_2, and it is indeed optimal for both of them to randomise between the two allocations in $\mathcal{M}'$. This completes the discussion for why market can improve surplus.

□

Bibliography

P. Aghion and J. Tirole. Formal and real authority in organizations. *Journal of Political Economy*, 105(1):1–29, 1997.

G. A. Akerlof. The market for ‘lemons’: Quality uncertainty and the market mechanism. *The Quarterly Journal of Economics*, 84(3):488–500, 1970.

S. N. Ali, N. Kartik, and A. Kleiner. Sequential veto bargaining with incomplete information. *arXiv preprint arXiv:2202.02462*, 2022.

R. Alsonso and N. Matouschek. Optimal delegation. *The Review of Economic Studies*, 75(1):259–298, 2008.

M. Armstrong and J. Vickers. A model of delegated project choice. *Econometrica*, 78(1): 213–244, 2010.

B. Balzer and J. Schneider. Belief management and optimal arbitration. *Working Paper*, 2019.

J. Bull and J. Watson. Hard evidence and mechanism design. *Games and Economic Behavior*, 58(1):75–93, 2007.

G. Calzolari and A. Pavan. On the optimality of privacy in sequential contracting. *Journal of Economic Theory*, 130(1):168–204, 2006.

G. Calzolari and A. Pavan. Sequential contracting with multiple principals. *Journal of Economic Theory*, 144(2):503–531, 2009.

G. Celik and M. Peters. Equilibrium rejection of a mechanism. *Games and Economic Behavior*, 73(2):375–387, 2011.

Y.-K. Che, W. Dessein, and N. Kartik. Pandering to persuade. *American Economic Review*, 103(1):47–79, 2013.

P. C. Cramton and T. R. Palfrey. Ratifiable mechanisms: Learning from disagreement. *Games and Economic Behavior*, 10(2):255–283, 1995.

R. Deneckere and S. Severinov. Mechanism design with partial state verifiability. *Games and Economic Behavior*, 64(2):487–513, 2008.

W. Dessein. Authority and communication in organizations. *The Review of Economic Studies*, 69(4):811–838, 2002.

P. Dworczak. Mechanism design with aftermarkets: Cutoff mechanisms. *Econometrica*, 88(6):2629–2661, 2020.

D. Gerardi, J. Hörner, and L. Maestri. The role of commitment in bilateral trade. *Journal of Economic Theory*, 154:578–603, 2014.

J. R. Green and J.-J. Laffont. Partially verifiable information and mechanism design. *The Review of Economic Studies*, 53(3):447–456, 1986.

Y. Guo and E. Shmaya. Regret-minimizing project choice. *Working Paper*, 2022.

B. Holmström. On the theory of delegation. *In Bayesian Models in Economic Theory,eds. M. Boyer and R. Kihlstrom*, pages 115–141, 1984.

W. Li. Discretion and dynamics in delegated project choice. *Working Paper*, 2022.

R. B. Myerson. Analysis of two bargaining games with incomplete information. *in A. Roth (ed.) Game Theoretic Models of Bargaining*, pages 115–147, 1985.

V. Nocke and M. D. Whinston. Dynamic merger review. *Journal of Political Economy*, 118(6):1200–1251, 2010.

V. Nocke and M. D. Whinston. Merger policy with merger choice. *American Economic Review*, 103(2):1006–33, 2013.

T. Philippon and V. Skreta. Optimal interventions in markets with adverse selection. *American Economic Review*, 102(1):1–28, 2012.

A. Rubinstein. Perfect equilibrium in a bargaining model. *Econometrica*, 50:97–109, 1982.

W. Samuelson. Bargaining under asymmetric information. *Econometrica*, 52(4):995–1005, 1984.

R. Strausz. Deterministic mechanisms and the revelation principle. *Economics Letters*, 79(3):333–337, 2003.

J. Tirole. Overcoming adverse selection: How public intervention can restore market functioning. *American Economic Review*, 102(1):29–59, 2012.

E. Van Damme, R. Selten, and E. Winter. Alternating bid bargaining with a smallest money unit. *Games and Economic Behavior*, 2(2):188–201, 1990.

Vita

Vasundhara Mallick

CONTACT INFORMATION

Department of Economics
The Pennsylvania State University
420 Kern Graduate Building, University Park, PA 16802

EDUCATION

Ph.D. in Economics, Pennsylvania State University, 2018–2024
Advisor: S. Nageeb Ali

Masters in Quantitative Economics, Indian Statistical Institute, Delhi, 2016–2018

B.A (honors) in Economics, University of Delhi, 2013–2016

FIELDS OF INTEREST

Economic Theory, Mechanism Design, Organisational Economics, Evidence Games

WORKING PAPERS

"How Markets Disrupt Mediated Trade," (Job Market Paper)

"The Art of Waiting," with Ece Teoman
Revise and Resubmit at *Journal of Law, Economics, and Organization*

HONORS AND AWARDS

Fellowship of the Indian Statistical Institute, 2016–2018
Graduate Fellowship, Penn State, 2018–present