

NORTHWESTERN UNIVERSITY

Essays on Information Acquisition

A DISSERTATION

SUBMITTED TO THE GRADUATE SCHOOL
IN PARTIAL FULFILLMENT OF THE REQUIREMENTS

for the degree

DOCTOR OF PHILOSOPHY

Field of Economics

By

Francisco Poggi

EVANSTON, ILLINOIS

June 2021

© Copyright by Francisco Poggi 2021

All Rights Reserved

Abstract

Essays on Information Acquisition

Francisco Poggi

This dissertation consists of three essays that deal with problems of information acquisition and aggregation and the incentives to acquire such information. The first of the essays is titled *The Timing of Complementary Innovations* and studies the problem of an agent that acquires information about the viability of a project while working on the complementary parts of the project. The second essay, titled *Market-Based Mechanisms* and coauthored with Quitzé Valenzuela-Stookey, looks at the general problem of designing mechanisms that use as an input the aggregate action of a large number of players. The third essay is titled *Liability Design with Information Acquisition* and is joint work with Bruno Strulovici. The chapter asks how to guarantee that agents acquire sufficient information about the risks that their actions might impose on others before acting using incentives that only when some damage occurs.

Acknowledgements

I am indebted to my advisors, Jeff Ely, Bruno Strulovici, and Wojcich Olszewski, for their guidance and support. To Eddie Dekel, Alessandro Pavan and Yingni Guo for their insightful comments. I want to thank the rest of the Theory Group at Northwestern for contributing to such an stimulating environment. To my classmates, coauthors, and friends Quitze and Ludvig for long discussions and mutual support. Finally, to my family: my wife Isabela, Andrea, Damián, Julia, Rafael, and Sofía.

Table of Contents

Abstract	3
Acknowledgements	4
Table of Contents	5
Chapter 1. The Timing of Complementary Innovations	7
1.1. Introduction	7
1.2. Model	15
1.3. Order-independence and the Efficient Allocation	17
1.4. Uncertain Rate of Success	22
1.5. Extensions	32
1.6. Conclusion	35
Chapter 2. Market-Based Mechanisms	
(Joint work with Quitzé Valenzuela-Stookey)	37
2.1. Introduction	37
2.2. The model	59
2.3. Characterization: one-dimensional state	72
2.4. Characterization: multi-dimensional state space	86
2.5. Properties and extensions	97

2.6. Applications	103
Chapter 3. Liability Design with Information Acquisition	
(Joint work with Bruno Strulovici)	112
3.1. Introduction	112
3.2. Model	116
3.3. Preliminary Analysis: Symmetric Information	119
3.4. Incentive Compatibility	122
3.5. Reducing Recklessness	130
3.6. Taxation Principle with Identifiable Information Acquisition	132
Bibliography	134
Appendix A. Appendix I	140
A.1. Omitted proofs from Section 1.3	140
A.2. Omitted proofs from Section 1.4	145
A.3. Equilibrium	149
A.4. Extensions: proofs	152
A.5. Appendix II	156
A.6. Proofs of Section 2.3	156
A.7. Bridgeability	187
A.8. Deriving R	192
A.9. Extension: no commitment	205

CHAPTER 1

The Timing of Complementary Innovations

1.1. Introduction

The aim of innovation policy is to orient scarce resources toward the most socially valuable R&D projects. The value of an innovation, however, is sometimes tied to the uncertain outcome of other developments: A new treatment for a medical condition is more valuable if there is a novel diagnosis method that helps identify the condition at an early stage, when the treatment is more effective; likewise, the novel diagnosis method is more valuable if there are new treatments for said condition; quantum software can be used to solve problems that classical computers cannot solve, but can only be implemented with quantum hardware.¹ These complementarities in innovation are important—and increasingly relevant—in many industries. Moreover, R&D projects carry high levels of uncertainty, both in terms of outcomes—the project may or may not prove successful—and in terms of costs—it is not clear how much time and how many resources will be required to complete the development.

This paper studies the development of complementary innovations, when innovation requires non-specific resources such as time, money, or attention, and the timing of innovation is driven by the allocation of these resources across the available projects. To fix ideas, consider two abstract complementary innovations, A and B . The first question is

¹Shor's quantum algorithm, a method for solving integer factorization problems in polynomial time, was written in 1994, four years before the first quantum computer prototype was developed.

what is the *efficient* way to develop these innovations. Should resources be concentrated on the development of A and then be switched to B if and only when A is successfully developed? Or vice versa? Or should both A and B be developed in parallel? Moreover, when should a project be abandoned or put on hold? This question of efficient development of innovations is relevant for innovation policy, but also for firms or venture capitalists that want to develop innovations in a way that is consistent with profit maximization.

The second question of the paper is whether the efficient allocation of resources can be implemented with private incentives. In many industries, R&D is carried out in a decentralized way, by firms or individuals with private interests. Today, we observe startups and established companies alike investing resources to develop quantum software that can only be implemented with hardware that does not yet exist, and it is not clear that it ever will.² Is the allocation carried out by private agents efficient? Moreover, how does the environment (for example, the level of concentration of the industry, the appropriability of innovations, or the requirements for patentability) affect the allocation of resources to the development of complementary innovations?

The first contribution of this paper is to introduce a tractable model that features key aspects of the R&D process when the development of innovations requires non-specific resources such as time, money, or attention. A unit of a resource (attention from now on) is allocated over two complementary projects at each point in time. For each of the projects, a success arrives discretely in the form of a breakthrough. More precisely,

²The current quantum computers seem to be insufficiently powerful to be useful. The highest integer that has been factorized using Shor's algorithm is $21 = 7 \times 3$. Moreover, it is not clear that it will be possible to develop a sufficiently powerful quantum computer, at least in the near future.

a success arrives when the cumulative amount of attention paid to a project reaches a certain level. Successes are observable, but the amount of attention required to succeed is unknown. Moreover, success amounts are independent across projects, so that working on one of the developments does not provide any information about how demanding will the other one be. The only relationship between the projects comes through the complementarity in payoffs.³ The joint value of the innovations is realized when the development stage is endogenously terminated and depends on the set of projects that was completed by that time. This paper is particularly concerned with complementary innovations, where the marginal value of completing each project is increasing in the set of projects that was successfully completed. Finally, the cost of development takes the form of a constant flow throughout the development stage.

The second contribution of the paper is to introduce techniques to solve the dynamic problem of attention allocation. More precisely, I use properties of the solution to prove an equivalence between the original allocation problem and a simpler recursive problem. For complements, a success in one project makes it more attractive to keep working on the remaining one. Proposition 1 shows that when this is the case, the dynamic problem of attention allocation is equivalent to a recursive problem in which the decision maker pledges a certain amount of attention to each of the remaining projects, and commits to abandon the projects if no success is obtained after the attention is allocated. Thus, all that matters for efficiency is how much total attention the DM is willing to pledge to each project, and not the specific order in which the agent allocates the attention. This is useful because it reduces the strategy space to a few parameters. The intuition is that

³In the extended model of Section 1.5.1, the assumption of independence is weakened.

given the complementarities in payoff, the amount that the planner is willing to work on a project if there is no new success is the minimum she is going to work on that project on any optimal strategy. Since the planner would work this amount independently of the outcome of other projects, *when* she does it is not payoff-relevant.

I analyze specifically the canonical case where projects are perfect complements and the rate of success of each project is constant over time but unknown. The beliefs about these rates evolve with the outcomes of the development process. In particular, lack of success is evidence in favor of the rates being relatively low, or in other words “the project i being relatively more challenging.”

If the rate of success for each project were known, the timing of development is irrelevant. Any project that is worth pursuing is worth completing; therefore, the order of completion is not going to affect the final expected payoff.⁴ In contrast, when the rate of success is uncertain, the order of development is relevant, since it affects the arrival of information about the unknown parameters. An initial failure to develop A not only reduces the prospects of ever completing A , but also decreases the expected returns from completing B . The problem therefore does not fit in the classical experimentation framework and thus there is no general result (such as the Gittins index) that governs the optimal dynamic allocation.⁵

Consider the case where project A is of uncertain feasibility, that is, the success rate is either zero or λ_A , and project B has a known success rate λ_B . In this case, it is efficient to first work on project A : there is no learning by working on B , so there is no efficiency loss

⁴A similar logic holds when the completion time for each project is deterministic.

⁵The problem fits the framework of *restless multi-armed bandits*, for which there is no general index solution.

in back-loading all development of B . Front-loading the development of A increases the speed of learning, which is valuable because of the option given by the stopping decision. The intuition from the previous case applies more generally: the efficient allocation of resources reflects the optimal learning process about the potential of the joint project. For symmetric projects, Proposition 2 partitions the parameter space in two. For projects with high uncertainty about the rate of success and high costs, it is efficient to work in a sequence starting with the less promising project. For other parameters, it is optimal to work on the projects simultaneously. This result is generalized to asymmetric projects in Proposition 3.

The final contribution of the paper is the study of inefficiencies that arise due to decentralization and when whether the efficient allocation can be implemented with private incentives. The allocation of private R&D resources depends on several factors: who assigns these resources, the appropriability of the innovations—which is determined by the legal and patent systems—and how informed the agents are about a given project’s successes. In Section 1.4.2, I study private allocation of resources in the extreme case of a decentralized allocation.

A *decentralized allocation* myopically allocates the resources to the projects maximizing the immediate expected return of development. In particular, the decentralized allocation does not consider how the allocation changes the dynamics of beliefs about the feasibility and difficulty of the projects.⁶ The return from innovation, in turns, captures all the

⁶The decentralized allocation is equivalent to the equilibrium outcome of an industry that consists of a continuum of firms that race to develop the innovations, and each of whom controls an equal portion of the total unit of resource available at each moment in time. The firms don’t consider the informational externalities that their actions generate.

marginal value of the innovation and a portion of the value from subsequent innovations. We consider policies that alter this proportion or *appropriation level*.

There are two inefficiencies associated with decentralization. First, an imperfect appropriability might induce underdevelopment since innovators don't consider the full value of the innovations. This inefficiency is present even fixing the timing of developments. With endogenous timing, a new type of inefficiency might appear, due to inefficient timing of development.

With substitute projects, Bryan and Lemus (2017) show how decentralization and competition biases the allocation of resources toward fast, easy projects to the detriment of harder but cost-efficient ones. Although one might think that this effect is also a concern with complements, Proposition 4 shows that this is not the case: the decentralized allocation of resources is not biased toward projects just because they are thought to be easier. In particular, if the rates of success for both projects have the same support of beliefs, the decentralized allocation of resources is efficient when the appropriation is perfect.

For projects where the beliefs have different supports, the decentralized allocation is inefficient even when appropriability is perfect. In particular, the allocation of resources is biased toward projects where learning is slower. I show that these inefficiencies disappear, however, if the stakes are sufficiently high.⁷

The remainder of the paper is as follows. In the next section, I discuss the relevant literature. Section 1.2 introduces the model. Section 1.3 shows when is it possible to solve the allocation problem by looking at a auxiliary problem. In Section 1.4, we focus on a

⁷This contrasts again with the case of substitutes, where higher stakes magnify the race effects. See Bryan et al. (2020).

canonical case of perfect complements with uncertain rates of success. In Section 1.4.1 the efficient allocation is characterized. In Section 1.4.2 we identify the inefficiencies generated by the private allocation of resources by a decentralized industry. Section 1.6 concludes.

1.1.1. Related Literature

The main contribution is to the literature that studies complementary innovations. Scotchmer and Green (1990) and Ménérier (2008) ask what is the optimal inventive requirements for a patent in the context of complementary innovations. Biagi and Denicolò (2014) study the optimal division of profits with complementary innovations. Fershtman and Kamien (1992) study the effects of cross licensing in the incentives to innovate. Bryan and Lemus (2017) study the direction of innovation in a general setting that accounts for complementary innovations. Bryan et al. (2020) focus on the effect of a *crises*—a proportional increase in the payoff from innovations—in the direction of innovation with partial substitutes. In these papers there is no learning in the development stage since the process of information arrival is memoryless.

A particular type of complementary innovation is that sequential developments or cumulative innovation. Papers that study sequential developments include Gilbert and Katz (2011) and Green and Scotchmer (1995). Moroni (2019) studies a contracting environment with sequential innovations. In these papers the timing of innovation is exogenously given. I focus on complementary innovations in which the timing is determined endogenously by the allocation of resources. To the best of my knowledge, my paper is the first one to combine an endogenous timing of development with a learning process in the development stage in the context of innovation.

Another important related literature studies the problem of dynamic information acquisition from multiple sources. With Poisson information structure, both Nikandrova and Pans (2018) and Che and Mierendorff (2017) study an agent that acquires information before an irreversible decision. Nikandrova and Pans (2018) studies the case of independent processes while Che and Mierendorff (2017) studies processes that are negatively correlated. Mayskaya (2019) also studies irreversible decision with Poisson structure in a general setting. Ke and Villas-boas (2019) study problem of independent information sources where the agent learns about the state by observing a Brownian process. Klabjan et al. (2014) study the problem of sequential acquisition of information about the attributes of an object.

Liang et al. (2018) and Liang and Mu (2020) compare the performance of optimal strategies and a different strategy. Liang et al. (2018) asks the question of how well does a strategy that neglects all dynamic considerations and acquire information in a myopic way performs with respect to the optimal information acquisition strategy. Liang and Mu (2020) compare efficient information acquisition to what results from the choices of short-lived agents who do not internalize the externalities of their actions.

The paper shares many key elements with the theory of scheduling in operations research. This literature is concerned with the problem of specifying an order in which jobs or tasks should be completed. Although there are papers in this literature that incorporate uncertainty in the amount of resources that each task demands, the type of objective functions is typically different.

1.2. Model

A decision maker (DM) can work on two research projects, A and B . The DM decides when to stop the research activity and, before that, in which way to allocate resources across the projects. Time is continuous and each instant before stopping, the DM allocates a unit of resource across the projects that were not completed so far. $\mathcal{A} := [0, 1]^2 \cup a_0$ is set of available actions at each time, where a_0 is the stopping action.

Let $\alpha_i(t)$ be the amount of resource allocated to project i at time t . Each project is completed when the cumulative resources allocated to it $X_i(t) := \int_0^t \alpha_i(\tilde{t}) d\tilde{t}$ reaches a certain amount τ_i . Project completion is observable but τ_i is unknown. Formally, the agent observes the stochastic process $S_t := \{i : \tau_i \leq X_i(t)\}$ that represents the set of projects that was completed so far. The completion times of the projects are independent with F_i the cdf of project i .

At each point in time, the agent decides either to stop or how to allocate a unit of the resource across the incomplete projects. The resource is scarce: $\sum_{i \in S(t)^c} \alpha_i(t) \leq 1$ for all t .

If the DM stops at time T , the payoff derived from the innovations is $q(S_T)$. As a normalization, if no project was completed the payoff is zero.

Assumption 1 (free disposal). q is increasing in the inclusion order, i.e.

$$q(S) \leq q(S') \quad \text{for all } S \subseteq S'.$$

We are interested in complementary innovations.

Definition 1. *The projects are complements if the function q is supermodular, that is,*

$$q(A) + q(B) \leq q(\{A, B\})$$

The projects are perfect complements if $q(A) = q(B) = 0$.⁸

Finally, using the resources is costly: there is a constant flow cost of c during the development stage, that is independent on which project the DM works on.⁹ There is no discounting.¹⁰ Thus, the payoff of an DM that stops at time t and completed projects S by that time is $q(S(t)) - c \cdot t$. The DM is an expected-payoff maximizer.

Strategies. A *strategy* is a map from the set of histories \mathcal{H} to the set of actions \mathcal{A} . A *stationary strategy* only uses part of the information contained in the history, namely the cumulative resources and the set of completed projects. In particular, in a stationary strategy the actions of the agent do not depend on the order in which resources were allocated so far. Notice that since no resources can be spent on completed projects, by knowing the set of completed tasks S and the vector of cumulative resources X , the agent can recover the completion times: $\tau_i = X_i$ for $i \in S$. Let $\mathcal{H}^\circ := 2^{\{A, B\}} \times \mathbb{R}_+^2$ be the set of stationary histories. Formally, a stationary strategy consists on a vector field for each subset of developments.

⁸Substitutes are defined by q being submodular, and perfect substitutes by the property that $q(A) = q(B) = q(\{A, B\})$.

⁹This assumption is innocuous since we can normalize the time unit for different projects, by changing the distribution of τ , so that the cost is the same for all projects.

¹⁰Alternatively, we could have had a discount factor instead of the linear cost in time. The qualitative features of the solution remain unchanged for this alternative version of the model.

Definition 2. A stationary strategy is a function $x : \mathcal{H}^\circ \rightarrow \mathcal{A}$ such that $x_i(S, X) = 0$ for all $i \in S$ and

$$\sum_{i \in K \setminus S} x_i(S, X) \leq 1 \quad \forall (S, X) \in \mathcal{H}$$

Let \mathcal{X} be the set of stationary strategies.

1.3. Order-independence and the Efficient Allocation

The problem of the decision maker is to choose a strategy to maximize their expected payoff. Since all the information about the underlying uncertainty is embedded in the stationary history, it is without loss of optimality to focus on stationary strategies. For an initial state (S, X) , strategy $x \in \mathcal{X}$, and vector τ of completion times, there is a deterministic extra time that the agent that follows strategy x will work $\tilde{T}(x, \tau)$ and deterministic set of projects that will be completed by that time $\tilde{S}(x, \tau)$. The *allocation problem* can be expressed as:

$$V(S, X) = \max_{x \in \mathcal{X}} \mathbb{E}_\tau \left[q(\tilde{S}(x, \tau)) - c \cdot \tilde{T}(x, \tau) \mid (S, X) \right]$$

A different problem is as follows. The agent decides how many resources to *pledge* to each of the remaining projects, and then allocates them in some order. If none of the projects is completed after all the pledged resources are allocated, the agent has to stop. If at least one of the projects is completed, more resources can be allocated to the remaining ones. We are going to refer to this alternative problem as the *order-independent problem*. The name stems from the fact that for a given pledge, the order in which the resources are allocated will not affect the outcome.

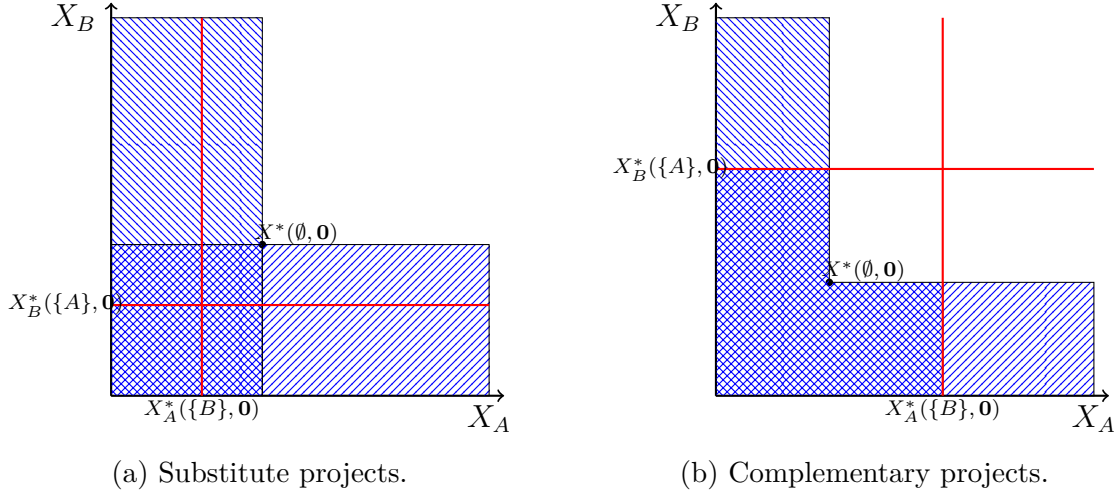


Figure 1.1. Solution to the order-independent problem. The agent completes project A (project B) when the vector of completion times falls in the area with north east (west) pattern.

Formally, the order-independent problem can be recursively defined as follows. Let $\hat{S}_X(\tau)$ be the set of completed projects at X , that is $\hat{S}_X(\tau) := \{i \in K : \tau_i \leq X_i\}$ and $D_{X,\hat{X}}(\tau)$ an indicator function that takes value 1 if there are no new completed projects from X to \hat{X} , i.e. $D_{X,\hat{X}}(\tau) := \mathbf{1}_{\{\hat{S}_{\hat{X}}(\tau) = \hat{S}_X(\tau)\}}$.

$$\hat{V}(S, X) = \max_{\hat{X} \geq X} \mathbb{E}_\tau \left[W(X, \hat{X}, \tau) - c \sum_{i \in K} (\min\{\hat{X}_i, \tau_i\} - x_i) \mid (S, X) \right]$$

where $W(X, \hat{X}, \tau) := D_{X,\hat{X}}(\tau) \cdot q(S) + [1 - D_{X,\hat{X}}(\tau)] \cdot \hat{V}(\hat{S}_{\hat{X}}(\tau), \hat{X} \wedge \tau)$.

Let $X^*(S, X)$ be the solution to the order-independent problem. It must be that $V \geq \hat{V}$ since the DM can always choose a strategy that implements the solution to the order-independent problem. We are interested in conditions on q and F such that these two problems equivalent, i.e. $V = \hat{V}$.

When completing a project induces the agent to work more on the remaining one, the order in which the agent works on the projects before the first success irrelevant modulo the cumulative work at the stopping decision. This is formalized in the following proposition.

Proposition 1. *If the projects are complements, then $V = \hat{V}$. If the projects are not complements, there exists a family of distributions $\{F_i\}$ such that $V \neq \hat{V}$.*

PROOF. in the Appendix A.1.1. Here is a sketch that provides the intuition of the result.

For any strategy, there is a (potentially infinite) amount of resources \hat{X} that the agent allocates to each project before stopping, conditional on that no project is ever completed. For complementary projects, the amount the agent allocates on a project for an *optimal* strategy must be less than the amount that the DM would be willing to allocate to the project if the other one was completed: $\hat{X}_i \leq X_i^*(\{j\}, X)$. This means that a pledge is not binding: if the agent pledges \hat{X} , independently of the outcomes of the process, he was going to put more than \hat{X}_i resources on project i .

Figure 1.1 illustrates the intuition for Proposition 1. When projects are complements, success increases the amount the agent is willing to pledge to a project. So $X_i^*(\emptyset, \mathbf{0}) \leq X_i^*(\{j\}, \mathbf{0})$, as in Figure 1.1b. This implies that the agent never regrets a pledge made. If the projects are substitutes, then $X_i^*(\emptyset, \mathbf{0}) < X_i^*(\{j\}, \mathbf{0})$, as in Figure 1.1a. At any point X with $X_i > X_i^*(\{j\}, \mathbf{0})$, the agent that completes project j would regret his pledge to project i .

Proposition 1 implies that for complements, and only for complements, it is possible to solve the dynamic optimal allocation by finding the *optimal pledge*. To emphasize the generality of the result, notice that no assumption was made on the distributions of completion times F_i . Thus, the result holds even for discrete time with arbitrary costs.¹¹

Let $h_i := F'_i/(1 - F_i)$ be the *completion rate* of project i . When the completion rate is decreasing, failures depress the prospects of each project, we can bound the stopping frontier—the set of states at which the agent stops—using the following lemma.

Lemma 1. *If h_i is decreasing for one of the projects and the projects complements, then $V(\cdot, X)$ is supermodular for all $X^*(\emptyset, X) \neq X$.*

SKETCH OF THE PROOF: Since the projects are complements, completion of one of the projects weakly increases the willingness to work on the remaining one. If at (\emptyset, X) the DM wants to stop, it must be that

$$h_j(\emptyset, X)(V(\{j\}, X) - V(\emptyset)) \leq c$$

But the agent would be weakly willing to work on project B if project A was completed.

Thus, since

$$h_j(\emptyset, X)(V(\{i\}, X) - \underbrace{V(\emptyset, X)}_{=0}) \leq c \leq h_j(\{i\}, X)(V(\{i, j\}, X) - V(\{i\}, X))$$

By independence, $(V(\{i\}, X) \leq (V(\{i, j\}, X) - V(\{i\}, X))$.

¹¹To see this just consider an F_i with mass probabilities at different times. The difference between the mass points can be interpreted as the cost of working on the project for an extra period.

Lemma 1 says that supermodularity of $V(\cdot, X)$ holds for all X where the agent wants to stop before the first success. $V(\cdot, X)$ is not supermodular for all X . For instance, for q modular we have

$$V(\{A, B\}, X) = q(\{A, B\}) = q(A) + q(B) \leq V(A, X) + V(B, X).$$

Where the last inequality holds strictly for all X where the DM would like to continue working on any of the projects. However, since for the modular case the stopping problem can be thought as independent across projects, at a stopping point $v(\{i\}, X) = q(i)$ for both projects.

Change of variables. To solve the order-independent problem, it is convenient to change variables and work with the probabilities of success. Thus, instead of choosing a pledge, the DM chooses the probability with which she wants to succeed in each project, conditional on the (lack of) success on the other project. Let $p_i = F_i(\hat{X}_i(\emptyset, \mathbf{0}))$ and $\bar{p}_i = F_i(\hat{X}_i(j, \mathbf{0}))$, the order-independent problem can be written as:

$$\max_{\bar{p}_A, p_A, \bar{p}_B, p_B} (p_A \bar{p}_B + p_B \bar{p}_A - p_A p_B) q(\{A, B\}) + \sum_{i=A, B} p_i (1 - \bar{p}_j) q(i) - \sum_{i=A, B} [p_i C_j(\bar{p}_j) + (1 - p_i) C_j(p_j)]$$

Where $C_i(p)$ is the expected cost associated with completing a project with probability p .

$$C_i(p) := \int_0^{F^{-1}(p)} (1 - F(\tau)) c \, d\tau$$

When the hazard rate of project i is decreasing, the cost function is convex, and the solution to the problem can be characterized by the first order conditions: $C'_i(\bar{p}_i) = q(\{A, B\}) - q(j)$ and $C'_i(p_i) = q(i) + \frac{\bar{p}_j - p_j}{1 - p_j} q(\{A, B\})$.

Sometimes, the solutions to the order-independent problem is such that all resources are pledged to one of the projects. We are going to say that it is efficient to work on the projects *in sequence* if for every X there exists a project i such that $X_i^*(\emptyset, X) - X_i = 0$.

In the next section, we focus on a family of canonical problems and use the results from this section to answer qualitative aspects of the solution—when is it optimal to work on the projects in sequence and when it is not—and to analyze the effects of industry concentration on the allocation of resources.

1.4. Uncertain Rate of Success

The focus of this section is on a set of canonical problems: perfect complements with constant but unknown completion rate. The agent knows that the rate takes one of two possible values, $\lambda_i \in \{\lambda_i^L, \lambda_i^H\}$ and, as before, the rates are independent across projects, with $p_i := \Pr(\lambda_i = \lambda_i^H)$.¹² As resources are allocated to a project and this is not completed, the agent becomes more pessimistic about its difficulty. Let $p_i(X)$ be the probability of the rate for project i being λ^H when no project was completed after resources X where allocated. We are going to use $\delta_i := \frac{1}{2}(\lambda_i^H - \lambda_i^L)$ and $\bar{\lambda}_i := \frac{1}{2}(\lambda_i^H + \lambda_i^L)$, and normalize time so that $\frac{1}{2}(\bar{\lambda}_A + \bar{\lambda}_B) = 1$.

Definition 3. *The projects have the same support if $\lambda_A^H = \lambda_B^H$ and $\lambda_A^L = \lambda_B^L$. If the projects have the same support, we say that project i is the most promising project at the state (\emptyset, X) if $p_i(X) > p_j(X)$.*

¹²Thus, $F_i = 1 - p_i e^{-\lambda_i^H} - (1 - p_i) e^{-\lambda_i^L}$.

1.4.1. Efficient Allocation

Observation: If project A has a known rate of success, it is efficient to work on the projects in sequence. The efficient sequence starts with the project of uncertain rate (project B).

The reason is that there is no learning by working on the project with known rate. So, for any strategy with $X_i^*(\emptyset, X) > 0$, the expected return of the same strategy with $X_i^*(\emptyset, X) = 0$ is weakly larger.

The next proposition tell us that the nature of the optimal strategy depends on measure that is increasing in the normalized cost and the uncertainty about the underlying success rate.

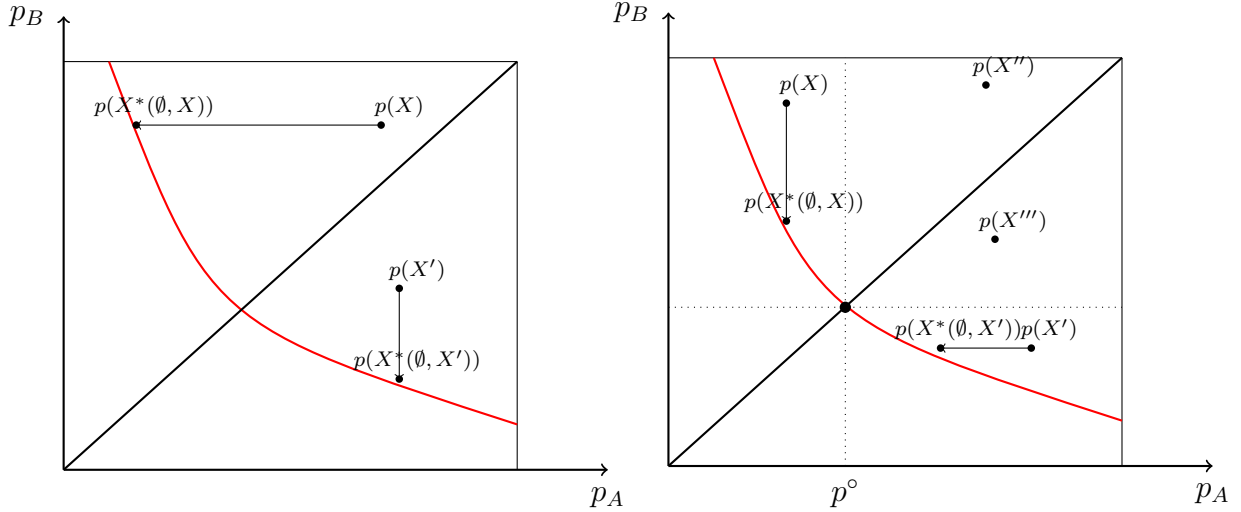
Proposition 2. *For symmetric projects, let $g := 2\frac{c}{\gamma} + \delta^2$.*

- *If $g > 1$, it is efficient to work on the projects in sequence, starting from the least promising project.*
- *If $g < 1$, it is efficient to work more on the more promising project:*

$$X_i^*(\emptyset, X) - X_i > X_j^*(\emptyset, X) - X_j \quad \Leftrightarrow \quad p_i(X) > p_j(X)$$

Moreover, there exists a $p^\circ \in (0, 1)$ such that if $p_i(X) > p^\circ$ for both projects then $p_i(X^(\emptyset, X)) = p^\circ$ for both projects. If $p_i(X) < p^\circ$ for one of the projects then it is efficient to work on the projects in sequence, starting with the most promising one.*

Figure 1.2 shows the optimal allocation of resources before the first success for different priors, in belief space. The red curve represents the boundary of the set of points at which



(a) When $g > 1$ it is optimal to work on the projects in sequence, starting from the least promising one. (b) When $g < 1$ it is optimal to work more on the more promising project.

Figure 1.2. Optimal allocation for symmetric projects in the first stage.

the agent is willing to stop before the first success $p(X^*(\emptyset, \mathbb{R}_+^2))$. In Figure 1.2a, when $g > 1$, it is optimal to work only on the project with lower prior. In Figure 1.2b, when $g < 1$, to the left of the 45° line the initial prior $p_i > p^\circ$ for $i = A, B$, then it is optimal to work on both projects before stopping in the first stage. More precisely, for an initial state (\emptyset, X) , $X_i^* = p_i^{-1}(p^\circ) - X$. Thus, there are multiple stationary strategies that are optimal. All these are payoff-equivalent by Proposition 1.

The result says that it is efficient to concentrate the resources (and therefore work in sequence) when the cost of development is sufficiently high, or the difference between the high and low rates is sufficiently large for both projects. The intuition is that in this case, having a single project that is difficult is sufficiently bad to abandon the joint project, so by concentrating the resources the DM gets to learn fast if this is the case. In the cost of development is sufficiently low, or the difference between the high and low rates is low

for both projects, then it is optimal to work on the project simultaneously. In this case, having a single project that is difficult is not sufficient to stop.

We can formalize this intuition by interpreting the result in terms of optimal information acquisition. There are three possible states: both projects are easy ($\lambda_A = \lambda_B = \lambda^H$), both are hard ($\lambda_A = \lambda_B = \lambda^L$) or one is easy and the other one hard ($\lambda_A \neq \lambda_B$). For the decision problem to be interesting it must be that the DM would be willing to work on the projects if he knew that both are easy, and he does not want to work on the projects when both of them are hard.

Suppose that the DM would be willing to work on the projects if he knew that one was difficult and the other one was easy. Then the partition of the state space that is relevant for decision making is whether there is at least one easy project (continue) or both projects are hard (abandon).

The probability of the event ‘at least one of the projects is easy’ is $p^{\text{OR}} = p_A + p_B - p_A \cdot p_B$. By assigning extra resources dX_i to project i and not succeeding, the change in p^{OR} is

$$\frac{dp^{\text{OR}}}{dX_i} = (1 - p_j) \frac{dp_i}{dX_i} = -p_i(1 - p_i)(1 - p_j)2\delta$$

The fastest way to learn about the relevant state is to work on the project with highest p , and therefore to work on the projects simultaneously.

If the DM does not want to work when one of the project is hard and the other one is easy, the relevant state is whether there is a hard project or not. There is no hard project with probability $p^{\text{AND}} = p_A \cdot p_B$. By working on project i for a period dt and not

succeeding the change in p^{AND} is

$$\frac{dp^{\text{AND}}}{dX_i} = p_j \frac{dp_i}{dX_i} = -p_i p_j (1 - p_i) 2\delta$$

The fastest way to learn about the relevant state is to work on the project with lowest probability of success, and therefore to work on the projects in sequence.

When does the DM want to continue working when one of the projects is difficult and the other one is easy? When the expected cost of completing both projects is less than the payoff from doing so, i.e.

$$\underbrace{\frac{c}{1+\delta} + \frac{c}{1-\delta}}_{\text{Expected cost}} < \underbrace{\gamma}_{\text{benefit}}$$

Rearranging we can see that this is equivalently to $g < 1$. Proposition 2 can be extended to asymmetric projects as follows.

Proposition 3. Let $g_i := 2 \frac{c/\gamma}{\lambda_i} + \frac{1}{4} \left(\frac{\delta_i}{\lambda_i} \right)^2$,

- (1) If $g_i > 1$ for both projects, then it is efficient to work on them in sequence.
- (2) If $g_i < 1$ for both projects, and $\lambda_i^L > \lambda_j^H$ then it is efficient to work on the projects in sequence, starting with the most promising one.
- (3) If $g_i < 1$ for both projects and $\lambda_i^H > \lambda_j^L$ for $i \neq j$, then there exists a $p \in (0, 1)$ such that if $p_i(X) > p$ for both projects, then for $X^* = X^*(\emptyset, X)$:

$$\frac{h_A(X^*)}{V(K \setminus \{A\}, X^*)} = \frac{h_B(X^*)}{V(K \setminus \{B\}, X^*)}$$

PROOF. in the Appendix A.2.1. Extra conditions for the case where $g_A < 1 < g_B$ can be found in Appendix A.4.3. \square

1.4.2. Decentralized Allocation

So far, we focused on the efficient allocation: how would resources be allocated by a single forward-looking decision maker that internalizes the social value of the innovations. Research and development is rarely carried out by such decision maker. The resources in the economy might be controlled by different agents, with private incentives. The level of concentration and the appropriability of the innovations affect the incentives of the agents and, ultimately, the allocation of resources to different projects.

Imperfect appropriation. Even fixing a sequence of innovations, decentralization might generate inefficiencies due to imperfect appropriation. The agent who develops an innovation might not capture how his innovation increases the value of subsequent innovations. Efforts to develop an innovation might stop inefficiently early. A way to fix this inefficiency is to give compensate early developers for the subsequent development of complementary innovations.

Compensating innovators for subsequent innovations, however, opens the door to a different type of inefficiency when the timing of development is endogenous. Innovators might be tempted to work on easy innovations to capture rents from the posterior development of difficult complementary innovations. Thus, the decentralized resource allocation might be inefficient when it is efficient to start with the difficult innovations.

Decentralized allocation. In this section, we consider a family of strategies that, at each state, maximize a flow payoff. The flow payoff depends on the reward that is expected from

innovation. The reward function is parametrized by the share of value that appropriation level α . These reward functions consider both the current marginal value of the innovation as well as the potential marginal value, in the case that complementary technologies are developed.

Definition 4. *A strategy x is the decentralized allocation for appropriation α if for each state (S, X) maximizes*

$$\sum_{i=A,B} x_i [h_i(X) \cdot R_i^\alpha(S, X) - c]$$

Where

$$R_i^\alpha(S, X) := \underbrace{[q(S \cup \{i\}) - q(S)]}_{\text{immediate}} + \alpha \underbrace{[V(S \cup \{i\}, X) - q(S \cup i)]}_{\text{potential}}$$

We can decompose an innovation's worth in two: an *immediate worth*, how useful it is given the current technology, and a *potential worth* in the form of future uses. The parameter α captures the level of appropriability of the subsequent innovation rents. When $\alpha = 0$, the decentralized allocation assigns resources to the project that creates higher expected immediate worth net of the cost of development. When $\alpha > 0$ the allocation internalizes, at least partly, the potential uses of the innovation in the future.¹³

The following result when is the allocation efficient for symmetric projects.

Proposition 4. *If the projects are symmetric, the decentralized allocation is efficient if and only if $\alpha = 1$.*

¹³The decentralized allocation with $\alpha = 1$ turns out to be the equilibrium allocation in a game with a continuum of firms that compete in the development stage when the first firm to succeed gets the whole surplus of the development process. This micro-foundation is explored in Appendix A.3.

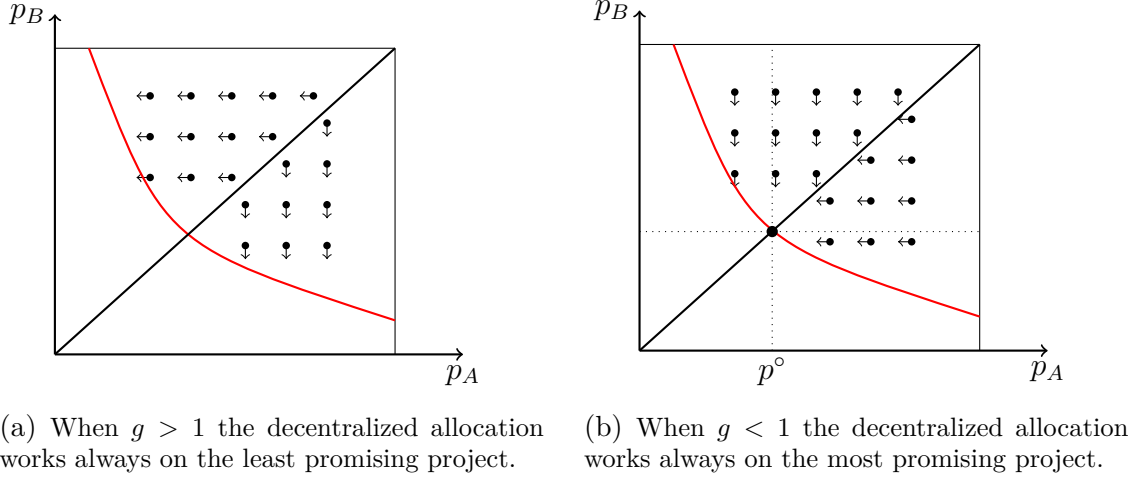


Figure 1.3. Decentralized allocation for $\alpha = 1$ with symmetric projects in the first stage.

When $\alpha = 1$, the decentralized allocation chooses to work always on the project with highest hazard-to-value ratio $\frac{h_i(X)}{V(\{j\}, X)}$, which is constant on X_j . We know, from Proposition 2, that when projects are symmetric and $g > 1$, it is efficient to work on them in sequence, always starting with the least promising one. $g > 1$ implies that the hazard rate $h_i(X)$ decreases faster than the value $V(\{j\}, X)$ so for $X'_j = X_j$ and $X'_i > X_i$

$$h_i(X)V(\{i\}, X) \geq h_j(X)V(\{j\}, X) \quad \Rightarrow \quad h_i(X')V(\{i\}, X') \geq h_j(X')V(\{j\}, X')$$

and this implies that the decentralized allocation for $\alpha = 1$ also works on the projects in sequence starting from the least promising one.

Proposition 4 indicates that, for symmetric projects, the intuition that was developed at the beginning of this section is flawed: there is no trade-off between the inefficiencies caused by imperfect appropriability and the ones caused by the endogenous timing of development. The reason is that when the easy project is completed, the remaining

project is difficult for which there is a smaller chance of success. Thus the value of completing an easy project is smaller than the value of completing a difficult one, what compensates for the differential in the rate of success.

For asymmetric projects however, the value of completing a more difficult project might not perfectly compensate for the lower probability of success. The next proposition shows that, even for asymmetric projects, whenever it is efficient to work more on the more promising project, the decentralized allocation of resources is efficient.

Proposition 5. *If $g_i < 1$ for both projects, the decentralized allocation is efficient for $\alpha = 1$.*

When $g_i < 1$ the value decreases faster than the hazard rate. Proposition 5 shows that when this holds for both projects, it is efficient to work always on the project with highest hazard-to-value ratio. This is exactly what the decentralized strategy does.

The decentralized allocation does not consider the information generated by the allocation of resources. Decentralization will thus bias the allocation of resources towards projects where there is less learning. The next proposition considers the case where project A has a known rate of success. For project B the rate that is unknown.

Proposition 6. *If project A has a known success rate, the decentralized allocation for $\alpha = 1$ is inefficient if and only if $\lambda_B^H < \lambda_A$ and p_B is large enough.*

PROOF. The decentralized allocation allocates resources to the project with highest hazard-to-value ratio, where $v_i(X_i) = V(K \setminus \{i\}, X)$. The observation from below says that it is efficient to work in sequence starting from project B .

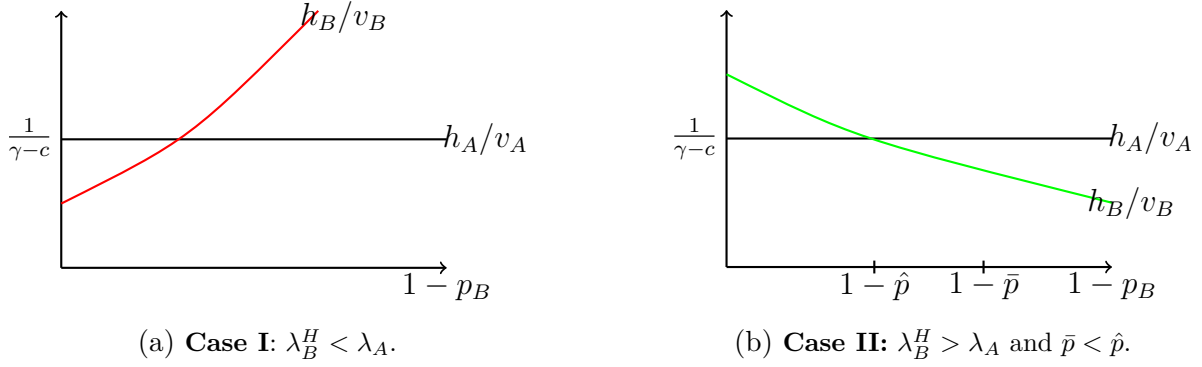


Figure 1.4. The decentralized allocation is inefficient if h_B/v_B is lower than h_A/v_A for a belief for which it is efficient to work.

Lets call \hat{p} the posterior at which is efficient to stop working (interior if the proble is not trivial). h_A/v_A is constant and equal to $\lambda_A/(\gamma - c/\lambda_A)$.

If h_B/v_B is below h_A/v_A when the belief is close to 1, then there is an initial belief p such that it is optimal to work on B but the decentralized allocation works instead on A . This is illustrated in Figure 1.4a. h_B/v_B when p is close to 1 is $\lambda^H/(1 - c/\lambda^H)$, so the condition is

$$\frac{\lambda^H}{1 - c/\lambda^H} \leq \frac{\lambda_A}{\gamma - c/\lambda_A} \quad \Leftrightarrow \quad \lambda^H \leq \lambda_A$$

If h_B/v_B is higher than h_A/v_A at $p = 1$, the only way there could be an inefficiency is if h_B/v_B is decreasing and the stopping belief \bar{p} is higher than the belief at which h_B/v_B and h_A/v_A intersect, \hat{p} . This situation is illustrated in Figure 1.4b. This, however, is never the case:

By Lemma 1, at the efficient stopping point X^* , $V(\{A\}, X^*) + V(\{B\}, X^*) \leq \gamma$. But $V(\{B\}, X^*) = \gamma - \frac{c}{\lambda_A}$, so $V(\{A\}, X^*) < c/\lambda_A$. Moreover at the stopping point, $h_B \cdot V(\{B\}, X^*) = c$, so

$$h_B(X^*) \cdot V(\{B\}, X^*) = c \geq \underbrace{h_A(X^*)}_{\lambda_A} \cdot V(\{A\}, X^*)$$

Rearranging we get

$$\frac{h_B(X)}{V(K \setminus \{B\}, X)} > \frac{h_B(X^*)}{V(K \setminus \{B\}, X^*)} \geq \frac{h_A(X^*)}{V(K \setminus \{A\}, X^*)} \quad \forall X : X_B < X_B^*$$

Where the first inequality holds since $g_B < 1$. □

The intuition that if a project is thought to be easier this would attract more attention to it is partially flawed. As the previous result shows, inefficiencies will show up if one of the projects has a higher rate of success than the other in every state, but they also require that the efficient to start project is thought to be as relatively easy.

1.5. Extensions

In this section, I explore potential departures from the model introduced in Section 1.2. In particular, I extend results relaxing independence, considering more than two projects.

1.5.1. Relaxing Independence

So far we assumed that the projects are independent. Proposition 1 can be generalized more generally to affiliated projects. Let the vector of completion times τ be drawn from a distribution with density function f .

Definition 5. *The projects are affiliated if for every $\tau, \hat{\tau} \in \mathbb{R}_+^2$*

$$f(\tau \wedge \hat{\tau}) \cdot f(\tau \vee \hat{\tau}) > f(\tau) \cdot f(\hat{\tau})$$

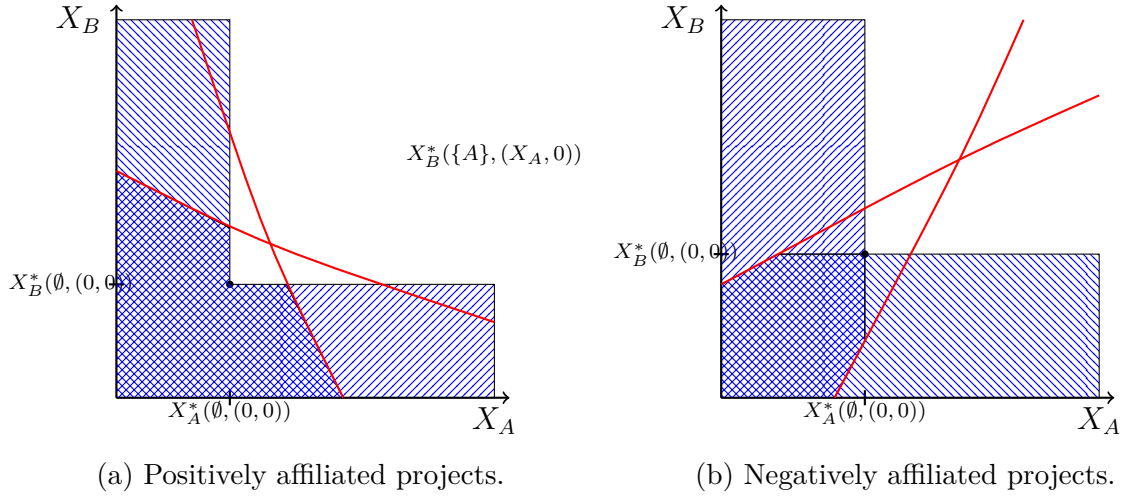


Figure 1.5. Dependent completion times. The agent completes project A (project B) when the vector of completion times falls in the area with north east (west) pattern.

Proposition 7. *If the projects are complements and affiliated, then $V = \hat{V}$.*

The result of the proposition hinges on the success in one project incentivizing the agent to work more on the remaining one. For projects that are affiliated, a success in one of them is good news about the distribution of the completion times of the other, what makes it more attractive to continue working on it.

If the completion times are not affiliated, like in the Figure 1.5b, it can be that a very early success in one of the projects leads to optimally lowering the amount that the DM wants to work on the remaining project. Thus, pledging could be costly and the two problems are not equivalent.

1.5.2. More than two projects

The DM can work on a finite set of projects $K := \{1, 2, \dots, k\}$. As before, the DM allocates a unit of resource across the projects that were not completed so far $S(t)^c$. Again, we define complementarity by the supermodularity of the value of innovations.

Definition 6. *Projects in the set K are complements if q is supermodular*

$$q(A \cup B) + q(A \cap B) \geq q(A) + q(B) \quad \forall A, B \subseteq K$$

If the hazard rate is weakly increasing for all projects, the allocation problem and the order-independent problem.

Proposition 8. *If the projects are independent and h_i is weakly increasing for all $i \in K$ then $V = \hat{V}$ for all q .*

The reason is that when the hazard rate is increasing, if a project is worth allocating any resource, then it must be worth completing. Thus, for the optimal strategy, the set of completed projects is the same for all realizations of τ .

For more than two projects, however, we cannot claim that the problems are equivalent even with complementary, affiliated projects. The reason is that complementarity of the projects does not guarantee monotonicity of the solution to the problems that was used to prove the equivalence with two projects.

Claim 1. *For $k > 2$, independent and complementary projects is not sufficient for $\hat{V} = V$.*

We show this claim by means of the following counterexample:

Example 1. Let $K = \{A, B, C\}$. Suppose $q(\{A, B\}) = \gamma < q(\{A, B, C\}) = 1$. $q(S) = 0$ for any subset. And suppose C is either feasible or infeasible, and that you can learn instantly about it. $\lambda_L^C = 0, \lambda_H^C = \infty$. The optimal strategy is to learn about C , and then doing the optimal thing for A and B (that might be different depending on whether C is completed or not).

In the case where

$$c < \frac{\lambda_L \lambda_H}{\lambda_L + \lambda_H} < \frac{c}{\gamma}$$

then by results when C is completed it is optimal to work simultaneously, $Y_i(\{C\}, 0) > 0$ for $i = A, B$. But when C fails, it is optimal to work in sequence, so $Y_i(\emptyset, 0) = 0$ for $i \in \{A, B\}$.

1.6. Conclusion

Innovation is one of the main determinants of long-term economic growth. Thus, understanding the trade-offs in different approaches to innovation as well as the inefficiencies associated with economic environments is of central importance.

This paper makes substantive contributions to the understanding of the problem of development of complementary innovations:

First, the problem of efficient development of complementary innovations features different challenges than that of substitute innovations: failures in one development affect the expected returns from complementary innovations. With complementary innovations, however, successes make it more attractive to continue working on the remaining developments, what simplifies the problem substantially.

Second, allocating resources to innovation projects in an efficient way involves developing complementary innovations with a specific timing: sometimes it is efficient to develop in sequence and sometimes it is efficient to develop multiple innovations simultaneously. Sequential development is more likely to be efficient for high cost of development and higher uncertainty about the completion rates of the projects.

Third, an important part of innovation is carried out by the private sector. The timing of innovation is partly determined by the investment decisions of agents whose objectives are typically misaligned from the social welfare. Complementary innovations, generate investment dynamics that are different than for substitutes. In particular, the allocation of resources is not simply biased towards the easy and fast component to the detriment of the hard but cost-effective ones.

CHAPTER 2

Market-Based Mechanisms**(Joint work with Quitzé Valenzuela-Stookey)****2.1. Introduction**

One of the fundamental insights of information economics, going back at least to Hayek (1945), is that market outcomes can aggregate dispersed information. As a result, policy makers facing uncertainty often use, or are encouraged to use, market outcomes, such as prices in financial markets, to inform their decisions. However, the use of market outcomes to inform policy making is complicated by the fact that the policy maker's own action may have a significant impact on the market in question. Market participants anticipate the policy maker's action, and this influences the market outcome. This creates a feedback loop between actions and market outcomes, which constrains the policy maker's ability to learn from the market. Market-based policies may also be vulnerable to manipulation by market participants. Moreover, the dependence of market outcomes on expectations of the policy maker's endogenously determined action can lead to multiplicity of equilibria, and potentially non-fundamental market volatility (Woodford, 1994).

This paper studies the general problem of using market outcomes to inform decision making in settings with feedback effects. To fix ideas, consider the following example. An international lender such as the IMF or World Bank must decide on the size of an emergency loan to be extended to a country experiencing a crisis. The severity of the

crisis is determined by a number of factors, such as anticipated changes in the balance of payments, prospects of domestic manufacturers, and the government’s capacity for reform. The lender is unaware of the precise severity of the crisis; this is the unknown state. Traders in the market for bonds issued by the country have private information regarding the state.¹ Suppose that, anticipating that bond prices will be an informative signal of the state, the lender publicly commits to a decision rule which specifies the loan amount as a function of bond prices.² Thus, by observing bond prices, traders can infer the size of the loan which will be extended to the country.

The difficulty, from the lender’s perspective, is that the information revealed by prices depends on the joint distribution of prices and states, which is an equilibrium object. In particular, traders’ demand will depend on the anticipated loan amount, as well as their private information. As a result, the lender’s choice of decision rule will itself shape the information revealed by bond prices. This effect must be accounted for by the lender when choosing the decision rule mapping bond prices to loan amounts. The results of this paper allow us to fully characterize the set of joint distributions of states, bond prices, and loan amounts that the lender can implement by using such a decision rule. More importantly, we characterize what joint distributions can be implemented as the unique equilibrium outcome while also preventing manipulation by bond traders. We return to this example in detail in Section 2.1.1, where we use our results to characterize optimal

¹While it is not necessary at this point to specify explicitly how the bond market functions, the reader can imagine, for concreteness, that traders submit demand schedules to a market maker, who chooses the bond price to clear the market.

²Bond prices are one of a number of factors in the Debt Sustainability Analysis used by the IMF and World Bank when evaluating potential borrowers. More generally, there is growing interest in explicitly conditioning relief on measurable outcomes, for example via state-contingent debt instruments (Cohen et al., 2020).

policy. We show how features of the environment, in particular the size of the spillovers from the country to the global economy, determine whether the lender’s first-best can be achieved, and shape how responsive the loan amount is to the bond price under the lender’s optimal policy.

The use of market outcomes to inform policy decisions occurs in a number of other settings. Many cap-and-trade policies for regulating carbon emissions include mechanisms for responding to the price and/or excess supply of emissions credits (Flachsland et al., 2020). Bank regulators may use market prices of bank securities to inform an intervention decision (Greenspan, 2001). Central banks condition monetary policy on macroeconomic indicators such as the unemployment rate or the rate of inflation. A fall in a company’s stock price can prompt shareholder action to replace top management (Warner et al., 1988). Moreover, the slow and disjointed response to the current COVID-19 crisis has prompted additional interest in “rules-based” policy, in which policy is conditioned on measurable outcomes in a pre-determined way. For example, it is argued that state-contingent debt instruments, in which payments are conditioned on variables such as GDP or commodities prices, should be used to reduce the need for protracted and costly sovereign debt restructurings (Cohen et al., 2020).

The model

Before detailing the primary contributions of the paper, we outline the model and discuss the challenges faced by the principal. The example of the international lender illustrates the components of the general model. A principal (the lender) will need to choose an action (the loan amount), but faces uncertainty regarding a payoff-relevant state of nature (crisis severity). There is a market populated agents (the bond traders)

who may have some private information about the state. The behavior of agents in the market depends on their beliefs about the state, as well as the action that they expect the principal to take. The joint behavior of agents determines an aggregate market outcome (the bond price).

We focus primarily on a principal who has commitment power. To exploit information revealed by the market, the principal publicly commits in advance to a decision rule mapping the market outcome to their action. This decision rule is the principal's design instrument.

It remains only to describe how the market functions, i.e. how the market outcome is determined. For simplicity, we refer to the market outcome as the price, although the theory allows for more general interpretations. The price is the equilibrium outcome of some underlying market game played by the agents. Our analysis applies to a broad class of markets, include asset markets and settings in which experts forecast a macroeconomic variable such as the unemployment rate. Of particular interest is the noisy rational expectations equilibrium model, of the type introduced by Grossman and Stiglitz (1980) and Hellwig (1980), which is the workhorse model of asymmetric information in asset markets, and which we study in detail in Section 2.4.2. In order to analyze these markets in a unified framework, we derive a reduced-form representation of equilibrium in the market game. This approach is detailed below.

Analysis

Starting from some primitive market game, we break or analysis of the principal's design problem into three steps. The first step is to redefine the problem in outcome space. While the principal's design instrument is the decision rule, we assume that their

payoffs depend on the decision rule only through the equilibrium outcomes that it induces: the joint distribution of the action, state, and price.³ This distribution can be described by the equilibrium mapping from the state to the principal’s action and the price. Therefore, rather than directly studying the choice of decision rules (the mapping from the price to the principal’s action), we instead focus on the mappings from the state to the principal’s action and the price that are induced in equilibrium. We refer to these equilibrium objects as the action and price functions respectively. We then ask which action and price functions are *implementable*, i.e. can be induced as equilibrium outcomes by some decision rule. Given implementable price and action functions, the implementing decision rule can be immediately identified. There are a number of benefits of working directly in the space of price and action functions, rather than the space of decision rules.

The second step is to derive a tractable way to model many different types of markets within an unified framework. We do this by showing, for a wide range of markets, how to derive a reduced-form representation of market equilibrium. This is done in a way that facilitates a type of state-by-state analysis. The complication that must be dealt with in this step is that the decision rule used by the principal can affect agents’ behavior in two ways. The first is a direct forward-guidance effect: the decision rule shapes what action agents anticipate the principal taking, conditional on the price. However in markets where agents also use the price to learn about the state (for example, as in Grossman and Stiglitz (1980)) the decision rule will also have an informational effect on agents: the decision rule will shape the equilibrium price function, and thus determine what inferences agents draw

³The model accommodates a principal who cares directly about the price (for example, if the “price” is the rate of inflation managed by a central bank), as well as a principal who’s payoff depends only on the joint distribution of the state and their own action, in which case the price is purely an instrument for learning about the state.

from the price. This informational effect is somewhat subtle, as the inferences agents draw from a given price realization can in general depend on global features of the equilibrium price function, and thus global properties of the principal's decision rule.

Whereas the typical approach in the literature is to explicitly solve for equilibria in specific settings, our approach allows us to analyze policy even when an explicit solution is not available. This additional flexibility is extremely valuable in applications, such as that of Section 2.6.2.

The final step in the analysis is to use the reduced-form representation of the market to characterize the set of implementable price and action functions. We focus in particular on characterizing what is implementable subject to the constraints imposed by concerns about equilibrium multiplicity and market manipulation.

Robustness concerns

In many cases, implementability alone is too weak a condition for practical policy-making purposes. First, the principal may be concerned about market manipulation. Manipulation will be particularly easy if the decision rule is discontinuous; agents will be able to induce a discrete change in the principal's action by triggering arbitrarily small perturbations to the price. This suggests that in order to prevent manipulation, it is necessary to use a continuous decision rule. In fact, requiring continuity everywhere is stronger than needed. Discontinuous decision rules can be used, provided the discontinuities occur away from any prices that could arise in equilibrium. We call such decision

rules *essentially continuous*.⁴ We show that essential continuity characterizes robustness to small manipulations.⁵

Second, the principal may be concerned about indeterminacy of outcomes when the proposed decision rule admits multiple equilibria. Non-fundamental volatility is well documented in environments such as asset markets where expectations play an important role in determining outcomes. Such volatility can arise when agents in the market coordinate on one or another self-fulfilling belief. There is therefore great interest in designing policies for which a unique equilibrium outcome exists, as discussed in Woodford (1994). This is especially true in problems, such as managing inflation, in which stability is a paramount concern. Moreover, conditioning policy decisions on prices often exacerbates equilibrium multiplicity (Bernanke and Woodford, 1997).

We say that a decision rule is *robust to multiplicity* if it induces a unique outcome in all states. In other words, there is a unique equilibrium map from states to the price and the principal's action. A decision rule is *weakly robust to multiplicity* if it induces a unique outcome in all but a zero measure set of states. Our primary focus will be on unique implementation; in other words, implementation via decision rules that are robust to multiplicity. Moreover, we show in Section 2.5.2 that it is without loss of optimality to

⁴For the purposes of the discussion in the introduction, the reader can think of the decision rules as continuous. This captures the primary content of the assumption. The reasons why essential continuity, as opposed to continuity everywhere, is the correct condition is discussed further in Section 2.2.5. How continuity, as opposed to essential continuity, affects the results is explained by Lemma 2. Continuity arises as a necessary condition for preventing manipulation in related settings. For example, Duffie and Dworczak (2020) show that financial benchmarks which prevent manipulation must be continuous functions of prices and trades.

⁵In most applications we model agents in the market as infinitesimal, in which case each agent has no price impact. However we understand this as an idealization of a model in which agents are small but not infinitesimal, and have a very small price impact. As discussed in Section 2.2.5, essential continuity can be interpreted as the limit of the conditions needed to prevent manipulation as agents' individual price impact goes to zero.

restrict attention to decision rules that are robust to multiplicity if the principal takes a worst-case or adversarial view of multiple equilibria. Thus this restriction may be useful even when non-fundamental volatility is not a primary concern.

We say that an action and price function pair is *continuously uniquely implementable (CUI)* if it is implementable by an essentially continuous decision rule that is robust to multiplicity, and *continuously weakly uniquely implementable (CWUI)* if the implementing decision rule is essentially continuous and weakly robust to multiplicity.⁶ We are primarily interested in CUI outcomes; the weaker notion is used to clarify key features of the results. Our main results concern the characterization of the set of CUI and CWUI price and action functions. Before describing these results, we outline the general contributions of the paper.

Contribution

The current paper makes four major contributions relative to the existing literature. First, we provide a general framework for studying market-based interventions in environments with feedback effects. By focusing on implementable price and action functions, rather than directly on the decision rule, we are able to shed new light on the general structure of the problem. Moreover, we show in a broad class of settings, including the canonical noisy REE model (Section 2.4.2), how the effect of the decision rule on the beliefs of market participants about the state can be separated from its role in shaping their expectations of the principal's action. This separation significantly reduces the complexity of analyzing market-based interventions.

⁶Again, we do not in fact require continuity everywhere. The precise continuity requirement is discussed in Section 2.2.5.

Second, we fully characterize the set of price and action functions which the principal can induce in equilibrium. More importantly, given concerns about manipulation and equilibrium multiplicity, we characterize what price and action functions can be induced as the unique equilibrium by an essentially continuous decision rule. In other words, we characterize the feasible set of policy outcomes under robustness to multiplicity and manipulation. We view this as our primary contribution. This characterization admits a tractable representation, which significantly simplifies the analysis of optimal policy in applications (see Section 2.1.1 for a simple illustration). Existing analyses optimize over the space of decision rules. Generally, this approach requires one to impose restrictions on the environment and/or the admissible decision rules which make it possible to solve for equilibrium in closed form. In contrast, optimizing in the space of implementable action and price functions increases tractability, and we are able to identify qualitative features of optimal policy even without a closed form solution. Moreover, we are able to highlight the cost imposed by restrictions on the decision rule which are sometimes used in the literature for tractability purposes (Section 2.6.2).

On a conceptual level, this characterization also reveals a surprising interaction between unique implementation and the constraints imposed by concerns about market manipulation. Singly, neither set of constraints imposes a substantive restriction on the implementable set. However jointly they have important implications for what the principal can achieve (discussed below). To our knowledge, we are the first to consider these constraints jointly.

Third, we show that essential continuity and robustness to multiplicity imply a natural notion of robustness to model misspecification (Section 2.5.1). This means the principal's

payoff is not highly sensitive to their limited understanding of market fundamentals. Finally, the results also allow us to analyze optimal policy when the requirement of unique implementation is relaxed. In particular, we use our characterization of the implementable set to show that if the principal takes a worst-case approach to equilibrium multiplicity then the restriction to unique implementation is generally without loss of optimality. We also discuss optimal policy under alternative criteria for evaluating multiple equilibria.

Main results

We first show how to map a wide range of problems into our general framework (Section 2.2.2). Within this framework, the key results concern the characterization of CUI and CWUI action functions. When the state is one-dimensional, we show, under general conditions discussed in Section 2.3, that all essentially continuous decision rules which are weakly robust to multiplicity induce a price that is monotone in the state (Theorem 1).⁷ In other words, monotonicity of the price is a necessary condition for CUI and CWUI. A version of this result extends to multi-dimensional state spaces, discussed in Section 2.4. To summarize the basic reasoning behind this result, assume that for any fixed principal action, the equilibrium price is strictly increasing in the state. Then for a price function to be implementable and non-monotone it must be discontinuous; it cannot be that different actions are taken in different states yet induce the same price, since the decision rule is measurable with respect to the price. The assumption that the decision rule is essentially continuous does not mean that the price and action functions

⁷Readers familiar with the mechanism design literature may suppose that monotonicity of the price in this setting is related to the monotonicity of feasible allocations that arises in many mechanism design problems, and which is generally the consequence of single-crossing payoffs. This is not the case; our environment does not share the important features of classical mechanism design problems. Price monotonicity reflects an entirely different set of factors.

need be. However, we show that if there are discontinuities and non-monotonicities in the price functions then there are multiple equilibria. This is due to the fact that while the price function may be discontinuous, the decision rule must be essentially continuous. The key step is to show that continuously “bridging the gaps” where the price function is discontinuous creates multiplicity.

We then fully characterize the set of CUI and CWUI price and action functions. In general, an action function will be CUI if and only if it is continuous and induces a monotone price function (Theorem 2, which contains additional technical conditions which apply to extreme states when the state space is closed). The key feature of this characterization is monotonicity of the price function: we show that under general conditions an action function is CWUI if and only if it induces a monotone price function (Section 2.3.2). Additional nuances arise when the state space is multi-dimensional, which we cover in section Section 2.4.

The characterizations of CUI and CWUI price functions significantly simplify the problem of finding optimal decision rules. The set of action functions which induce increasing price functions is much smaller than the set of all action functions. Moreover, when the first-best is not CUI (or CWUI), the optimal policy can often be solved for via a simple flattening procedure, involving only a few scalar parameters. This is illustrated in Section 2.6.1.

This paper also addresses another important aspect of robustness of the decision-making protocol. The principal in general has limited information about the fundamentals of the economy. In particular, the principal may not know precisely the relationship

between states, prices, and the expectations of market participants regarding the principal's action. It is therefore desirable for the principal to use a decision rule that is robust to such uncertainty; small perturbations to the fundamentals should not lead to drastic changes in the equilibrium joint distribution of states, prices, and actions. We say that the decision rule is *robust to structural uncertainty* when it induces a map from fundamentals to outcomes that is suitably continuous. We show that any essentially continuous decision rule that is robust to multiplicity is robust to structural uncertainty. In other words, given CUI, the principal gets robustness to structural uncertainty for free (Theorem 3).

When non-fundamental volatility of market outcomes is not a primary concern, the principal may be willing to tolerate equilibrium indeterminacy, provided all equilibria give the principal a sufficiently high payoff. We therefore consider decision making when the requirement of unique implementation is relaxed. Under general conditions, we show that any essentially continuous decision rule which admits multiple equilibria induces at least one equilibrium that could in fact be implemented (weakly) uniquely by an appropriate modification of the decision rule (Proposition 17). This result has a number of important implications. Most notably, suppose the principal takes a strict worst-case view of multiplicity, i.e. evaluates decision rules based only on the worst equilibrium that they could induce. The above result then implies that it is without loss of optimality to restrict attention to CWUI outcomes. We discuss further relaxations of unique implementation in Section 2.5.2.

Applications

In Section 2.6.2, we discuss the distinctive features of settings in which the principal attempts to “move against the market”. For example, central banks often use open market

operations (the principal’s action) to reduce the interest rate (the price) during severe crises, while, absent interventions, interest rates would be increasing in the severity of the crisis (the state). We show that in such settings, it is necessary for the principal to use a non-monotone decision rule (not to be confused with a non-monotone price function) in order to induce a unique equilibrium. This highlights the cost of placing ex-ante restrictions on the decision rule, for example restricting attention to linear decision rules, as is common in the literature for reasons of tractability. In the central bank example, restricting attention to monotone decision rules implies that the bank cannot induce lower interest rates when the crisis is more severe without also being vulnerable to non-fundamental volatility. This restriction comes from a surprising interaction between equilibrium multiplicity, monotonicity of the principal’s decision rule, and bounds on the set of actions available to the principal (e.g. the size of asset purchases/sales). By allowing for more general decision rules we show that the central bank can uniquely implement essentially any decreasing function from the state to the interest rate. This application demonstrates the value of our characterization of the entire feasible set, beyond simply facilitating the search for optimal policies.

Our results are also useful for identifying optimal decision rules. In Section 2.6.1, we apply the results to the problem of a government considering a bailout of a firm or industry. The government uses the firm’s stock price to inform its decision.⁸ We show that the government’s first best policy is CUI if and only if the positive social externalities from bailing out the company are high. The optimal decision rule involves a gradual transition from a large to a small bailout as the firm’s stock price increases. We also characterize

⁸A related application is performance pricing in debt contracts, whereby the interest rate is conditioned on the borrower’s financial ratios, e.g. interest coverage, or credit ratings (Grochulski and Wong, 2018).

the optimal CUI policy when first-best is not feasible, which in this case involves a rapid reduction of the level of support as a function of the stock price, and show which policies are optimal when the uniqueness requirement is relaxed.

Related literature

A detailed description of our contribution is contained in an earlier section of the introduction. Here we outline the relationship to the existing literature. This paper is part of the literature related to the two-way feedback between financial markets and the real economy, beginning with Baumol (1965). For a survey of this literature see Bond et al. (2012). Among other contributions, this literature documented the possibility of multiplicity of equilibria (see, among others Dow and Gorton (1997), Bernanke and Woodford (1997), and Angeletos and Werning (2006)). The current paper brings a design approach to policy making in these settings, formalizing the problem of policy design under commitment in a general setting and providing a full characterization of feasible policy outcomes while accounting for manipulation and equilibrium multiplicity concerns.

This paper is also related to the literature studying market-based intervention in the presence of feedback effects *without* commitment. Bond et al. (2010) study a problem similar to the emergency lending example of Section 2.1.1, but where the principal does not have commitment power. They identify that there cannot be an equilibrium in which the principal's first-best is achieved exactly in the situation in Figure 2.1b, when the induced price function would be non-monotone. In the language of the current paper, this is because the induced price and action functions violate the necessary measurability

condition for implementability; the action must be measurable with respect to the price.⁹ However, we show that if the principal is concerned with equilibrium multiplicity and manipulation, then non-monotonicity of the price is problematic even if there is no violation of measurability, for example as depicted in Figure 2.2a. Identifying that monotonicity of the price is necessary for unique implementation under robustness to manipulation is one of our primary contributions.

The most closely related work in the literature without commitment is Siemroth (2019), which studies a noisy REE market with a principal who learns from the asset price, similar to the setting in Section 2.4.2, and identifies conditions under which a rational expectations equilibrium exists when the principal lacks commitment. In contrast, we fully characterize unique implementability and solve for optimal policies under commitment.¹⁰ Moreover, Siemroth (2019) restricts attention to equilibria in which the price function is continuous (not to be confused with continuity of the decision rule). This is a substantive assumption, as it implies that the equilibrium, when it exists, is unique. In a noisy REE model *without* feedback effects, Pálvölgyi and Venter (2015) and Breon-Drish (2015) show that in general multiple equilibria are possible. Uniqueness holds only within the class of equilibria with continuous price functions. Multiplicity that arises even without feedback effects, for example if the principal does not condition on the price, can be called *fundamental multiplicity*. However, in settings with feedback effects there may also be equilibrium multiplicity caused by the endogeneity of the principal's

⁹Bond et al. (2010) then observe that if the principal has access to a signal with a sufficiently narrow bounded support around the true state, they will be able to differentiate between high and low states which induce the same price, and thus overcome the measurability problem.

¹⁰Other important differences between the current paper and Siemroth (2019) are discussed in Section 2.4.2.

action. Eliminating fundamental multiplicity by imposing continuity of the (endogenous) price function also eliminates multiplicity caused by action endogeneity. Moreover, it does so by imposing restrictions on the principal's policy.¹¹ The present paper's contribution is in characterizing the set of implementable outcomes; we do not restrict this set ex-ante by imposing continuity of the price function. Instead, we characterize the restrictions on the set of implementable outcomes imposed by robustness to multiplicity and manipulation.

Other papers have noted that policy based on market outcomes may be vulnerable to manipulation. Goldstein and Guembel (2008) study manipulation by strategic traders when firms use share prices in secondary financial markets to guide investment decisions. In Lee (2019) a regulator uses stock-price movements of affected firms to determine whether or not to move forward with new regulation. In Lee (2019), the discontinuous nature of the policy considered opens the door to manipulation. Motivated by these concerns, we characterize robustness to manipulation in the limit as agents in the market becomes small, and consider policies that are robust to manipulation in this sense.

This paper relates most directly to the literature on policy making under commitment in the presence of feedback effects. Important contributions include Bernanke and Woodford (1997), Ozdenoren and Yuan (2008), Bond and Goldstein (2015), Glasserman and Nouri (2016), Boleslavsky et al. (2017), and Hauk et al. (2020). Bernanke and Woodford (1997) show how the use of inflation forecasts to inform monetary policy can reduce the informativeness of forecasts. In the language of our paper, this occurs when the induced

¹¹Without commitment, one could argue that if the principal's best response is suitably continuous, it is natural to focus on equilibria with continuous price functions. This is not the case with commitment however; the principal may wish to commit to a policy that induces a discontinuous price function, even if the first-best action function is continuous.

market-outcome function (in this case the inflation forecast) violates the necessary monotonicity condition. Bernanke and Woodford (1997) restrict attention to linear decision rules, and show that equilibrium multiplicity can arise. Our analysis shows that non-monotone decision rules may in fact be *necessary* to prevent multiplicity (Section 2.6.2). Bond and Goldstein (2015) focus on the how market-based interventions affect the efficiency of information aggregation by prices. In contrast to the current paper, traders in Bond and Goldstein (2015) care about the state only insofar as it allows them to predict the government’s action. As a result, information aggregation is highly dependent on the decision rule.

Glasserman and Nouri (2016) show how equilibrium multiplicity issues that arise in a static setting may not be present in a dynamic trading model. In a static problem nearly identical to that depicted in Figure 2.1a, they show that equilibrium multiplicity will arise if the principal uses a discontinuous threshold rule. Restricting attention to such rules, they show that in a dynamic version of the model there may be a unique equilibrium. We observe that in this type of problem, in which the price function is monotone, the multiplicity issue can also be resolved by allowing for gradual adjustment of the principal’s action. Our main focus however is on identifying what conditions are *necessary* for unique implementation.

Hauk et al. (2020) develop variational techniques for identifying optimal decision rules in settings with feedback effects. These techniques complement our results, which simplify the problem of identifying optimal policies by characterizing the feasible set in the space of action and price functions, rather than the space of decision rules.

In general, our primary contribution relative to this literature is the complete characterization of implementable outcomes, taking into account the practical concerns of equilibrium multiplicity and manipulation. Moreover, by providing a tractable framework for studying flexible market-based policy design in a general setting, we avoid the artificial restrictions imposed by some of the simplifying assumptions used in the literature.

The remainder of the paper is organized as follows. In Section 2.1.1 we illustrate the key results in the context of the emergency lending example introduced above. Section 2.2 introduces the model, and discusses the various robustness notions considered. Section 2.3 presents the main characterization results when the state space is one-dimensional, while Section 2.4 extends the model to multi-dimensional state spaces, and discusses the special case of noisy REE in detail. Section 2.5.1 covers robustness to structural uncertainty, and Section 2.5.2 discusses optimal policy when the unique implementation restriction is relaxed. Section 2.6 explores the applications to monetary policy and bailouts.

2.1.1. A brief illustration: emergency lending

Before detailing the model, we illustrate the key results in the context of the problem, discussed at the start of the introduction, of an international lender such as the IMF or World Bank deciding on the size of an emergency loan to extend to a country experiencing a crisis. The lender is unaware of the precise severity of the crisis, which is represented by an unknown state $\theta \in [\underline{\theta}, \bar{\theta}]$; lower states represent greater severity. Dispersed information regarding the state may be at least partially reflected in the price of government bonds. For simplicity, imagine that all traders in the bond market know the true state (this

assumption is purely for illustrative purposes; it does not affect the results discussed here and is not required in the general model).

Let $\pi(a, \theta)$ be the value of government bonds, i.e. their ex-post payout, if the lender extends a loan of size $a \in [0, \bar{a}]$ and the state is θ . For any loan amount a , bond values are increasing in θ . A large emergency loan leads to higher bond prices when the crisis is severe, as it reduces the probability of default in the short term. However bondholders may also worry that the increase in the country's debt burden could have adverse long-term affects. For example, the increase in the debt burden may lead to debt overhang and push the country down the back side of the debt Laffer curve, as investors worry that long-term growth will be negatively affected by the higher taxes needed to service the increased debt burden (Cordella et al., 2010). If the current crisis is mild, this effect may dominate, in which case bond prices will react negatively to the lender's intervention.¹² These considerations are captured by the following two assumptions on bond values:

- (1) There exists $\theta^* \in [\underline{\theta}, \bar{\theta}]$ such that $\pi(\cdot, \theta)$ is increasing for $\theta \leq \theta^*$ and decreasing for $\theta > \theta^*$.
- (2) $\pi_2(a, \theta)$ is decreasing in a .

The lender would like to extend emergency relief only when the crisis is severe.¹³ For simplicity, assume there exists a state θ^\bullet such that the lender's payoff is increasing in a when $\theta \leq \theta^\bullet$, and decreasing in a when $\theta > \theta^\bullet$. As a result, the lender would ideally like to extend the maximal loan amount \bar{a} if and only if $\theta \leq \theta^\bullet$, and otherwise extend no loan.

We refer to this policy as the first-best action function. A higher θ^\bullet corresponds to a more

¹²Indeed, Cordella et al. (2010) find that the strongest empirical evidence of a negative relationship between debt and growth is for countries with relatively good policies and institutions.

¹³It could be that the lender does not wish to make a loan if the crisis is too severe, and the loan is unlikely to be repaid. Preferences of this sort are covered in Section 2.6.1.

interventionist policy on the part of the lender. The lender is likely to be interventionist if the country is very poor, in which case the short-run welfare losses from government austerity are large, or if the country is central to the global economy, because in this case a recession there will have large spillover effects on other countries. Figure 2.1a illustrates an interventionist first-best action function in which $\theta^\bullet > \theta^*$. The solid lines denote the bond values as a function of the state under the two extreme actions 0 and \bar{a} . The dashed blue line is the price function P^* induced by the first-best action function. Note that for each price p there is at most a single state θ such that $P^*(\theta) = p$. It is therefore possible to choose a decision rule mapping prices to actions that implements the first-best action function. In fact, Corollary 4 implies that this first best in Figure 2.1a will be continuously weakly uniquely implementable (there will only be multiple equilibrium actions in state θ^\bullet). In this case, the decision rule which uniquely implements the first-best involves a gradual reduction in the level of support as the bond price increases over an intermediate range.

Figure 2.1b illustrates a conservative first-best policy. In this case the lender is unwilling to intervene in some states in which bondholders would like the government to receive a large emergency loan. This is likely the most realistic scenario for middle-income countries. In this case the first-best action function cannot be implemented by a market-based decision rule. This is due to the fact that for prices in (p', p'') the price function does not reveal enough information: upon observing such a price the decision maker cannot tell if the state is below θ^\bullet , in which case the action \bar{a} should be taken, or above θ^\bullet , in which case the action should be 0. In other words, the action is not measurable with respect to the induced price.

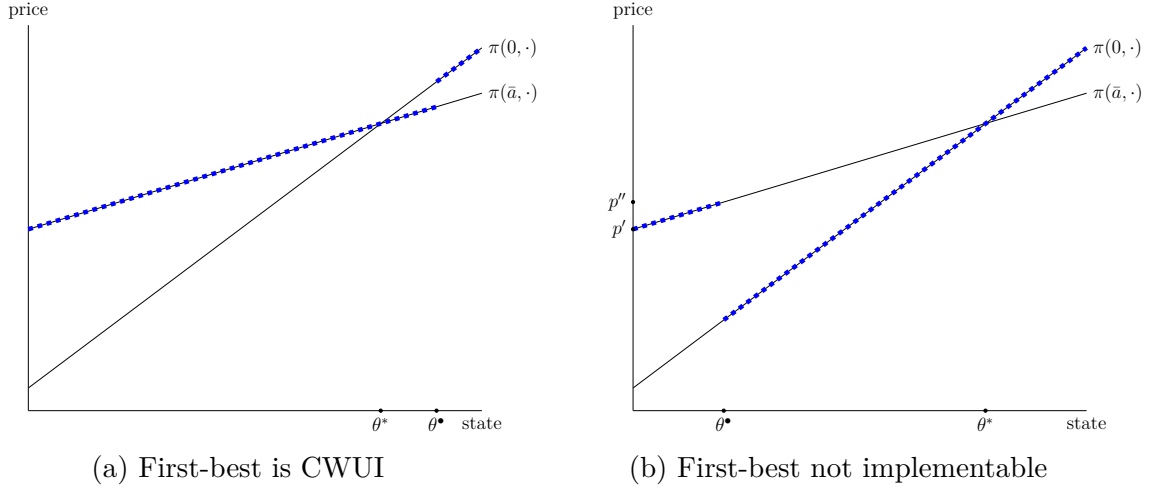


Figure 2.1. First-best

Consider the modification of the conservative first-best action function illustrated in Figure 2.2a, which is a natural way to eliminate the measurability problem discussed above. This requires making an intermediate loan for states in (θ', θ'') , where the lender would prefer not to intervene at all. Given this modification, for any price p there is a unique state θ such that $P^*(\theta) = p$, and so this action function is implementable. In fact, it will be implementable with a continuous decision rule.

Unfortunately, it is not possible to continuously and *uniquely* implement a policy resembling that of Figure 2.2a. In fact, any continuous decision rule M that implements this action function will induce at least one equilibrium in which large loans are made for all states in $(\theta^\bullet, \theta')$. Theorem 1 shows that this is precisely because of the non-monotonicity in the induced price function. Roughly speaking, the intuition is that when M is continuous there cannot be discontinuities in the function from prices to states which specifies, for each price p , the set of states in which p could be an equilibrium price.

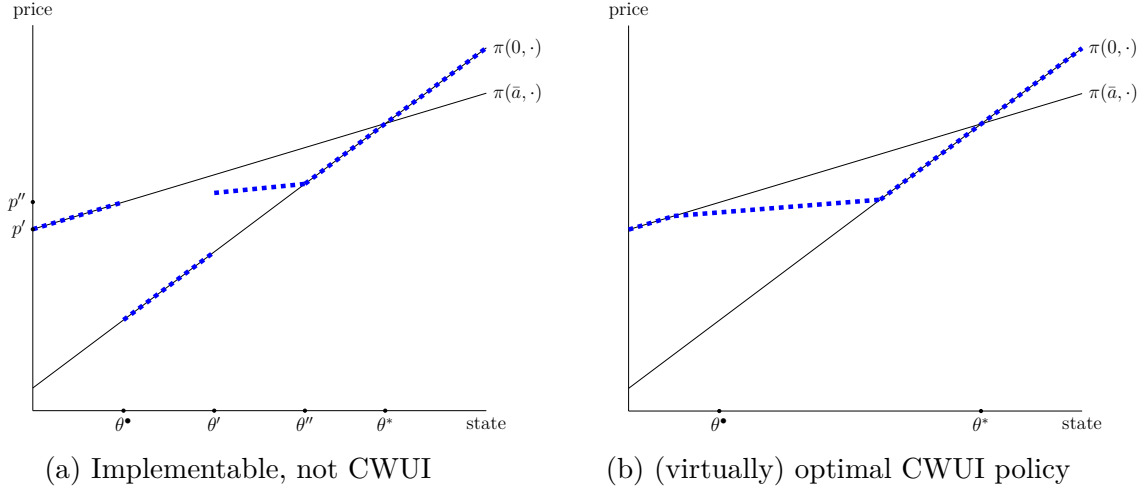


Figure 2.2. Implementable policies

The (virtually) best CWUI action function will have a monotone price, as illustrated in Figure 2.2b; the lender trades off lower than desired lending for some states below θ^\bullet with higher than desired lending for states above θ^\bullet .¹⁴ In fact, the optimal policy will always involve both types of loss (see Section 2.6.1 for a discussion). Lending as a function of θ decreases gradually in order to maintain price monotonicity. In fact, this policy will be CUI (Theorem 2). Moreover, Proposition 16 implies that if decision rule M^* is optimal within the set of essentially continuous and decision rules that are robust to multiplicity, i.e. those that induce an equilibrium of the form in Figure 2.2b, then M^* will be in the set of optimal decision rules even if the uniqueness requirement is relaxed and the principal takes a worst-case approach to multiplicity. In contrast to the decision rule implementing the first-best when the lender is interventionist, the virtually-optimal decision rule in the

¹⁴An action and price function pair is *virtually implementable* if they can be approximated arbitrarily well by an implementable pair. We must consider virtual optimality in this setting, since the price function must be strictly increasing, but can be arbitrarily close to flat.

conservative case will entail a rapid reduction in the level of support as the bond price increases over an intermediate range.

2.2. The model

The baseline model consists of the following primitive objects.

- (1) The state space, denoted by $\Theta \subseteq \mathbb{R}^N$.
- (2) A convex set \mathcal{A} of principal actions, which is a subset of a Banach space.
- (3) A convex set $\mathcal{P} \subseteq \mathbb{R}$ of aggregate outcomes.

For clarity, we will refer to the aggregate outcome as the price, although the model applies to many situations in which the aggregate outcome is not a price, as is discussed below. The state may contain dimensions that are not directly payoff relevant for the principal. For example, in a noisy REE model of an asset market, as in Grossman and Stiglitz (1980), the state will include the supply shock, in addition to the payoff relevant state. More generally, the state will represent the entire profile of agent's private signals, as in Jordan (1982). There are three periods, and the timing of interaction is as follows.

- 0. The principal publicly commits to a decision rule $M : \mathcal{P} \rightarrow \mathcal{A}$ specifying an action for each price.¹⁵
- 1. The price is determined.
- 2. If the price is p , the principal takes the action $M(p)$.

It only remains to describe how the price is determined in period 1. As this is a key feature of the model, it is discussed in some detail in Section 2.2.2, and expanded upon in

¹⁵In Appendix A.9 we briefly discuss the version of the model in which the principal cannot commit.

Section 2.4.2 and Appendix A.8. Before detailing the approach to modeling the market, we introduce the concept of implementation in this setting.

2.2.1. Implementation

We make no assumptions on the preferences of the principal, other than that they do not depend directly on the announced decision rule M . From an ex-ante perspective, the principal cares only about the joint distribution of states, actions, and prices induced in equilibrium by their decision rule. This joint distribution can be summarized by the maps from states to actions and prices that are induced in equilibrium. We refer to $Q : \Theta \rightarrow \mathcal{A}$ and $P : \Theta \rightarrow \mathcal{P}$ as the *action function* and *price function* respectively. The goal is then to characterize what action and price functions are implementable, meaning that they are the equilibrium outcomes given some decision rule $M : \mathcal{P} \rightarrow \mathcal{A}$. The equilibrium concept will of course depend on the market in question. This is the subject of the subsequent section.

2.2.2. Price formation

The analysis of equilibrium price formation is one of the central technical challenges posed by the endogeneity of principal actions. This is especially true when market participants learn about the state from the price, as is the case in the majority of applications. When this occurs, the principal's decision rule will affect the behavior of agents in the market in two ways. The first is the direct forward guidance effect: the decision rule determines what action agents expect the principal to take. However the decision rule will also have an indirect informational effect on agents. To understand this second effect, suppose first

that the principal's action is fixed. Consider, as an example, a rational expectations equilibria in an asset market with asymmetric information, along the lines of Grossman and Stiglitz (1980). The equilibrium is characterized by a price function, and investors draw inferences regarding the state not only from their private information, but also from the realized price. The information conveyed by the price is determined by the coarsest partition of the state space with respect to which the price function is measurable. Thus, the information revealed to investors in any given state depends on *global* properties of the equilibrium price function. This is also the case when the principal's action is price contingent, as agents continue to learn from the price. However the principal's decision rule shapes the price function, and hence the information revealed by the price to investors (as well as to the principal). As a result, the equilibrium price in any given state can depend on the equilibrium prices and principal actions in other states, rather than simply the equilibrium action in that state. This dependence on global properties is what complicates the problem. Fortunately, we are able to identify general conditions under which the forward guidance and informational effects of the principal's decision rule can be concisely summarized.

In order to facilitate a wide range of applications, we take a modular approach to modeling price formation. We show that the equilibrium outcome in a broad class of market models can be summarized as a function of the state and anticipated principal action. This allows us to derive a reduced-form representation of the equilibrium price. We first present a general reduced-form representation. The reduced form summarizes the equilibrium outcomes of the game through which the price is determined. We then discuss the derivation of the reduced-form representation in various price formation games

which are consistent with this reduced-form approach, including rational expectations equilibrium asset market models with and without noise traders. Once the reduced form has been identified, we can use this representation to characterize the implementable set.

2.2.2.1. Price formation: reduced form. We show that a wide range of markets the equilibrium price in a given state is uniquely determined by the equilibrium principal action in that state, independent of the global properties of the equilibrium. This is formalized in the following definition.

Definition 7. *The market **admits a reduced-form representation** if \exists a function $R : \mathcal{A} \times \Theta \rightarrow \mathcal{P}$ such that for any Q, P, M , the pair (Q, P) are equilibrium outcomes given M iff for all $\theta \in \Theta$*

$$\begin{aligned} (1) \quad Q(\theta) &= M(P(\theta)) && \text{(commitment)} \\ (2) \quad P(\theta) &= R(Q(\theta), \theta) && \text{(market clearing)} \end{aligned}$$

We refer to R as the *market-clearing function*. This function is best understood as representation of equilibrium in the market game. However, we can also interpret this representation as stating that if the state is θ and all agents believe that the principal will take action a then the price will be $R(a, \theta)$. The defining feature of environments with feedback is that if the principal has announced decision rule M and the price is p then all agents anticipate that the principal will take action $M(p)$. As a result, agents correctly anticipate the principal's action conditional on the price. The key feature of this reduced form representation is that the function R does not depend on which decision rule M is used. This gives the separation between the forward-guidance and information-aggregation roles of the price. We discuss in the next section how this separation can

be obtained in various settings. Throughout, we maintain the assumption that R is continuous (which assumption can be easily derived from conditions in an underlying price-formation game).

A benefit of being able to summarize the equilibrium price via the function R is that the principal does not need to know the details of the market micro-structure in order to design policy. R represents the equilibrium relationship between the action, state, and price. Since it does not depend on the decision rule M , it can be estimated using data from a market in which the principal's action is not conditioned on the price, or in which some other decision rule was used. Thus a principal contemplating the introduction of a market-based decision rule can use historical aggregate data to estimate the function R and design the decision rule, without being subject to the Lucas critique (Lucas et al., 1976) that a change in the policy regime will change the relationship between the fundamentals (the state and anticipated action) and aggregate outcomes.

2.2.2.2. Price formation: micro-foundations. A formal treatment of various micro-foundations is presented in Appendix A.8. The central example of noisy REE in asset markets is analysed in detail in Section 2.4.2. Here we will simply discuss informally two examples.

Asset market. Consider an environment in which there is fixed supply of a single asset and a continuum of traders.¹⁶ The asset pays a dividend that is a function of the state and the principal's action. Each trader receives a private signal that is partially informative about the state. After observing their signals, traders submit demand schedules to a market

¹⁶In Section 2.4.2 we show how the results can be extended to markets with aggregate supply shocks/noise traders.

maker, who then chooses a price to clear the market. Traders base their demand on *a)* the market price, *b)* the anticipated action of the principal, and *c)* their belief about the state. The latter is a function of both their private signal and information conveyed by the asset price. Since the principal's action is a function of the price, there is no ambiguity about the action given the observed price. A rational expectations equilibrium (REE) in this environment consists of a price function $\tilde{P} : \Theta \rightarrow \mathcal{P}$ such that markets clear in each state θ given

- The anticipated action $M(\tilde{P}(\theta))$,
- The inferences made from the price given the function \tilde{P} .

In Appendix A.8 we show that in such an environment, under some assumptions on information and payoffs, this market admits a reduced-form representation. This is despite the fact that the inferences that investors draw from the price function \tilde{P} depend on global properties of \tilde{P} .

Expert forecasts.

In many situations agents may not observe the aggregate outcome when making the decisions that will, taken together, determine the aggregate outcome. For example, the unemployment rate in a given month is the result of the decisions of firms and workers who act without observing the realized unemployment rate. If, in such a situation, the principal makes a decision that is relevant for agents based on the aggregate outcome then agents will need to predict the action that the principal will take. In many such settings, expert forecasts play an important role in agent decision making.¹⁷

¹⁷In other settings policy may be conditioned explicitly on expert forecasts or ratings. For example, Bernanke and Woodford (1997) discuss conditioning monetary policy on expert forecasts of inflation, and

Suppose an economist receives a signal $\theta \in \Theta$ about the underlying state of the economy ω , and reports publicly their expectation \hat{p} of the unemployment rate p . At the end of the month, the government observes p and chooses $a \in \mathcal{A}$ according to $M(p)$. The action here could be, for example, the amount of money to put into an employment subsidy program. The realized unemployment rate will depend on firm's expectations about a and the underlying state θ . Assume that firms trust the economist's forecast; they take it as an accurate prediction of the unemployment rate. Firms then make their personnel decisions. The realized unemployment rate will be given by $J(M(\hat{p}), \theta)$.

The economist recognizes the effect that their forecast has on firm behavior, and thus on the realized unemployment rate. The economist will take this into account when making their prediction. Thus their expectation of the unemployment rate will be given by

$$\hat{p} = \mathbb{E}[J(M(\hat{p}), \omega) | \theta] \equiv R(M(\hat{p}), \theta).$$

Such a fixed point exists when \mathcal{A} is compact and M continuous. Note that R here is a function of the economist's signal θ , rather than the underlying state ω .

2.2.3. Implementation using reduced form

Having established the existence of a reduced-form representation of market equilibrium, we do not need to specify here the solution concept used to derive this reduced form from the underlying market game. Provided that such a reduced-form representation exists, we simply need to ensure that the subsequent analysis is consistent with this representation.

We can define implementability explicitly in terms of the reduced-form representation.

performance pricing provisions in debt contracts may make interest rates contingent on credit ratings (Asquith et al., 2005).

Definition 8. (Q, P) is *implementable* if there exists $M : \mathcal{P} \rightarrow \mathcal{A}$ such that

$$(1) P(\theta) = R(M \circ P(\theta), \theta) \quad \forall \theta \in \Theta \quad (\text{market clearing})$$

$$(2) Q = M \circ P. \quad (\text{commitment})$$

The market clearing condition requires that the realized price be consistent with the anticipated action given decision rule M . The commitment condition simply says that the principal is in fact using decision rule M .

Implementability can be equivalently defined without making explicit reference to the implementing decision rule M .

Observation 1. (Q, P) is *implementable* iff

$$(1) P(\theta) = R(Q(\theta), \theta) \quad \forall \theta \in \Theta \quad (\text{market clearing})$$

$$(2) Q(\theta) \neq Q(\theta') \quad \Rightarrow \quad P(\theta) \neq P(\theta'). \quad (\text{measurability})$$

Here the measurability condition guarantees that there exists a P measurable function M that induces action function Q . Clearly if this condition is violated there can exist no such M . Given an implementable (Q, P) , the implementing decision rule can be easily identified. Measurability implies that the set $Q(P^{-1}(p))$ is either empty or singleton; this defines M on $P(\Theta)$.¹⁸

For any continuous decision rule M , an equilibrium exists (for any $\theta \in \Theta$, the function $a \mapsto M(R(a, \theta))$ has a fixed point by the Schauder fixed point theorem.) However it is possible to define discontinuous decision rules for which no equilibrium exists. Non-existence of equilibrium is a manifestation of incompleteness in the description of the

¹⁸For the purposes of implementing (Q, P) , it is sufficient specify M on $P(\Theta)$. However, concerns about equilibrium multiplicity constrain M outside of $P(\Theta)$, as discussed in Section 2.3.

model. The true meaning of equilibrium non-existence will depend on the nature of the fundamental game played by agents in the market, i.e. on the micro-foundation for the function R . For example, Bond et al. (2010) show that non-existence in a setting with feedback can be translated to a breakdown in trade: the market-maker abstains from making markets because they would lose money by doing so. We will not explicitly model market outcomes when an equilibrium fails to exist. Our focus is on the set of outcomes that can be implemented in equilibrium.

Observation 1 gives a characterization of the set of implementable (Q, P) . However it is not, on its own, a very useful characterization for two reasons. First, it does not point to any general qualitative features of implementable mechanisms. Second, it ignores practical considerations which may constrain the principal in choosing a decision rule. When such constraints are taken into account a more meaningful characterization of the set of implementable mechanisms can be given.

2.2.4. Uniqueness

Our primary focus is on issues related to multiplicity of equilibria. The approach to multiple equilibria depends on the type of analysis being conducted. From an implementation perspective, the question is how to induce a given (Q, P) as equilibrium outcomes. In the implementation literature, this means that (Q, P) should be the unique equilibrium outcomes induced by some decision rule. The mechanism design perspective, on the other hand, is that the principal can choose from any of the equilibria induced by a given decision rule M . From this perspective, the goal of the principal is simply to induce (Q, P) as *an* equilibrium outcome.

We will consider both perspectives in this paper. To begin, we will take the implementation perspective that outcomes must be induced uniquely. We will then show how the results obtained can be related to a more permissive attitude towards multiplicity. We say that (Q, P) are *uniquely implementable* if they are the unique equilibrium outcomes given some decision rule M . In this case, we say that M is *robust to multiplicity*. Equivalently, a decision rule M is *robust to multiplicity* if $\{p \in \mathcal{P} : p = R(M(p), \theta)\}$ is singleton for all θ .

It will also be helpful to consider a slightly weaker notion of robustness to multiplicity. A decision rule M is *weakly robust to multiplicity* if $\{p \in \mathcal{P} : p = R(M(p), \theta)\}$ is singleton for almost all θ . This definition of robustness makes most sense when the principal maximizes expected utility and has an absolutely continuous prior H . If instead H has atoms then the definition should be modified so that the requirement of a unique price holds almost everywhere under H . There is no difficulty in accommodating this modification into the analysis, although it requires rewording some of the results.

2.2.5. Manipulation

Manipulation of the principal's decision via the market price may be a concern, even in markets in which agents are small. An agent may manipulate the price by buying/selling the asset, releasing false information, or other means.¹⁹ A discontinuous decision rule will be particularly vulnerable to manipulation. Suppose there is a discontinuity near a price which realizes in equilibrium. An agent will be able to induce a significant change in the principal's action by manipulating the price, even if their individual price impact

¹⁹Goldstein and Guembel (2008) discusses manipulation of this sort.

is small. While in most applications considered in this paper agents in the market are modeled as infinitesimal, in which case they have no individual price impact, we view this as an idealization of a situation in which agents have a small but non-zero price impact. In such a model, assuming agents' cost of manipulation is proportional to the induced price change and there exist at least some agents who could benefit from moving the principal's action in any given direction, Lipschitz continuity of the decision rule in the neighborhood of any equilibrium price will be necessary to prevent manipulation. The Lipschitz constant is determined by the costs and benefits of manipulation. In the limit, as agents become small and their cost of having a given price impact increases, we simply require continuity of the decision rule in the neighborhood of any equilibrium price.²⁰

A related concern is that if M is discontinuous then the set of equilibrium outcomes may be overly sensitive to the model fundamentals, in particular to the function R , about which the principal may well have imperfect knowledge. Indeed, Lemma 9 shows that if M has a discontinuity at some price which could occur in equilibrium then the equilibrium outcomes will respond discontinuously to changes in R . Decision rules for which the equilibrium outcomes respond continuously to perturbations of R , which we refer to as *robust to structural uncertainty*, are discussed in Section 2.5.1. These results can also be used to model manipulation which translates into perturbations to R .

The concerns about manipulation and model misspecification discussed above suggest that we should restrict attention to continuous decision rules. However the restriction

²⁰Requiring Lipschitz continuity of the decision rule, rather than simply continuity, would not substantively change the analysis. Continuity may be insufficient to prevent manipulation in some settings; additional restrictions may be required in specific applications, and will imply refinements of the set of admissible decision rules.

to everywhere-continuous decision rules is stronger than is needed to address these concerns. As Theorem 3 shows, it is enough to have continuity in the neighborhood of any equilibrium price to guarantee robustness to structural uncertainty. Similarly, if a discontinuity in M occurs at a price which is far from any which could arise in equilibrium then manipulation via a small price impact will not be possible. We therefore allow for discontinuities in the decision rule, provided they do not occur near equilibrium prices. Formally, for any decision rule M , let $\bar{P}_M = \cup_{\theta \in \Theta} \{p \in \mathcal{P} : R(M(p), \theta) = p\}$ be the set of possible equilibrium prices given M , and let $cl(\bar{P}_M)$ be its closure.

Definition 9. *A function $M : \mathcal{P} \rightarrow \mathcal{A}$ is **essentially continuous** if it is continuous on an open set containing $cl(\bar{P}_M)$.*

In other words, an essentially continuous decision rule can have discontinuities only where there are no nearby equilibrium prices. Let \mathcal{M} be the set of essentially continuous decision rules. Throughout, we will restrict attention to decision rules in \mathcal{M} . We will at times refer to this as a continuity requirement; although it does not imply that M must be everywhere continuous, it has the same intuitive content. Discontinuities in M are only needed when the principal attempts to “move against the market”, as discussed in Section 2.6.2, and then are only needed above the highest equilibrium price and below the lowest equilibrium price (see Lemma 2 and Lemma 11).

One can interpret the restriction to essentially continuous decision rules as a way to ensure consistency between the model with small, but not atomistic, agents, and the model with infinitesimal agents. For tractability we generally want to work in the limiting

model in which agent's are infinitesimal, but we do not wish to artificially disregard any manipulation concerns by doing so.

2.2.6. Unique implementation

We first analyse the problem of unique implementation. Unique implementation is a primary objective in many market-based decision settings in which non-fundamental volatility is a fundamental concern, such as the management of inflation by a central bank.

Definition 10. (Q, P) is *continuously uniquely implementable (CUI)* if it is implementable uniquely by an $M \in \mathcal{M}$.

In other words, (Q, P) is continuously uniquely implementable if there exists $M \in \mathcal{M}$ such that:

- (1) $Q = M \circ P$
- (2) $P(\theta)$ is the unique solution to

$$p = R(M(p), \theta)$$

for all θ .

(3)

$$Q(\theta) \neq Q(\theta') \quad \Rightarrow \quad P(\theta) \neq P(\theta')$$

There are two differences between implementability and CUI; the uniqueness requirement in condition 2 and the continuity requirement that $M \in \mathcal{M}$. Continuity, as discussed above, reflects manipulation concerns. If condition 2 holds for almost all θ , rather than all θ , then we say that (Q, P) is *continuously weakly uniquely implementable (CWUI)*.

There is no substantive difference between the two notions, but it is sometimes easier to state results for the weaker notion.

We will sometimes refer to an action function Q as CUI, by which we mean that there exists a P such that the pair (Q, P) is CUI, in similarly for price functions P .

At times, it will be convenient to discuss approximate, rather than exact, implementation. As is standard, we say that (P, Q) is virtually implementable if it can be approximated arbitrarily well by some implementable (\hat{P}, \hat{Q}) . Say that Q' is an ε -approximation of Q if the set $\{\theta : Q(\theta) \neq Q'(\theta)\}$ has measure less than ε .

Definition 11. (P, Q) is *virtually CUI* if for any $\varepsilon > 0$ there exists an ε -approximation of Q that is CUI.

The characterization of CUI (and virtually CUI) outcomes will be one of the main results of this paper. It turns out that this characterization is also central to understanding optimal decision rules even when the uniqueness constraint is relaxed.

2.3. Characterization: one-dimensional state

We turn now to the main results of the paper. To begin, assume that the state space is an interval in \mathbb{R} , with endpoints $\underline{\theta}$ and $\bar{\theta}$.²¹ We then extend the results to the multi-dimensional case. Some preliminary definitions and results are first needed. The following assumption on price formation will be maintained for most results.

Definition 12. R is *weakly increasing in θ* if $\theta \mapsto R(a, \theta)$ is weakly increasing for all $a \in \mathcal{A}$

²¹It is convenient to state some results for open Θ , and others for closed Θ . Unless otherwise specified, the results apply to both cases.

We say that R is *strictly increasing* in θ if $\theta \mapsto R(a, \theta)$ is strictly increasing for all $a \in \mathcal{A}$. Note that the order used on Θ is irrelevant. All results that assume that R is weakly increasing continue to hold under the weaker assumption that R is comonotone in θ ; $R(a, \theta'') \geq R(a, \theta')$ implies $R(a', \theta'') \geq R(a', \theta')$ for all a', a'' . Similarly results that assume that R is strictly increasing continue to hold as long as there exists some order on Θ such that $\theta \mapsto R(a, \theta)$ is strictly increasing for all a . Both strictly and weakly increasing R can be justified by natural assumptions on primitives in many micro-foundations, and weakly increasing R is satisfied in all applications that we have come across.

For any decision rule M , define $\theta_M(p) = \{\theta \in \Theta : R(M(p), \theta) = p\}$. The following useful properties of θ_M are proven in Appendix A.6.1. If R is weakly increasing in θ , $\theta_M(p)$ is a convex and compact valued correspondence. Convexity follows immediately from the fact that R is weakly increasing in θ . Continuity of R in θ implies that $\theta_M(p)$ is compact valued. Moreover, if M is continuous at p' then $p \mapsto \theta_M(p)$ is upper hemicontinuous at p' . When R is strictly increasing in θ , $\theta_M(p)$ will be a function from \mathcal{P} to $\Theta \cup \emptyset$, and will be continuous wherever M is continuous.

The defining feature of CUI outcomes is a monotone price. This condition is necessary under weakly increasing R , and essentially sufficient under additional mild conditions, discussed in the next section.

Theorem 1. *Assume R is weakly increasing in θ . If $M \in \mathcal{M}$ is weakly robust to multiplicity then it induces a monotone price function.*

PROOF. In Appendix A.6.2

□

In other words, if $M \in \mathcal{M}$ induces a price function P that is non-monotone then there will be multiple equilibria. Two features of Theorem 1 are worth emphasising. First, the induced equilibrium price function P need not be increasing; it may be monotonically decreasing, even when R is strictly increasing in θ . Second, monotonicity of the induced P is not simply a consequence of measurability. This would be the case if we required Q to be continuous. However Q need not be continuous; continuity of M does not imply that of Q . However, continuity of Q will be an implication of Theorem 1 and robustness to multiplicity (Theorem 2).

To understand Theorem 1, notice first that any implementable non-monotone P must be discontinuous. If not then there would be a violation of measurability, i.e. prices at which the decision maker would need to take different actions depending on the state in order to implement P (this is easiest to see when $\theta \mapsto R(a, \theta)$ is strictly increasing, in which case there is at most one state satisfying $R(M(p), \theta) = p$ for any $p \in \mathcal{P}$). Figure 2.3 illustrates a discontinuous and non-monotone price function (the solid blue line). The dotted line plots $\theta_M(p) = \{\theta \in \Theta : R(M(p), \theta) = p\}$ for prices in the intervals (p_1, p_2) and (p_2, p_3) , for the case of a continuous M . As discussed above, θ_M is a convex and compact-valued, correspondence, and is upper hemicontinuous wherever M is continuous (θ_M is a function if $\theta \mapsto R(a, \theta)$ is strictly increasing). Moreover, by definition of a REE, P must be a selection from the graph of θ_M . Together, these properties of θ_M imply that there will be multiple equilibria for states above $P^{-1}(p_1)$, as shown in Figure 2.3. The proof of Theorem 1 formalizes this argument, using a generalization of the intermediate value theorem for convex, compact-valued, and upper-hemicontinuous correspondences, and extends the result from continuous M to all of \mathcal{M} .

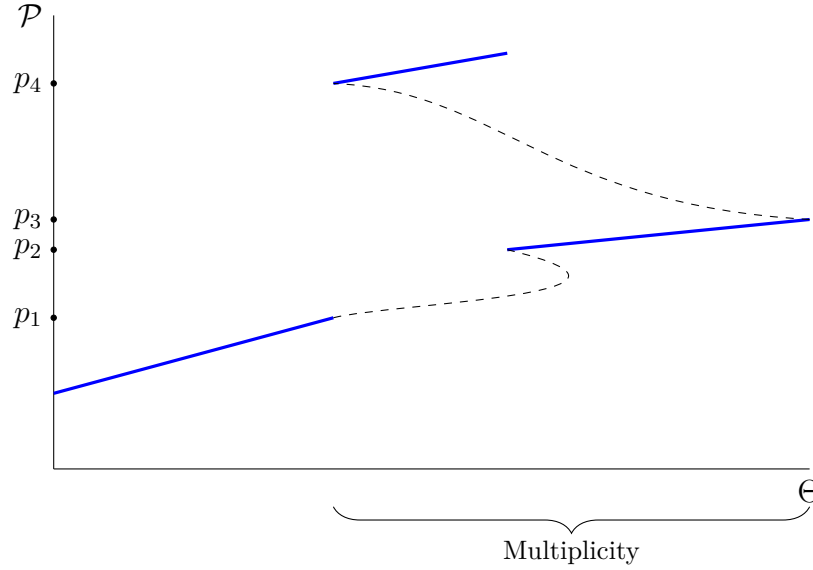


Figure 2.3. Multiplicity with a non-monotone price

2.3.1. Implementable action functions

In most situations the principal cares about the actions that they take. The price is determined in equilibrium by the action function, as $P(\theta) = R(Q(\theta), \theta)$. Therefore the joint distribution of state, price, and action is fully pinned down by the action function. The question is which action functions Q are CUI.²²

To understand the sufficient conditions for CUI, assume first that $\theta \mapsto R(a, \theta)$ is strictly increasing for all a . If Q is a continuous action function and $\theta \mapsto P(\theta) := R(Q(\theta), \theta)$ is strictly monotone then there will be no multiplicity for prices in $P(\Theta)$: define $M(p)$ on $P(\Theta)$ as the unique function satisfying $M(R(Q(\theta), \theta)) = Q(\theta)$. M is continuous since Q is continuous, and there can be no multiplicity involving prices in $P(\theta)$ since $|\{\theta \in \Theta : R(a, \theta) = p\}| \leq 1$ for all p under strict monotonicity of $\theta \mapsto R(a, \theta)$. While

²²In Appendix A.6.9 we explore the case in which the principal only cares about the price function, i.e. does not have direct preferences over actions.

it remains to define M on $\mathcal{P} \setminus P(\Theta)$ so as to avoid multiplicity, this argument suggests that strict monotonicity of the induced price function and continuity of Q are sufficient for CUI. With a minor caveat, this is indeed the case. The difficulty is that continuity of Q is not implied by continuity of M . It turns out however that Theorem 1 implies that continuity of Q is necessary a condition for CUI. This gives the characterization of CUI when R is strictly increasing, which requires only minor modifications when R is weakly increasing. Before stating the theorem, we require the following definitions.

Definition 13. *Local upper monotonicity* is satisfied at (a, θ) if there exists $\varepsilon > 0$ and a continuous function $m : [R(a, \theta), R(a, \theta) + \varepsilon] \rightarrow \mathcal{A}$ with $m(R(a, \theta)) = a$ such that $R(m(p), \theta) > p$ for all $p \in (R(a, \theta), R(a, \theta) + \varepsilon]$.

Definition 14. *Local lower monotonicity* is satisfied at (a, θ) if there exists $\varepsilon > 0$ and a continuous function $m : [R(a, \theta) - \varepsilon, R(a, \theta)] \rightarrow \mathcal{A}$ with $m(R(a, \theta)) = a$ such that $R(m(p), \theta) < p$ for all $p \in (R(a, \theta) - \varepsilon, R(a, \theta)]$.

While these conditions may appear dense, they essentially require that $a' \mapsto R(a', \theta)$ does not have weak local extremum at a (maximum for upper, minimum for lower). This is always necessary, it is sufficient when R is smooth; if $a' \mapsto R(a', \theta)$ is continuously differentiable at a then local upper (lower) monotonicity is satisfied at (a, θ) if and only if a is not a weak local maximum (minimum) of $a' \mapsto R(a', \theta)$. These conditions will be relevant only when the price function is decreasing.

Theorem 2. *Assume R is strictly increasing in θ . Then*

- (1) If Q is CUI then it is continuous on the interior of Θ and induces a strictly monotone price function. Moreover, if P is decreasing then local upper monotonicity is satisfied at $(Q(\underline{\theta}), \underline{\theta})$ and local lower monotonicity is satisfied at $(Q(\bar{\theta}), \bar{\theta})$.
- (2) If Q is continuous and induces a strictly increasing price function then it is CUI. If it induces a strictly decreasing price function, satisfies local upper monotonicity at $(Q(\underline{\theta}), \underline{\theta})$, and satisfies local lower monotonicity at $(Q(\bar{\theta}), \bar{\theta})$, then Q is CUI.

PROOF. Proof in Appendix A.6.3. □

The only difference between the necessary conditions for CUI of Theorem 2 part *i*, and the sufficient conditions of part *ii* concerns discontinuities at the extreme states $\underline{\theta}$ and $\bar{\theta}$. In most applications this gap can be closed: the necessary and sufficient conditions for CUI will be strict monotonicity of P and continuity of Q on the interior of Θ . This will hold when any discontinuity at $\underline{\theta}$ is lower-bridgeable, and any discontinuity at $\bar{\theta}$ is upper-bridgeable, as defined in Appendix A.7. The stronger notion of bridgeability, discussed in Section 2.3.2, is also sufficient. Since discontinuities at the extreme states are of no substantive interest we defer discussion of these conditions, but note that if Θ is open, Theorem 2 can be stated more succinctly.

Corollary 1. *Assume R is strictly increasing in θ and Θ is open. Then Q is CUI iff it is continuous and induces a strictly monotone price.*

It is worth pointing out that any CUI Q that induces an increasing price can be implemented with a continuous M .

Lemma 2. *A continuous and CUI Q can be uniquely implemented by a continuous M if and only if it induces an increasing price.*

Discontinuities in M (which must occur outside of the set of equilibrium prices under the restriction to essentially continuous decision rules) are only useful for implementing decreasing price functions. This is shown in the construction of M given in the proof of Theorem 2. When the price function P is decreasing, it is necessary to have two discontinuities in the decision rule, one above $P(\underline{\theta})$ and one below $P(\bar{\theta})$, in order to prevent the price function from “bending back”. This is discussed in greater detail in Section 2.6.2.

Relaxing the assumption of strictly increasing $\theta \mapsto R(a, \theta)$ to weakly increasing, we obtain a similar characterization to Theorem 2. It is necessary however to add an additional condition to account for actions for which the induced price is constant over an interval of states. Let $r(a, p) = \{\theta \in \Theta : R(a, \theta) = p\}$. Under strictly increasing R , $r(a, p)$ contains at most one state for all $a \in \mathcal{A}, p \in \mathcal{P}$. Under weakly increasing R however, $r(a, p)$ may be a non-degenerate interval.

Let $P(\theta) = R(Q(\theta), \theta)$, and suppose $r(Q(\theta'), P(\theta'))$ is non-degenerate. If $Q(\theta'') \neq Q(\theta')$ for some $\theta'' \in r(Q(\theta'), P(\theta'))$ then clearly there will be multiplicity, since $R(Q(\theta'), \theta'')$ is an REE price in state θ'' , by definition of r . The only modifications needed to extend Theorem 2 are those that rule out such instances of multiplicity.

Corollary 2. *Assume R is weakly increasing in θ . Then*

- *If Q is CUI then it is continuous on the interior of Θ , induces a weakly monotone price function, and for all $\theta, \theta' \in r(Q(\theta), P(\theta))$ we have $Q(\theta') = Q(\theta)$. Moreover if P is decreasing then local upper monotonicity is satisfied at $(Q(\underline{\theta}), \underline{\theta})$ and local lower monotonicity is satisfied at $(Q(\bar{\theta}), \bar{\theta})$.*

- If Q is continuous, induces a weakly increasing price function, and for all θ , $\theta' \in r(Q(\theta), P(\theta))$ we have $Q(\theta') = Q(\theta)$, then Q is CUI. If it induces a weakly decreasing price function and additionally satisfies local upper monotonicity at $(Q(\underline{\theta}), \underline{\theta})$ and local lower monotonicity at $(Q(\bar{\theta}), \bar{\theta})$, then Q is CUI.

Under the condition that for all θ , $\theta' \in r(Q(\theta), P(\theta))$ we have $Q(\theta') = Q(\theta)$, Corollary 2 follows from the same argument as Theorem 2. Again, the gap between the necessary and sufficient conditions is closed under mild assumptions, discussed in Section 2.3.2.

2.3.2. Continuous weakly unique implementation

The necessity of continuity of Q in Theorem 2 and Corollary 2 is an implication of Theorem 1 and the requirement of uniqueness for all θ . The substantive characteristic of CUI outcomes however is monotonicity of P . This section formalizes this assertion. We show that in many common settings continuous *weakly* unique implementability will be characterized fully by monotonicity of P .

We begin with some preliminary observations. First, since $\theta \mapsto R(Q(\theta), \theta)$ must be monotone, by Theorem 1, any discontinuity in P must be a jump discontinuity, and P can have at most countably many discontinuities. Moreover, Q can be discontinuous at θ only if P is as well: otherwise it would not be possible for Q to be implemented by an M that is continuous at $P(\theta)$. Thus Q can also have at most countably many discontinuities. Finally, recall that M must be continuous on $cl(P(\Theta))$ (this is simply a rephrasing of the definition of essential continuity). This implies that the one-sided limits of Q , denoted by $\lim_{\theta \nearrow \theta'} Q(\theta)$ and $\lim_{\theta \searrow \theta'} Q(\theta)$, must exist for all θ' .

Roughly, Q can have discontinuities only if it can be well approximated by a continuous Q' . For any two actions $a', a'' \in \mathcal{A}$, a *path* from a' to a'' is a continuous function $\gamma : [0, 1] \rightarrow \mathcal{A}$ such that $\gamma(0) = a', \gamma(1) = a''$. Say that there exists a *monotone path* from a' to a'' at θ if there exists a path γ from a' to a'' such that $x \mapsto R(\gamma(x), \theta)$ is strictly monotone.

Definition 15. A discontinuity in Q at θ' is **bridgeable** if there exists a monotone path from $\lim_{\theta \nearrow \theta'} Q(\theta)$ to $\lim_{\theta \searrow \theta'} Q(\theta)$ at θ' .

As alluded to above, a necessary condition for a discontinuity at θ to be bridgeable is $\lim_{\theta \nearrow \theta'} R(Q(\theta), \theta) \neq \lim_{\theta \searrow \theta'} R(Q(\theta), \theta)$. We say that the environment is *fully bridgeable* if for every θ , this condition is also sufficient for bridgeability. Finally, say that the environment is *continuously bridgeable* if for any $\theta^* \in \Theta$ there exists $\varepsilon > 0$ such that if the gap between a' and a'' is bridgeable at θ^* and $R(a'', \theta) \neq R(a', \theta)$ for all $\theta \in [\theta^*, \theta^* + \varepsilon]$ then there exists a sup-norm continuous function $\sigma(\cdot | a', a'') : [\theta^*, \theta^* + \varepsilon] \rightarrow \mathcal{A}^{[0,1]}$ such that $\sigma(\theta | a', a'')$ is a monotone path from a' to a'' for all $\theta \in [\theta^*, \theta^* + \varepsilon]$. Say that the environment is *continuously fully bridgeable* if it is full bridgeable and continuously bridgeable. Bridgeability, and the related notions, will be discussed further following the statement of the results.

Proposition 9. Assume $\theta \mapsto R(a, \theta)$ is strictly increasing for all $a \in \mathcal{A}$. Then Q is CWUI iff

- (1) Either $P(\theta) := R(Q(\theta), \theta)$ is strictly increasing; or it is strictly decreasing and satisfies local upper monotonicity at $(Q(\underline{\theta}), \underline{\theta})$ and local lower monotonicity at $(Q(\bar{\theta}), \bar{\theta})$.
- (2) Any discontinuity in Q on the interior of Θ is bridgeable.

(3) If Q is discontinuous at $\underline{\theta}$ then local upper monotonicity is satisfied at

$(\min\{\lim_{\theta \searrow \underline{\theta}} Q(\theta), Q(\underline{\theta})\}, \underline{\theta})$ and local lower monotonicity at $(\max\{\lim_{\theta \searrow \underline{\theta}} Q(\theta), Q(\underline{\theta})\}, \underline{\theta})$

(4) If Q is discontinuous at $\bar{\theta}$ then local upper monotonicity is satisfied at

$(\min\{\lim_{\theta \nearrow \bar{\theta}} Q(\theta), Q(\bar{\theta})\}, \bar{\theta})$ and local lower monotonicity at $(\max\{\lim_{\theta \nearrow \bar{\theta}} Q(\theta), Q(\bar{\theta})\}, \bar{\theta})$.

PROOF. In Appendix A.6.4. □

If the principal's payoffs are invariant to changes on zero measure sets then conditions *iii.* and *iv.* can be ignored for the purposes of choosing optimal policies; we can restrict attention to Q that are continuous at the endpoints. When the environment is fully bridgeable the type of discontinuities in Q that are allowed can be more easily characterized.

Corollary 3. Assume $\theta \mapsto R(a, \theta)$ is strictly monotone for all $a \in \mathcal{A}$, and the environment is fully bridgeable. Then Q is CWUI iff

(1) $P(\theta) := R(Q(\theta), \theta)$ is strictly monotone. Moreover if P is decreasing then local upper monotonicity is satisfied at $(Q(\underline{\theta}), \underline{\theta})$ and local lower monotonicity is satisfied at $(Q(\bar{\theta}), \bar{\theta})$.

(2) If Q is discontinuous at θ then so is P .

Finally, it will be useful to know when condition *ii* in Corollary 3 is redundant. This will be the case when any discontinuities that violate this condition can be well approximated. Say that Q has a *degenerate discontinuity* at θ if Q is discontinuous at θ and P is not. The environment is *correctable* if for and $\varepsilon > 0$, any strictly monotone Q , and

any θ at which Q has a degenerate discontinuity, there exists a monotone Q' that has no degenerate discontinuities in $(\theta - \varepsilon, \theta + \varepsilon)$ and such $Q' = Q$ on $\Theta \setminus (\theta - \varepsilon, \theta + \varepsilon)$. Sufficient conditions for correctability are discussed in Appendix A.7.

Corollary 4. *Assume $\theta \mapsto R(a, \theta)$ is strictly monotone for all $a \in \mathcal{A}$, and the environment is fully bridgeable and correctable. Then Q is virtually CWUI iff*

- (1) $P(\theta) := R(Q(\theta), \theta)$ is strictly monotone. Moreover if P is decreasing then local upper monotonicity is satisfied at $(Q(\underline{\theta}), \underline{\theta})$ and local lower monotonicity is satisfied at $(Q(\bar{\theta}), \bar{\theta})$.
- (2) The set of states at which Q is discontinuous has zero measure.

PROOF. In Appendix A.6.5. □

Bridgeability of discontinuities and correctability of the environment are less transparent conditions than monotonicity of the price, and so it will be useful to know general conditions under which they are satisfied. If \mathcal{A} is a subset of \mathbb{R} then clearly a discontinuity at θ with left limit \underline{a} and right limit \bar{a} is bridgeable iff $a \mapsto R(a, \theta)$ is strictly monotone on $[\min\{\underline{a}, \bar{a}\}, \max\{\underline{a}, \bar{a}\}]$.²³ When the action space is multi-dimensional the condition becomes more difficult to check, but also easier to satisfy. This is because there will be a continuum of paths between two actions, as opposed to only one in the one-dimensional case. Proposition 30 shows that a weak monotonicity condition is sufficient for full bridgeability.

A full discussion of bridgeability, correctability, and related notions is contained in Appendix A.7. This section gives general conditions under which every discontinuity is

²³This does not mean that $a \mapsto (a, \theta)$ is monotone in the same direction in every state; it could be increasing in some states and decreasing in others.

bridgeable. In most applications encountered in the literature it is easy to verify that the environment is continuously fully bridgeable and correctable. Even when it is not, the states at which these conditions fail are readily identifiable.

An alternative way to understand the conditions of Proposition 9 is in terms of approximations to Q . Say that Q' is a continuous ε -approximation of Q if Q' is continuous and $\lambda(\{\theta \in \Theta : Q(\theta) \neq Q'(\theta)\}) < \varepsilon$, where λ is Lebesgue measure.²⁴

Proposition 10. *Assume $\theta \mapsto R(a, \theta)$ is strictly increasing for all $a \in \mathcal{A}$ and the environment is continuously fully bridgeable. Then if Q is CWUI there exists a continuous ε -approximation Q' that is CUI, for any $\varepsilon > 0$.*

PROOF. In Appendix A.6.6. □

In other words, Proposition 10 says that the space \mathcal{Q} of continuous Q which induce a strictly monotone price is dense in the space of CWUI Q . Proposition 10 can help simplify the problem of solving for an optimal policy. Any such $Q \in \mathcal{Q}$ will be CUI, and by Proposition 10 there is no loss of optimality, provided the principal's payoffs are continuous.

Finally, we can extend the characterization by relaxing strictly increasing $\theta \mapsto R(a, \theta)$ to weakly increasing. Say γ is a *proper monotone path* from a' to a'' at θ if it is a monotone path, and moreover $r(\gamma(x), R(\gamma(x), \theta)) = \theta$ for all $x \in [0, 1]$. A discontinuity in Q at θ is *properly bridgeable* if there exists a proper monotone path from $\lim_{\theta' \nearrow \theta} Q(\theta')$ to $\lim_{\theta' \searrow \theta} Q(\theta')$ at θ . The environment is fully properly bridgeable if all non-degenerate

²⁴Since Θ and \mathcal{A} are compact, if there is an ε -approximation for any ε then there is a sequence that approaches Q in the L^p norm, for any $p < \infty$.

discontinuities are properly bridgeable. Note that if $\theta \mapsto R(a, \theta)$ is strictly increasing for all a then proper bridgeability is equivalent to bridgeability.

Proposition 11. *Assume $\theta \mapsto R(a, \theta)$ is weakly increasing for all a . Then Q is CWUI iff*

- (1) $P := R(Q(\theta), \theta)$ is weakly monotone. Moreover if P is decreasing then local upper monotonicity is satisfied at $(Q(\underline{\theta}), \underline{\theta})$ and local lower monotonicity is satisfied at $(Q(\bar{\theta}), \bar{\theta})$.
- (2) Any discontinuity in Q on the interior of Θ is properly bridgeable.
- (3) A discontinuity in Q at $\underline{\theta}$ is lower bridgeable, and at $\bar{\theta}$ is upper bridgeable.
- (4) $Q(\theta) = Q(\theta')$ for all $\theta' \in r(Q(\theta), P(\theta))$ and all θ .

PROOF. In Appendix A.6.8. □

The first condition is necessary by Theorem 1. It is sufficient given the other two conditions. Under monotonicity, condition *iii.* guarantees that the measurability restriction is satisfied, so an implementing M can be found. Condition *ii.* guarantees that an $M \in \mathcal{M}$ can be found that implements Q .

2.3.3. Manipulation and multiplicity singly

The restrictions on the set of implementable pairs (Q, P) in the characterization results above, such as Theorem 2, come from the interaction between the constraints imposed by manipulation and multiplicity concerns. Individually, these constraints are not restrictive.

Suppose the principle would like to maintain unique implementation, but is not concerned with manipulation, i.e. is not restricted to use $M \in \mathcal{M}$. Then, under weak conditions on the environment, the set of uniquely implementable (Q, P) are exactly those satisfying the conditions of Observation 1. In other words, unique implementation alone does not impose restrictions on the implementable set.

Say that the environment is *omitable* if for an $p \in \mathcal{P}$ there exists an action $a \in \mathcal{A}$ such that $\{\theta : R(a, \theta) = p\} = \emptyset$. This is satisfied in all the examples covered in this paper.

Lemma 3. *Assume $\theta \mapsto R(a, \theta)$ is strictly increasing for all a and the environment is omitable. If the principal can use any discontinuous decision rule $M : \mathcal{P} \rightarrow \mathcal{A}$ then the set of uniquely implementable (Q, P) is exactly the set of (Q, P) satisfying the conditions of Observation 1.*

PROOF. Let (Q, P) satisfy the conditions of Observation 1. Since $\theta \mapsto R(a, \theta)$ is strictly increasing for all a , there can only be multiplicity if some price $p' \notin P(\Theta)$ clears the market in some state. If the environment is omitable we can define M to rule this out. \square

Similarly, any (Q, P) satisfying commitment and measurability will be virtually implementable via an essentially continuous decision rule. In other words, manipulation concerns alone do not impose a qualitative constraint on the feasible set. The pair (Q, P) cannot be implemented by $M \in \mathcal{M}$ if and only if $Q \circ P^{-1}$ is discontinuous at some $p' \in P(\Theta)$. However if this is the case, we can perturb Q so that the associated price function is discontinuous in a neighborhood of $P^{-1}(p')$. Thus we are able to eliminate the discontinuity in $Q \circ P^{-1}$ at p' .

Say that the environment satisfies *single-crossing* if for any a, a' , the functions $\theta \mapsto R(a, \theta)$ and $\theta \mapsto R(a, \theta')$ intersect at most once.²⁵

Lemma 4. *Assume $\theta \mapsto R(a, \theta)$ is strictly increasing for all $a \in \mathcal{A}$ and that the environment satisfies single crossing. Let (Q, P) satisfy the conditions of Observation 1 and assume the set of states at which Q is discontinuous has zero measure. Then for any $\varepsilon > 0$ there is an ε -approximation of (Q, P) that is implementable with an $M \in \mathcal{M}$*

PROOF. Proof in Appendix A.6.13.4. □

2.4. Characterization: multi-dimensional state space

The assumption of a one-dimensional state space precludes a number of interesting applications, and so we extend the results here to multi-dimensional state spaces. However, it is worth first noting a way in which the one-dimensional environment is not as restrictive as it may seem. This restriction only enters the analysis through the assumption that $\theta \mapsto R(a, \theta)$ is increasing (strictly or weakly) for all $a \in \mathcal{A}$. This is a purely order-theoretic assumption; it does not directly concern the dimensionality of the state space. We could alternatively allow the state space be an arbitrary subset of a Banach space, but assume that there exists a complete order \succsim on the state space, with respect to which we can assume $\theta \mapsto R(a, \theta)$ is increasing (provided the interval structure is preserved by \succsim , i.e. $\theta'' \succsim \theta' \Rightarrow \theta'' \succsim \alpha\theta'' + (1 - \alpha)\theta' \geq \theta'$ for all $\alpha \in [0, 1]$). In this case the state space can be effectively compressed to a single dimension, and the analysis for the one-dimensional case applies.

²⁵Similar results obtain under weaker conditions.

In this section we discuss the general extension of the results to settings in which it is not possible to immediately map the problem to one with a uni-dimensional state space. We will then pay particular attention to the common example of an asset market with aggregate supply shocks, as this model is commonly used in the finance literature on feedback effects.

Suppose that Θ is a compact subset of \mathbb{R}^N , endowed with the usual product partial order. Assume that $R : \mathcal{A} \times \Theta \rightarrow \mathbb{R}$ is continuous, and is increasing with respect to the partial order on Θ , i.e.

Strictly increasing R . $\theta'' > \theta'$ implies $R(a, \theta'') > R(a, \theta')$ for all a .²⁶

Define $\bar{R}(a, \theta) = \{\theta' : R(a, \theta') = R(a, \theta)\}$. That is, $\bar{R}(a, \theta)$ is the level set of $R(a, \cdot)$ corresponding to the price $R(a, \theta)$. Under the assumptions of continuous and strictly increasing R , $\bar{R}(a, \theta)$ is a one-dimensional curve in Θ for any a, θ . The problem can be reduced to one with a uni-dimensional state space if and only if $\bar{R}(a, \theta) = \bar{R}(a', \theta)$ for all a, a' and θ ; if this condition does not hold then there is no complete order on Θ with respect to which R is monotone for any a . Nonetheless, we will be able to characterize CUI in this setting.

Clearly if there is an equilibrium under M in which the principal takes action a in state θ then there is an equilibrium in which the principal takes action a for all states in $\bar{R}(a, \theta)$. Therefore robustness to multiplicity implies that for all θ , $Q(\theta') = Q(\theta)$ for all $\theta' \in \bar{R}(Q(\theta), \theta)$. Robustness to multiplicity also implies that $Q(\theta) \neq Q(\theta') \Rightarrow \bar{R}(Q(\theta), \theta) \cap \bar{R}(Q(\theta'), \theta') = \emptyset$; otherwise there would be multiple equilibrium actions for any states in the intersection. As a result of these two observations, the previous characterizations of

²⁶This can be easily relaxed to weakly increasing.

CUI action and price functions can be extended without much difficulty. The following result is analogous to Theorem 2 in the uni-dimensional case.

Proposition 12. *Assume strictly increasing R . Then Q is CUI if*

(1) $P(\theta) := R(Q(\theta), \theta)$ is strictly monotone (in the product partial order on Θ).

Moreover if P is decreasing then local upper monotonicity and local lower monotonicity are satisfied at the actions and states corresponding to the highest and lowest prices respectively.

(2) Q is continuous.

(3) For all θ , $Q(\theta') = Q(\theta)$ for all $\theta' \in \bar{R}(Q(\theta), \theta)$.

(4) $Q(\theta) \neq Q(\theta') \Rightarrow \bar{R}(Q(\theta), \theta) \cap \bar{R}(Q(\theta'), \theta') = \emptyset$.

These conditions are also necessary, except that it may be possible for Q to have discontinuities at the states associated with the highest prices (see Section 2.3 for discussion).

PROOF. Proof in Appendix A.6.13.5 □

Proposition 12 continues to hold when Θ is unbounded. This is useful to note because some applications, such as the noise REE model of the next section, will make use of an unbounded state space.

2.4.1. Reduced-form with multi-dimensional state space

The derivation of a reduced form representation for some markets, such as the REE asset market discussed in Appendix A.8, makes use of assumptions related to the monotonicity of agents' actions and information as a function of the state. However when the state space is multi-dimensional, there is generally no complete order on Θ such that these

assumptions hold. This complicates the derivation of the reduced-form representation. One approach to this problem is to derive a reduced-form representation as in the one-dimensional setting. However in some markets, such as the noisy REE market studied below, this will hold only under additional refinements to the solution concept. Alternatively, we can anticipate that our objective will in the end be *unique* implementation, and derive a reduced-form representation under this restriction.

Definition 16. *The market **admits a reduced-form representation under uniqueness** if \exists a function $R : \mathcal{A} \times \Theta \rightarrow \mathcal{P}$ such that for any Q, P, M , the pair (Q, P) are the unique equilibrium outcomes given M iff for all θ*

- (1) $Q(\theta) = M(P(\theta))$ (commitment)
- (2) $P(\theta) = R(Q(\theta), \theta)$ (market clearing)
- (3) $\{p : p = R(M(p), \theta)\}$ is singleton (uniqueness)

The only difference between this definition and that of the reduced form is that we require that (Q, P) are the unique equilibrium outcomes given M , and impose the uniqueness condition.

2.4.2. Noisy REE in asset markets

Asset markets are an important setting in which decision making under feedback effects occurs. Since Grossman and Stiglitz (1980) and Hellwig (1980), the noisy rational expectations model has been a workhorse model for studying asymmetric information in asset markets. This model adds shocks to aggregate supply, interpreted as noise or liquidity traders, to a rational expectations model of the asset market. The standard approach,

without feedback effects, is to assume joint normality of asset returns and aggregate demand shocks, and look for equilibria in which the price is linear in trader's private signals. Breon-Drish (2015) generalizes the noisy REE model to allow for non-normal distributions of states and supply shocks. This section extends results from Breon-Drish (2015) to a setting with feedback effects.

The setting is as follows. There is a single asset that pays an ex-post dividend of $\pi(a, \omega)$, where $\omega \in \Omega$ is referred to as the payoff-relevant state. We assume that π is continuous and is affine in θ for all a ; $\pi(a, \omega) = \beta_0^a + \beta_1^a \omega$. Each investor observes an additive signal $s_i = \omega + \varepsilon_i$, where $\varepsilon_i \sim N(0, \sigma_i^2)$, where σ_i^2 lies in a bounded set. The supply shock is a random variable z taking values in \mathcal{Z} . We assume that z has a truncated normal distribution. That is, z is the restriction of a normal random variable $\hat{z} \sim N(0, \sigma_Z^2)$ to the interval $[b_1, b_2]$, with $-\infty \leq b_1 \leq 0 \leq b_2 \leq \infty$ (note that this assumption accommodates un-truncated supply shocks as well). For simplicity, let $b_1 = -b_2$; this does not affect the results. The state θ consists of both the payoff-relevant state ω and the supply shock z .

There are a continuum of investors $i \in [0, 1]$, each with CARA utility $u(w) = -\exp\left\{-\frac{1}{\tau_i}w\right\}$. The ex-post payoff to an investor who purchases x units of the asset at price p when the principal takes action a is given by $-\exp\left\{-\frac{1}{\tau_i}x(\pi(a, \theta) - p)\right\}$, where τ_i lies in some bounded set. We assume that the distribution of private signals in the population is uniquely determined by the state ω (this is the usual “continuum law of large numbers” convention). Let $x_i(p|a, \mathcal{I}, s_i)$ be the demand of investor i when the price is p , the anticipated principal action is a , and the public information revealed by the price is that $(\omega, z) \in \mathcal{I}$, and i 's private signal is s_i . Aggregate demand can be written as $X(p|a, \mathcal{I}, \omega)$.

As in the one-dimensional case, equilibrium consists of a price function as well as a specification of the public information for off-path prices. To be precise, fixing a decision rule M , an equilibrium is characterized by a price function $P : \Omega \times \mathcal{Z} \rightarrow \mathcal{P}$ and an off-path inference function $\lambda : \mathcal{P} \setminus P(\Omega, \mathcal{Z}) \rightarrow 2^{(\Omega, \mathcal{Z})}$. The price function P is such that for all (ω, a) market's clear given the anticipated action and the information revealed by the price, that is: $X(P(\omega, z) | M(P(\omega, z), \mathcal{I}(\omega, z)), \omega) = z$, where $\mathcal{I}(\omega, z) = \{(\omega', z') : P(\omega', z') = P(\omega, z)\}$. For the off-path information, we assume only that it is consistent with market clearing (when possible), that is: $\lambda(p) \subseteq \{(\omega, z) : X(p | M(p), \lambda(p), \omega) = z\}$.

The approach in this setting is somewhat distinct from the previous analysis, in that we do not begin by deriving a reduced-form representation directly from the primitives of the model. Rather, we show that there exists a reduced-form representation that can be used to design policy *under uniqueness*. However, the search for truly unique implementation is hopeless in the noisy REE model studied here, since there are multiple equilibria even when there is no policy feedback, that is, fixing the principal's action (Pálvölgyi and Venter, 2015). We therefore focus here on a more limited, but still meaningful, notion of uniqueness. What we really want to rule out is multiplicity arising from the endogeneity of the principal's action. Therefore, in the context of the noisy REE model, we say that M is robust to multiplicity if there is a unique market clearing price in every state, fixing the inferences draw from prices, that is, fixing the public information sets associated with each price both on and off path. Another interpretation is that we require a unique market clearing price fixing the demand schedules submitted to the market maker by each agent. Multiplicity that violates this restriction requires no change in agent behavior, simply a change in the selection of the market clearing price by the market maker.

Relative to the model with a one-dimensional state space, the complication in this setting in which the state space $\Theta = \Omega \times \mathcal{Z}$ is two dimensional, is that there is no easy way to narrow down the space of possible public information sets that can be revealed by the price. This makes it difficult to derive a reduced-form representation of the market ex-ante, without strong restrictions on the set of possible equilibria. To deal with, this difficulty, we instead analyse directly the problem of characterizing what equilibria can be induced with a decision rule $M \in \mathcal{M}$ that is robust to multiplicity. Under these restrictions, we show that the market admits a reduced-form representation under uniqueness. To do this, we will first need some preliminary results. The first concerns the continuity of aggregate demand.

Lemma 5. *For and $\mathcal{I} \subseteq \Omega \times \mathcal{Z}$, $p \in \mathcal{P}$, and $a \in \mathcal{A}$, the function $\omega \mapsto X(p|a, \mathcal{I}, \omega)$ is Lipschitz continuous.*

PROOF. Proof in Appendix A.6.13.6. □

Note that since the distribution of signals in the population is uniquely determined by ω (following the usual continuum law of large numbers convention) it cannot be that any public information set \mathcal{I} contains states (ω', z') and (ω', z'') with $z'' \neq z'$, since the aggregate demand would be the same in both cases. Therefore, the distribution of ω conditional on \mathcal{I} cannot have atoms. The following lemma strengthens this observation slightly, by showing that in fact, given Lipschitz continuity of aggregate demand, the distribution of ω conditional on \mathcal{I} will be absolutely continuous.

Lemma 6. *For any $p \in \mathcal{P}$ and $a \in \mathcal{A}$, let \mathcal{I} be a set satisfying $\mathcal{I} \subseteq \{(\omega, z) : X(p|a, \mathcal{I}, \omega) = z\}$. Then the distribution of ω conditional on \mathcal{I} is absolutely continuous.*

PROOF. Proof in Appendix A.6.13.7. \square

Lemmas 5 and 6 did not make use of much of the structure that we have assumed; for example, CARA utility and truncated-normal noise distributions. Using these properties, we can establish further characteristics of public information sets. The following lemma says that any public information set, either one revealed on-path by the price or by the off-path inference function, must lie in a linear subset of $\Omega \times \mathcal{Z}$. In other words, and such \mathcal{I} must be a subset of some set of the form $\{(\omega, z) : k \cdot \omega - z = \ell\}$ for some $k > 0$ and ℓ .

Lemma 7. *Let \mathcal{I} satisfy $\mathcal{I} \subseteq \{(\omega, z) : X(p|a, \mathcal{I}, \omega) = z\}$ for some p, a . Then there exists $k > 0$ and ℓ such that $\mathcal{I} \subseteq \{(\omega, z) : k \cdot \omega - z = \ell\}$*

PROOF. Proof in Appendix A.6.13.8. \square

The following proposition identifies exactly which hyperplanes the public information sets can lie in.

Proposition 13. *Assume CARA utility, π affine in θ and continuous, additive normal signal structure and truncated-normally distributed supply shocks. Then there exists a unique (up to positive transformations) function $L^* : \Omega \times \mathcal{Z} \times \mathcal{A} \rightarrow \mathbb{R}$ defined by*

$$(2.1) \quad L^*(\omega, z|a) = \left(\frac{1}{\beta_1^a} \int_i \frac{\tau_i}{\sigma_i^2} di \right) \cdot \omega - z$$

such that for any M , if \mathcal{I} is the public information revealed at price p (in which case $\mathcal{I} \subseteq \{(\omega, z) : X(p|a, \mathcal{I}, \omega) = z\}$) then $L^(\omega'', z''|M(p)) = L^*(\omega', z'|M(p))$ for all $(\omega'', z''), (\omega', z') \in \mathcal{I}$*

PROOF. Proof in Appendix A.6.13.9. \square

We now wish to use these properties, in particular Lemma 7, to identify features of equilibrium. Proposition 13 identifies the hyperplane to which each information set belongs. Following Breon-Drish (2015), we refer to these hyperplanes as *linear statistics*. So in other words, the public information will always reveal *at least* the associated linear statistic. The following result says that in fact, under robustness to multiplicity and $M \in \mathcal{M}$, the equilibrium price function will reveal *exactly* the linear statistic, and no more.

Proposition 14. *Maintain the assumptions of Proposition 13. If $M \in \mathcal{M}$ is robust to multiplicity then the level sets of the equilibrium price function \tilde{P} are given by $\{(\omega, z) : L^*(\omega, z | M(p)) = \ell\}$ for some ℓ , where L^* is given by (2.1).*

PROOF. Proof in Appendix A.6.13.10 □

From (2.1) we can see how the principal's action affects information aggregation; the higher is β_1^a , i.e. the more sensitive the asset value is to the state, the smaller the coefficient on θ in the equilibrium statistic. As a result, the price is less informative about the state. This is because when β_1^a is high, each trader's private signal is less informative about the asset value. As a result, traders place less weight on their private information relative to the information revealed by the price. The linear statistics for a fixed action $a \in \mathcal{A}$ are pictured in Figure 2.4 (in this figure the sign of the supply shock has been reversed, so that prices are increasing in the usual Euclidean product order). The slope of the linear statistics, is $-\frac{1}{\beta_1^a} \int_i \frac{\tau_i}{\sigma_i^2} di$, which again illustrates that the price reveals more precise information about ω the lower is β_1^a .

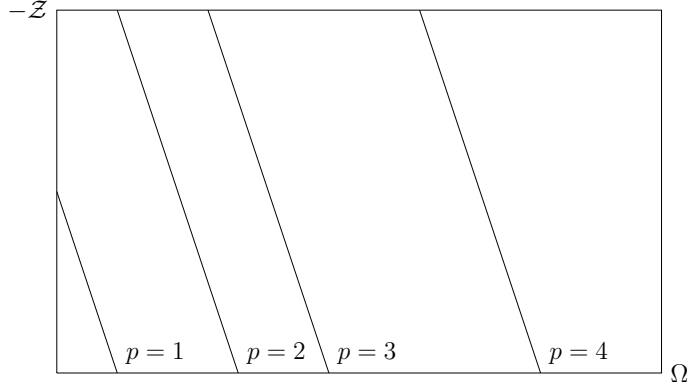


Figure 2.4. Linear statistics, fixed action $a \in \mathcal{A}$

The proof of Proposition 13 also yields an expression for $R(a, \theta, z)$, although for the current purposes it is sufficient to note simply that such a function exists and is strictly increasing (with the product partial order on $\Theta \times \mathcal{Z}$).

Corollary 5. *Assume CARA utility, π affine in θ , additive normal signal structure and truncated-normally distributed supply shocks. Then the market admits a reduced form representation under uniqueness, given by $R : \mathcal{A} \times \Theta \times \mathcal{Z} \rightarrow \mathcal{P}$. Moreover, $(\theta, z) \mapsto R(a, \theta, z)$ is strictly increasing for all a .*

There are three differences between the environment of Proposition 13 and that of Breon-Drish (2015) Proposition 2.1. First, the signal σ_i observed by each investor is given by the state plus noise, as opposed to the asset return plus noise as in Breon-Drish (2015). This is immediately handled by a suitable change of variables, given the assumption that $\theta \mapsto \pi(a, \theta)$ is affine for all a . Second, we allow here for the supply shock to follow a truncated normal distribution, where Breon-Drish (2015) considers only the un-truncated distribution. This requires generalizing Breon-Drish (2015) Proposition 2.1, which is relatively straightforward. Finally, and most importantly, the current setting features a

feedback effect, whereas asset returns follow a fixed distribution in Breon-Drish (2015). We show how the results for the fixed-action case imply the desired result when there is feedback.

The approach here is similar to Siemroth (2019). However that paper assumes that the asset value is additively separable in the state and the principal's action. This is more than a technical assumption; it implies, as the author demonstrates, that the information revealed by the price is the same in all equilibria, regardless of the principal's actions. In contrast, we show precisely how the relationship between the principal's action and the asset value affects the information content of the price. This connection between the principal's action and the type of information revealed by the price has important implications for equilibrium multiplicity, as discussed below (e.g. Lemma 8). Siemroth (2019) also restricts attention to equilibria in which the price function is continuous, which has substantive implications, as discussed in the introduction.

Given Corollary 5, we can apply Proposition 12 to the noisy REE setting. The problem of finding optimal policies is generally complicated by the additional restrictions *iii.* and *iv.*, relative to the uni-dimensional case. In some cases however, these constraints simplify the problem. For example, if the supports of the noise term and the state are unbounded, these conditions and the expression for L^* in (2.1) have the following implication.

Lemma 8. *In the noisy REE model with normally distributed supply shocks and unbounded Θ , any CUI action function must be such that $\beta_1^{Q(\theta,z)} = \beta_1^{Q(\theta',z')}$ for all $(\theta, z), (\theta', z') \in \Theta \times \mathcal{Z}$.*

In other words, Lemma 8 says that any CUI action function can only use actions for which the slope of the asset payoff with respect to the state is the same. This will not be true when the supply shocks are bounded; in this case additional action functions will be CUI.

2.5. Properties and extensions

In this section we discuss properties of CUI policies, and study optimal policy when the unique implementation requirement is relaxed.

2.5.1. Structural uncertainty

Another practical concern of the principal, aside from manipulation and multiplicity, is that the price may be influenced by uncertain factors other than the state in which the principal is interested. For example, the presence of noise/liquidity traders in an asset market could introduce aggregate uncertainty. As a consequence, the price may not be a deterministic function of the state and anticipated action. Additionally, the principal may simply have limited information about market fundamentals, which within the model translates into uncertainty about the function R . The principal will want to choose a decision rule that is robust to these types of uncertainty when the degree of uncertainty is small.²⁷

Assume throughout this section that Θ is closed. Endow the space of market-clearing functions $R : \mathcal{A} \times \Theta \rightarrow \mathbb{R}$ with the sup-norm. Let \mathcal{C} be the set of continuous functions on $\mathcal{A} \times \Theta$. For a given decision rule M and market clearing function R , let $\tilde{Q}_R(\theta|M) :=$

²⁷An alternative approach to additional dimensions of uncertainty is to model them explicitly. We show how this can be done in Section 2.4.

$\{a \in \mathcal{A} : M(R(a, \theta)) = a\}$. In words, $\tilde{Q}_R(\theta|M)$ is the set of actions that are consistent with rational expectations in state θ .

An open neighborhood of $\tilde{Q}_R(\cdot|M)$ is a set-valued and open-valued function $U : \Theta \rightarrow \mathcal{A}$ such that $\tilde{Q}_R(\theta|M) \subset U(\theta)$ for all θ . The map $R \rightrightarrows \tilde{Q}_R(\theta|M)$ is *uniformly continuous* at R if it is uniformly upper and lower hemicontinuous. That is, for any open neighborhood U of $\tilde{Q}_R(\cdot|M)$ and any open-valued function $V : \Theta \rightarrow \mathcal{A}$ such that $\tilde{Q}_R(\theta|M) \cap V(\theta) \neq \emptyset$ for all θ , there exists an open neighborhood N of R such that $\hat{R} \in N$ implies, for all $\theta \in \Theta$, *i*) $\tilde{Q}_{\hat{R}}(\theta|M) \subset U(\theta)$, and *ii*) $\tilde{Q}_{\hat{R}}(\theta|M) \cap V(\theta) \neq \emptyset$.

For any $S \subseteq \Theta$ let $\tilde{Q}_{R|S}$ be the restriction of \tilde{Q}_R to S . Say that $R \rightrightarrows \tilde{Q}_R$ is *almost uniformly continuous* at R if $\forall \varepsilon > 0 \exists S \subseteq \Theta$ with $\lambda(S) > 1 - \varepsilon$ such that $R \rightrightarrows \tilde{Q}_{R|S}(\theta|M)$ is uniformly continuous at R (where S replaces Θ in the definition of uniform continuity).

Definition 17. A decision rule M is **(weakly) robust to structural uncertainty** if $R \rightrightarrows \tilde{Q}_R$ is (almost) uniformly continuous at R .

The interpretation of this definition is that the decision rule should induce almost the same joint distribution of states and actions for small perturbations to the market clearing function. This in turn implies that the principal's expected payoff will be continuous in the function R . It turns out that continuous decision rules that are robust to multiplicity are robust to structural uncertainty.

Theorem 3. If $M \in \mathcal{M}$ is (weakly) robust to multiplicity then it is (weakly) robust to structural uncertainty.

PROOF. Proof in Appendix A.6.13.1

□

The important implication of Theorem 3 is that small changes in R lead to small changes in the principal's expected payoff. Formally, for any selection from $\theta \mapsto \tilde{Q}_R(\theta|M)$, i.e. any function $Q : \Theta \rightarrow \mathcal{A}$, such that $Q(\theta) \in \{\tilde{Q}_R(\theta|M)\}$ for all θ , abuse notation and write $Q \in \tilde{Q}_R(\cdot, M)$. Let the principal's expected payoff for a Q be given by

$$U(Q) = \int_{\Theta} u(\theta, Q(\theta)) dH(\theta)$$

where $u : \Theta \times \mathcal{A} \rightarrow \mathbb{R}$ is continuous and H is absolutely continuous with respect to Lebesgue measure.²⁸ Let $\mathcal{U}(R) = \{v \in \mathbb{R} : \exists Q \in \tilde{Q}_R(\cdot, M) \text{ with } U(Q) = v\}$ be the set of payoffs consistent with equilibria induced by M , given market clearing function R .

Proposition 15. *Let $M \in \mathcal{M}$ be a decision rule that is weakly robust to multiplicity for market clearing function R . Then $\mathcal{U}(R)$ is upper and lower hemicontinuous at R on \mathcal{C} .*

PROOF. in the Appendix A.6.13.2. □

The same conclusion holds *a fortiori* if M is robust to multiplicity.

If M is robust to multiplicity but has discontinuities on \bar{P}_M then it will not be robust to structural uncertainty under robustness to multiplicity. This is true when the discontinuity is not essential, i.e. when the left and right limits of M exist.²⁹ As discussed in Section 2.2, this further motivates the restriction to continuous decision rules.

Let $\theta_M(p|R) = \{\theta \in \Theta : R(M(p), \theta) = p\}$ be the set of states at which p could be an

²⁸Alternatively, we could dispense with absolute continuity and define robustness to multiplicity in terms of H .

²⁹Given that \mathcal{A} is compact, an essential discontinuity is one in which, roughly speaking, M oscillates with vanishing wavelength. The only potential benefit to the principal of using a discontinuous M is to avoid multiplicity, but an essential discontinuity does not help in this regard.

equilibrium price under M and R , and let $\bar{P}_M(R) := \{p \in \mathcal{P} : \theta_M(p) \neq \emptyset\}$ be the set of prices that could arise in equilibrium.

Lemma 9. *Assume that M satisfies robustness to multiplicity. If M has a non-essential discontinuity on $\bar{P}_M(R)$ then it is not robust to structural uncertainty.*

PROOF. Proof in Appendix A.6.13.3. □

This result can be generalized to other types of discontinuities, and to M that don't satisfy robustness to multiplicity.

2.5.2. Beyond uniqueness

When non-fundamental volatility is not a primary concern the principal may be willing to tolerate multiplicity. We maintain the assumption that the principal is not able to select freely from the set of equilibria. However, the principal may be willing to tolerate multiplicity, provided all possible equilibria are good from their perspective. The following two propositions are extremely useful when relaxing the requirement of unique implementation. They allow us to use the previous characterizations to study problems in which the strict uniqueness requirement is not imposed.

Proposition 16. *Assume R is weakly increasing in θ . If $M \in \mathcal{M}$ induces multiple equilibria then at least one has a monotone price function.*

PROOF. In Appendix A.6.10.1 □

Proposition 17. *Assume R is weakly increasing in θ and the environment is fully properly bridgeable. If $M \in \mathcal{M}$ induces multiple equilibria then at least one is characterized by (Q, P) that are CWUI.*

PROOF. In Appendix A.6.10.2 □

One implication Proposition 17 is that if the principal takes a strict worst case view of multiplicity then it is without loss of optimality to restrict attention to CWUI outcomes. That is, if the principal evaluates a decision rule M according to the worst equilibrium that it induces, then the principal may as well restrict attention to M that are weakly robust to multiplicity.

Proposition 17 also simplifies the problem of a principal who takes a less extreme approach to multiplicity than the strict worst-case preferences described above. Consider a principal who takes a lexicographic approach to multiple equilibria: the principal first evaluates a decision rule according to the worst equilibrium that it induces. Among those decision rules with the same worst-case equilibrium payoff, the principal chooses based on the best equilibrium that each induces (or indeed some other function of the remaining equilibria).³⁰ By Proposition 17 we know that the highest worst-case guarantee is exactly the maximum payoff over the subset of decision rules in \mathcal{M} that are weakly robust to multiplicity. Once this value has been determined, the goal of the principal is to choose the decision rule with the best equilibrium outcome, subject to not inducing any equilibrium with a payoff below this worst-case bound.

³⁰Such preferences are similar in spirit to these studied in the context of robust mechanism design (Börgers, 2017) and information design (Dworczak and Pavan, 2020).

Assume first that the principal's payoffs do not depend directly on the price; the principal cares only about the joint distribution of states and actions (a similar discussion will apply to other preferences). Assume that there is a unique optimal CWUI action function Q^* (similar discussion applies to virtual implementation), implemented uniquely by decision rule M^* (if there are multiple optimal CWUI action functions then Condition 1 in Proposition 18 below must hold for one of them). If this is the case then, by Proposition 17, the principal needs to choose a decision rule that implements Q^* as one of its equilibrium outcomes; if Q^* is not one of the equilibrium outcomes then there will be some other CWUI action function induced by the decision rule, which will be worse than Q^* by assumption. This pins down the decision rule for all prices in the range $\{R(Q^*(\theta), \theta) : \theta \in \Theta\}$; any optimal decision rule must coincide with M^* for such prices. Moreover, Q^* will be an equilibrium outcome of any such decision rule. This discussion implies the following.

Proposition 18. *Let Q^* be the set of optimal CWUI action functions. Then the constraints of the principal with lexicographic multiplicity preferences can be stated as follows: choose \hat{M} subject to*

- (1) $\exists Q \in Q^*$ such that $\hat{M}(R(Q(\theta), \theta)) = Q(\theta)$ for all $\theta \in \Theta$,
- (2) $\hat{M} \in \mathcal{M}$.

As will be illustrated in the application of Section 2.6.1, these constraints can greatly simplify the problem of finding optimal policies for a principal with lexicographic preferences over multiple equilibria.

2.6. Applications

2.6.1. Bailouts

We analyse here in greater detail an application similar to the emergency lending example discussed in the introduction. The purpose is to expand on the type of results discussed in the introduction. The government is considering a bailout for a publicly traded company, which it considers strategically important.³¹ The government chooses a level of support $a \in \mathcal{A} = [0, \bar{a}]$. The company's business prospects $\theta \in \Theta$, representing the demand environment, competition, future costs, etc., are unknown. Higher states represent better prospects; for each level of support the share price is a strictly increasing function of the state. We make two additional assumptions regarding the share price.³²

- (1) The slope of $\theta \mapsto R(a, \theta)$ is decreasing in a .
- (2) There exists a state θ^* such that $a \mapsto R(a, \theta)$ is strictly increasing for $\theta < \theta^*$ and strictly decreasing for $\theta > \theta^*$.

The first assumption represents the belief on the part of investors that government involvement in the firm will reduce upside when business prospects are good. This could be because the bailout involves the government taking a role in management, for example by gaining seats on the board, or carries negative stigma (Che et al., 2018). An alternative interpretation is that the bailout takes the form of forgivable loans. As a result of ex-post loan forgiveness, the effective amount owed is increasing in the state (which will be revealed *ex post*). The second assumption captures the fact that when business prospects are sufficiently bad, the bailout is necessary to sustain the operations of the

³¹Alternatively, the bailout could be for an entire industry, in which many of the firms are publicly traded.

³²These assumptions can be directly related to the asset dividends, as discussed in Appendix A.8.

business. When business prospects are sufficiently good however, the adverse effects of government intervention dominate. These features are derived from the discussion around recent bailouts, for example that of Lufthansa by the German government.³³

The government does not wish to give any support to the company if the state is below some threshold θ' . In such cases the business is not considered viable, and the government prefers to let it fail. On the other hand, if the state is above some threshold $\theta'' > \theta'$, the government would also like to offer no support. In this case the government believes that the business can survive without intervention. The government's payoff $u(a, \theta)$ is therefore decreasing in a for $\theta \notin (\theta', \theta'')$. The government would like to intervene when the state is in $[\theta', \theta'']$. In these states the government's payoff $u(a, \theta)$ is increasing in a . The government maximizes expected utility, and has an absolutely continuous prior H .

Figure 2.5 illustrates the situation in which $\theta^* \in [\theta', \theta'']$. The blue lines correspond to the price function P^* induced by the first-best action function Q^* . Assumption 2 on R (above) implies that the environment is continuously fully bridgeable and correctable. Since the price function induced by Q^* is strictly increasing, the first-best is CWUI by Corollary 4.

If $\theta^* < \theta''$ the government is relatively interventionist. In this case the government would like to intervene even in states in which investors would prefer no bailout. This will be the case when the strategic importance of the company is high, for example when the

³³In the Lufthansa case, one large shareholder, Heinz Hermann Thiele, threatened to veto the proposed bailout, which involved the government taking a 20% stake in the company and receiving seats on the board. Thiele was reportedly concerned that the government stake would make it harder to restructure and cut jobs. On the other hand, supervisory board chairman Karl-Ludwig Kley emphasised Lufthansa's dire prospects: "We don't have any cash left. Without support, we are threatened with insolvency in the coming days." Lufthansa shares rose 20% when Thiele announced that he would support the deal (Wissenbach and Taylor, 2020).

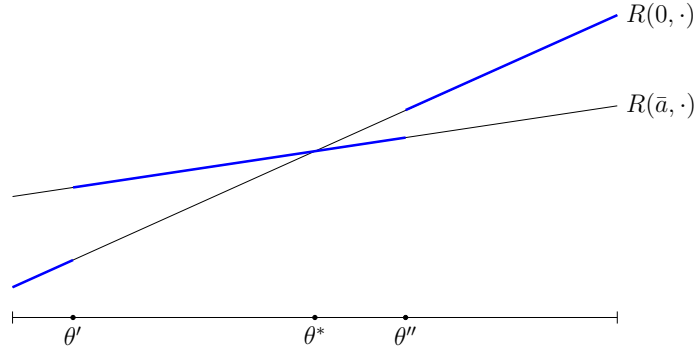


Figure 2.5. First-best is CWUI

company is involved in national security, employs a large number of workers, or engages in production which has large technological spillovers.

Although the first-best is CWUI when $\theta^* < \theta''$, the government must take care in choosing the appropriate implementing decision rule, so as to avoid multiplicity. There are a continuum of decision rules that implement the first-best as an equilibrium outcome. The decision rule for prices in $P^*(\Theta)$ is clearly determined by the desired action function. However the action function alone does not pin down the decision rule for prices in $\tilde{P} \setminus P^*(\Theta)$. Consider the prices in the range $(R(0, \theta'), R(\bar{a}, \theta'))$. For such prices, M must satisfy $p = R(M(p), \theta')$. If the government responds too much to price changes in this range, meaning that M increases faster than what this condition implies, then there will be equilibria in which action $a > 0$ is taken for states below θ' . Similarly if the government under-responds then there will be equilibria in which action $a < \bar{a}$ is taken for states above θ' . Similar restrictions apply to the discontinuity in P^* at θ'' .

Suppose instead that $\theta^* > \theta''$. In this case the government is *lassiez faire*; it does not wish to intervene in states (θ'', θ^*) in which investors would welcome a bailout. The price function associated with the first-best outcome is depicted in Figure 2.6. In this

case the price is non-monotone, and is therefore not CWUI. In fact, in this case it is not even implementable, as it violates measurability. The optimal CWUI outcome is found by flattening the price function to eliminate non-monotonicity.

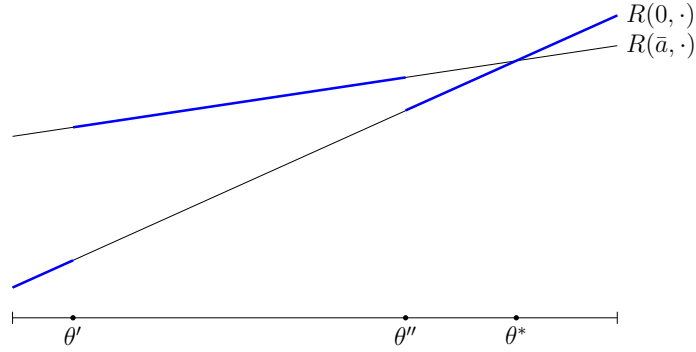


Figure 2.6. First-best not implementable

The price function for the virtually optimal decision rule is pictured in Figure 2.7 (it will only be virtually optimal since the price must be strictly increasing, but can have an arbitrarily small slope). It is fully characterized by a state $\hat{\theta}$ at which the flattening begins. For any $\hat{\theta} \in [\theta', \theta'']$ the government's payoff is given by

$$\int_{\underline{\theta}}^{\theta'} u(0, \theta) dH(\theta) + \int_{\theta'}^{\hat{\theta}} u(\bar{a}, \theta) dH(\theta) + \int_{\hat{\theta}}^{t(\hat{\theta})} u(\alpha(\theta, \hat{\theta}), \theta) dH(\theta) + \int_{t(\hat{\theta})}^{\bar{\theta}} u(0, \theta) dH(\theta),$$

where $\alpha(\theta, \hat{\theta})$ is defined by $R(\alpha(\theta, \hat{\theta}), \theta) = R(\bar{a}, \hat{\theta})$, and $t(\hat{\theta})$ by $R(0, t(\hat{\theta})) = R(\bar{a}, \hat{\theta})$. Here $\alpha(\theta, \hat{\theta})$ is decreasing in its first argument and increasing in the second, and $t(\hat{\theta})$ is decreasing. Assuming R and u are differentiable, the optimal $\hat{\theta}$ can be identified via the first order condition.

The virtually optimal can involve two types of loss for the government: under-support for the company for states below θ'' , or over-support for states above θ'' . In fact, the optimal policy will entail both types of loss. To see this, suppose $\hat{\theta} = \theta''$. There is a

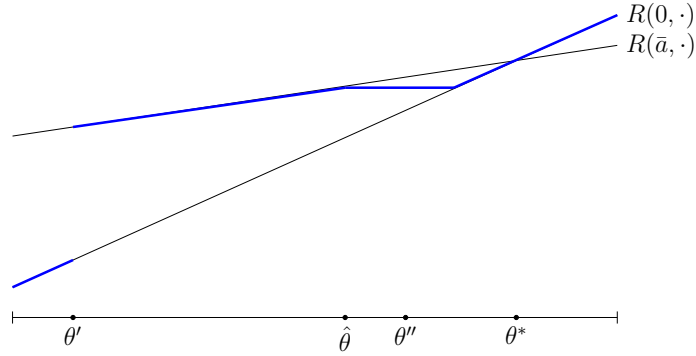


Figure 2.7. virtually optimal decision rule

first-order gain from lowering $\hat{\theta}$ by a small ε , since this means that less support needs to be offered on the entire interval $(\theta'', t(\theta''))$. The loss, which results from less support being offered on $(\theta'' - \varepsilon, \theta'')$, is second order. An analogous argument applies to raising $\hat{\theta}$ when $t(\hat{\theta}) = \theta''$.

Suppose that in the case of $\theta^* > \theta''$ the government is willing to tolerate some multiplicity, and takes the lexicographic approach described in Section 2.5.2. The question is whether or not the government can improve their upside while still guaranteeing the payoff given by the virtually optimal decision rule. Assume for simplicity that there is a unique $\hat{\theta}$ that defines the virtually optimal price function (if there are multiple such $\hat{\theta}$ the same analysis applies to any selection). Then, as shown in Section 2.5.2, the virtually optimal decision rule is pinned down on $P^*(\Theta)$. The only potential changes that could be made to the decision rule when allowing for multiplicity are on $(R(0, \theta'), R(\bar{a}, \theta'))$. It is easy to see from Figure 2.7 however, that changing the decision rule on this range will can only induce equilibria in which lower actions are taken on (θ', θ'') or higher actions are taken on $[\underline{\theta}, \theta')$. Neither of these modifications benefits the principal. Thus relaxing the unique implementation requirement does not change the virtually optimal policy.

2.6.2. Moving against the market

In this section we explore the distinctive features of a set of applications in which the principal would like to induce a *decreasing* price. As before, $\theta \mapsto R(a, \theta)$ is increasing. These are therefore situations in which the principal is working to move prices against the market. The following are two such applications.

Monetary policy in a crisis

During the financial crisis of 2008 and the ongoing Covid-19 recession, central banks have moved aggressively to lower interest rates. In this application the unknown state is the severity of the liquidity crisis faced by firms, and the market price is the interest rate. The action is the size of asset purchases through open market operations. The central bank's objective is to implement an interest rate that is decreasing in the state via their open market operations.

Grain reserves

Many developing countries manage grain reserves as a tool for stabilizing the grain price and responding to food shortages. The state here is the size of a demand or supply shock, the price is the grain price, and the action is the size of grain purchases/sales. Depending on the nature of the crisis and the structure of the grain market, the government may wish to implement a decreasing price. If the government has limited capacity to make direct transfers to households it may wish to implement transfers by lowering the grain price when there is a severe crisis. For example, suppose that grain is a Giffen good. If there is an employment crisis outside of agriculture the price of grain may rise,

absent government intervention.³⁴ In this case the government may wish to subsidize non-agricultural households by lowering the grain price.

Throughout this section, we maintain the assumptions that $\mathcal{A} = [\underline{a}, \bar{a}] \in \mathbb{R}$ and that $\theta \mapsto R(a, \theta)$ is strictly increasing for all a , and that $a \mapsto R(a, \theta)$ is strictly decreasing for all θ (that this function is decreasing as opposed to increasing is simply a normalization). A decreasing price function is possible if and only if $R(\underline{a}, \underline{\theta}) > R(\bar{a}, \bar{\theta})$. Figure 2.8 depicts such an environment.

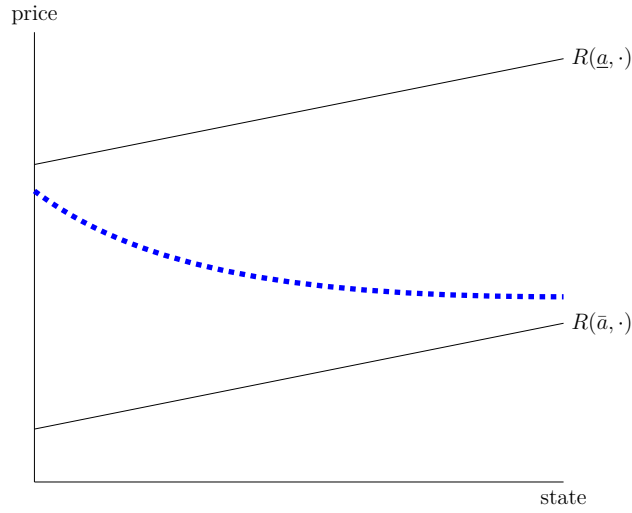


Figure 2.8. Decreasing price function

The following observation shows that implementing an increasing price function in this setting is easy.

Lemma 10. *If $a \mapsto R(a, \theta)$ is strictly decreasing for all θ then any strictly increasing $M \in \mathcal{M}$ induces an increasing and continuous price function as the unique equilibrium.*

PROOF. Proof in Appendix A.6.11. □

³⁴There is empirical evidence that food staples are Giffen goods for extremely poor households (Jensen and Miller, 2008).

An equilibrium exists for any increasing M by Tarski's fixed point theorem (in this simple setting a more direct proof can be given). That the price function will be increasing follows from the fact that $a \mapsto R(a, \theta)$ is decreasing and $\theta \mapsto R(a, \theta)$ is increasing. If P is increasing and M is increasing, there will be no equilibrium involving prices above $P(\bar{\theta})$ or below $P(\underline{\theta})$. Moreover, we show that M cannot have a discontinuity on $[P(\underline{\theta}), P(\bar{\theta})]$, which implies that P is continuous.

Decreasing price functions are more interesting in this setting. Non-monotonicity of M will be necessary to robustly implement a decreasing price.

Lemma 11. *Assume $a \mapsto R(a, \theta)$ is strictly decreasing for all θ , and let P be a decreasing price function. If $M \in \mathcal{M}$ uniquely implements P then*

- (1) $M(p)$ is decreasing and continuous on an open interval containing $(P(\bar{\theta}), P(\underline{\theta}))$,
- (2) M has discontinuities in $(P(\underline{\theta}), R(\underline{a}, \underline{\theta})]$ and $(R(\bar{a}, \bar{\theta}), P(\bar{\theta})]$.
- (3) There exist $p'' > p' > P(\underline{\theta})$ such that $M(p'') > M(p')$.
- (4) There exist $p' < p'' < P(\bar{\theta})$ such that $M(p'') > M(p')$.

PROOF. Proof in Appendix A.6.12. □

Lemma 11 shows that discontinuous and non-monotone M will be necessary to implement a decreasing price. The intuition comes from the fact that there are only two ways to guarantee that $\theta_M(p) = \emptyset$, i.e. that there are no equilibria with price p . Either M must specify an action that is too high, meaning $R(M(p), \bar{\theta}) < p$, or too low, so that $R(M(p), \underline{\theta}) > p$. If neither of these hold then there will be some θ such that $R(M(p), \theta) = p$, by continuity of R . The only way to ensure that there are no equilibria with prices in $[R(\underline{a}, \underline{\theta}), R(\underline{a}, \bar{\theta})]$ is to take a high enough action for such prices; it must

be that $R(M(p), \bar{\theta}) < p$ for all such prices. At the same time, M must be decreasing on $(P(\bar{\theta}), P(\underline{\theta}))$ in order to implement a decreasing P . This tension is what necessitates discontinuities and non-monotonicities in M .

Lemma 11 is important in applications because it highlights the danger of artificially restricting the class of permissible decision rules. If, for example, one restricts attention to monotone decision rules, it will not be possible to uniquely implement a decreasing price. It is nonetheless common practice in the literature to focus on monotone, or even linear, decision rules (see for example Bernanke and Woodford (1997)). Most papers which make this type of linearity assumption do so in models where the action space is unbounded. The fact that the action space is bounded here is an important driver of the non-monotonicity result in Lemma 11. However in reality there are often bounds on available set of actions. In the grain reserves example, the government cannot sell more grain than it has in reserve. Similarly, central banks in developing countries cannot make unlimited asset purchases without creating significant balance sheet risks (Crowley, 2015). Lemma 11 shows that such restrictions on the feasible actions can interact in surprising ways with conditions on the decision rules used to gain tractability. Our general framework allows us to avoid the need to impose such conditions.

CHAPTER 3

Liability Design with Information Acquisition**(Joint work with Bruno Strulovici)****3.1. Introduction**

In 2019, a California court sentenced paint maker Sherwin-Williams to pay hundreds of millions of dollars to address the dangers caused by lead paint. The sentence was remarkable because even though lead paint became banned in 1978, the suit concerned damage caused during the decades *before* the ban and centered on the accusation that paint makers were aware of the dangers caused by lead paint long before the ban was passed.

In essence, the court's argument was that Sherwin-Williams and other paint makers knew or should have known the dangers caused by lead paint.

While it is difficult for a regulator to guess a firm's private information, it is perhaps easier to assess due diligence: did paint makers research the risks of lead paint sufficiently well before marketing it?

Formally, the problem is not just one of private information, but also one of information acquisition: how can a regulator make sure that agents learn sufficiently well before taking actions?

One may model this question as a delegated Wald problem (Wald (1945)): the principal is a regulator who relies on an agent (the firm) to acquire information before deciding between launching a product and abandoning it.

If the regulator could unrestrictedly penalize a firm, she could force the firm to internalize any damage caused by the product and implement the socially-optimal level of information acquisition.

For various reasons, liability may be capped, however, which precludes the full transfer of damages to the firm. Moreover, the regulator may punish the firm only if some damage occurs, and choose a penalty that depends only on the information available to the regulator after the damage has occurred.

We analyze this problem in a Brownian version of the Wald Problem: the firm observes an arithmetic Brownian whose drift depends on the state of the world, i.e., on the riskiness of the product. Information acquisition is costly. The first-best policy is to acquire information until the riskiness of the product becomes sufficiently clear, launch the product if this riskiness is low, and abandon it if the riskiness is high.

We characterize all incentive-compatible liability rules when (i) the firm has initial private information, (ii) liability is capped, and the (iii) regulator can penalize the firm only when damage occurs. In general, the regulator may wish to propose at the outset a menu of contracts to the firm in order to extract some of the firm's private information. Indeed, this is the approach suggested by the Revelation Principle. In the present context, however, this approach may be difficult to implement, because it requires that the firm contracts with the regulator long before launching the product and, in fact, even before knowing whether it wishes to launch the product.

Fortunately, our first main result is that it is without loss generality for the regulator to focus on *tariff mechanisms*, which are mechanisms for which the firm does not report its private information and only pays a penalty if damage occurs. This result may be viewed as a Taxation Principle for situations in which transfers take place only after some contingencies (damage occurs), but not others, and builds on our companion paper (Poggi and Strulovici (2020)), which provides a general Taxation Principle with Non-Contractible Events.

With a tariff mechanism, a firm's decision to launch the product depends on its prior information, which affects the probability that the product causes damage. Our second main result is that any incentive-compatible tariff mechanism has the following property: firms whose initial private information assigns a higher probability of damage always acquire more evidence before launching their product. This monotonicity property is not an immediate consequence of incentive compatibility, and would in fact be violated if the regulator could impose evidence-based transfers to the firm regardless of whether a damage occurred.

Our third main result is to show that any launch thresholds that induce the firm to perform more due diligence than it would under a fixed penalty can be implemented by a monotonic tariff, i.e., a tariff whose penalty is decreasing in the strength of evidence acquired by the firm before launching the product.

We also show that for a general specification of the regulator's objective function, setting the tariff at its uniform ceiling induces too little due diligence compared to the social optimum, even when the social benefit from launching the product exceeds the firm's profit from doing so. This result holds under a cost-benefit ratio condition, which

stipulates that the social benefit from the product relative to the harm it may cause is smaller than the firm's profit relative to the maximum liability that it may face.

3.1.1. Literature

Liability is one way to address the inefficiencies that arise due to externalities derived from hazardous activities.

Shavell (1984) contrasts liability with regulation and argues that one of the advantages of liability is that it is potentially much cheaper: the administrative costs are incurred only if some harm occurs, and this may have a small probability. Kolstad et al. (1990) and Schmitz (2000) show that, in settings with heterogeneous agents, a combination of regulation and liability achieves sometimes a more efficient outcome than liability alone.

This paper contributes to the study of liability in the context of information acquisition. Shavell (1992) studies how different liability rules affect the incentives to acquire information about risks. By contrast, in our paper we allow the designer to commit to any liability rule, that is how much the agent must pay in case of harm as a function of what the principal observes. Moreover, we assume that in the case that a damage occurs, the planner observes all the evidence that the firm acquired — but not *how* this evidence was acquired.

Friehe and Schulte (2017) study information acquisition when firms require the approval of a regulator to launch a potentially risky product. In their model, the regulator cannot commit to an approval mechanism. Instead, the paper focuses on how the equilibrium outcome changes with different liability rules.

Finally, Henry et al. (2021) study the trade-off between acquiring information about risks of an activity before or when the activity is taking place. In their model, the planner commits to a liability rate per unit of time that the product is on the market. In this paper we allow the planner to commit to any liability rule.

The problem in this paper can also be seen as a contract problem with moral hazard and adverse selection. The agent (firm) has quasilinear preferences in the money, but transfers are only allowed if there is some harm.

3.2. Model

A firm must decide between launching a product and abandoning its development. If launched, the product may cause damage with positive probability. The firm has some private information about the product's riskiness and can acquire additional information ("due diligence"), before making a final decision.

A regulator wishes to encourage the launch of low-risk products and deter the launch of high-risk ones, as well as to encourage the firm to acquire sufficient information before making its decision.

The regulator faces two constraints. First, the firm has limited liability: the social cost caused by product damage is $L > 0$ and the firm's liability is capped at some lower level $l < L$. Second, the regulator can penalize the firm only if damage occurs. In particular, it cannot penalize firms that acquired too little information and took an overly risky decision unless such risk results in actual damage.

The timing of the game is as follows:

1. The firm is endowed with a prior $\theta \in \Theta \subset [0, 1]$ about the product's riskiness $y \in \{0, 1\}$, with $\theta = \Pr(y = 1)$.

2. The firm can acquire additional information about y according to a dynamic technology to be described shortly.
3. The firm decides between launching and abandoning the product.
4. If the firm launches the product, it causes some damage if the product was risky ($y = 1$) and doesn't if the product was safe ($y = 0$).
5. In case of damage, the firm pays a penalty $\psi \leq l$ set by the regulator.

The assumption that a risky product causes damage with probability 1 is without loss of generality: if this probability were less than 1, the same analysis would apply using expected damage and expected penalties.

Information structure: During the information-acquisition stage, the firms observes a process X given by

$$X_t = (-1 + 2y)t + \sigma B_t$$

where B is the standard Brownian motion. The drift of X depend symmetrically on the product's riskiness y : the drift is $+1$ if the product causes damage and -1 if it does not. Therefore, observing X gradually reveals y . This revelation is progressive due to the stochastic component of X .

The firm stops acquiring information at some time τ that is adapted to the filtration of X .

The regulator observes nothing about X except if some damage occurs, in which case she observes the last value X_τ taken by the process at the time of the firm's decision. X_τ is a measure of the firm's due diligence to assess the product's riskiness before launching it: in this Brownian model, it is well-known (though not immediate) that for each $t > 0$, the variable X_t is a sufficient statistic for the information about y contained by the entire

path $\{X_s\}_{s \leq t}$ of the process X until time t . Mathematically, the likelihood ratio of y associated with a path of X from time 0 to t is only a function of X_t .

Because the stopping time τ is chosen endogenously by the firm, which has private information about y , X_τ is not a sufficient statistic for y once the firm's strategic timing is taken into account. Our assumption that the regulator observes X_τ instead of the entire path $\{X_t\}_{t \leq \tau}$ captures the idea that the regulator does not perfectly observe all the decisions made by the firm during the information acquisition stage. Intuitively, the regulator observes the most informative signal about y contained by the path of X that is independent of the firm's private information.

Payoffs: The firm incurs a running cost c from acquiring information, and a profit π if it launches the product. Let $d = 1$ if the firm launches the product and $d = 0$ if it abandons it, and τ denote the time spent acquiring information. The firm's realized payoff is

$$d(\pi - y\psi) - c\tau$$

where π is the firm's profit from the launch in the absence of damage. The regulator's objective internalizes the entire damage caused by the product:

$$d(\beta - yL) - c\tau$$

where β is the social benefit from the launch in the absence of damage.

Throughout the paper, we make the following assumption:

Assumption 2 (Ordered Cost-Benefit Ratios). $l/\pi < L/\beta$.

This assumption captures the idea that the risk of damage is more severe for the regulator relative to the benefit of launching the product than it is for the firm. The assumption allows the social benefit from launching the product to exceed the firm's profit (i.e., $\beta > \pi$).

3.3. Preliminary Analysis: Symmetric Information

First Best: If the regulator knew the firm's type θ and could dictate the firm's strategy, the optimal strategy would consist in launching the product if the process X drops below some lower threshold x_θ^* and abandoning it if X exceeds some upper threshold $\bar{x}_\theta^* \geq x_\theta^*$.

Tariffs: A tariff is a function $\psi : \mathbb{R} \rightarrow \mathbb{R}$ mapping evidence x to a penalty $\psi(x) \leq l$.¹ Given a tariff ψ , a firm with prior θ chooses a stopping time τ and a launch/abandonment decision $d \in \{0, 1\}$ to maximize its expected utility

$$(3.1) \quad E[d(\pi - y\psi(X_\tau)) - c\tau \mid \theta].$$

It is straightforward to check that the solution to this problem consists of cutoffs $\underline{x}_\theta^\psi < \bar{x}_\theta^\psi$ such that the firm acquires information until X reaches either of the cutoffs.

Limited liability affects incentives in two ways. First, since the firm does not fully internalize damages, it is willing to take riskier decisions than is socially optimal for a given belief about the product's safety. Second, the value of information is different. For example, if the tariff is $\psi \equiv 0$, the firm has no incentive to acquire any information and always launches its product immediately.

¹We allow negative tariffs, which amount to a subsidy for the firm and may be used to reward firms that performed unusually careful inspections before launching their products.

To appreciate the consequences of limited liability, suppose that the regulator sets the tariff uniformly equal to the allowed maximum: $\psi(x) \equiv l$ for all $x \in \mathbb{R}$. In this case, the firm launches the product if X drops below some cutoff \underline{x}_θ^l and abandons it if X reaches some upper cutoff \bar{x}_θ^l .

This maximum penalty may motivate the firm to perform due diligence before launching the product, but the amount of due diligence is always strictly suboptimal, as the next result shows.

Proposition 19 (recklessness). *$\underline{x}_\theta^* < \underline{x}_\theta^l$ for all $\theta \in \Theta$.*

PROOF. We fix some prior $\theta \in \Theta$ throughout the proof and let x^* and x^l denote the socially-optimal and firm-optimal launch thresholds, respectively, when $\psi \equiv l$, given prior θ .

Given a current evidence level x , the firm's expected payoff if it launches the product at x is:

$$u(x) = \pi - p(x)l$$

where $p(x) = \Pr(y = 1|x, \theta)$. The regulator's expected payoff if the firm stops at x is:

$$v(x) = \beta - p(x)L.$$

Assumption 2 implies that

$$(3.2) \quad v(x) = \frac{L}{l}(u(x) - k)$$

where $k = \pi - \beta l/L > 0$.

Thus, the “launch-payoff functions” faced by the regulator and the firm are related by equation (3.2), and both parties face a running cost c before launching or abandoning the product and a payoff normalized to zero if the product is abandoned. Proposition 19 then follows from two observations:

Observation 1: Consider two launch-payoff functions \hat{u}, u . If $\hat{u} = \alpha u$ with $\alpha > 1$, then the optimal launch threshold for \hat{u} is lower than the optimal launch threshold for u .

Observation 2: Consider two launch-payoff functions \hat{u}, u . If $\hat{u} = u - \hat{k}$ with $\hat{k} > 0$, the optimal launch threshold for \hat{u} is lower than the optimal launch threshold for u .

Once we justify these observations, Proposition 19 follows from (3.2) by applying Observation 2 to $u - k$ and u and Observation 1 to $v = L/l(u - k)$ and $u - k$, using the fact that $L/l > 1$.

To prove Observation 1, notice that if $\hat{u} = \alpha u$ with $\alpha > 1$, the dynamic optimization problem with launch payoff \hat{u} and running cost c is equivalent to the problem with launch payoff u and running cost $\hat{c} = c/\alpha < c$, since the problems become identical up to the scaling factor α . With a lower running cost \hat{c} , the continuation interval $(\underline{x}(\hat{u}), \bar{x}(\hat{u}))$ contains the continuation interval $(\underline{x}(u), \bar{x}(u))$ with running cost c . In particular, the launch thresholds are ranked: $\underline{x}(\hat{u}) \leq \underline{x}(u)$.

To prove Observation 2, consider the optimal continuation interval (x^l, \bar{x}) when the launch-payoff function is u and let $\tau = \inf\{t : X_t \notin (x^l, \bar{x})\}$. Fixing any $x \in (x^l, \bar{x})$, acquiring information is optimal when starting at x , which means that

$$(3.3) \quad u(x) \leq f(x)u(x^l) - cE_x[\tau]$$

where $f(x)$ is the probability that $X_\tau = x^l$ (as opposed to \bar{x}) and $E_x[\tau]$ is the expected value of τ when the process X starts at x . For the launch-payoff function $\hat{u} = u - k$ with $k > 0$, (3.3) implies that

$$\hat{u}(x) < f(x)\hat{u}(x^l) - cE_x[\tau].$$

This shows that stopping at x to launch the product is strictly dominated by the strategy that consists in launching the product if X reaches x^l and abandoning it if X reaches \bar{x} . This implies that the optimal launch threshold with \hat{u} is lower than x^l and proves Observation 2. ■ □

Intuitively, Proposition 19 captures the idea that the regulator values more than the firm having a safer product conditional on launch. Remarkably, however, this result holds even when the social benefit from launching the product exceeds the firm's profit from doing so.

Although the uniform tariff $\psi \equiv l$ brings the firm closest to fully internalizing the damage that its product might cause, the regulator might choose a different tariff, for example, to reward the firm if it acquired more information. The next section studies the firms' incentives in more details.

3.4. Incentive Compatibility

Suppose that the regulator can contract with the firm after the firm has received its initial private information and before it takes any action, and that the regulator has full commitment power.

Definition 18. *A direct liability mechanism is a menu $M = (\{\tau_\theta, d_\theta, \psi_\theta\}_{\theta \in \Theta})$ such that for all $\theta \in \Theta$:*

- (i) *The stopping time τ_θ is measurable with respect to the filtration $\{\mathcal{F}_t^X\}_{t \geq 0}$ generated by X ;*
- (ii) *The decision d_θ is measurable with respect to the information at time τ , i.e., to the σ -algebra $\mathcal{F}_{\tau_\theta}^X$;*
- (iii) *The tariff $\psi_\theta : \mathbb{R} \rightarrow \mathbb{R}$ is uniformly bounded above by l .*

Since the regulator has full commitment power, the Revelation Principle guarantees that it is without loss of generality to focus on direct liability mechanisms.

Given a direct liability mechanism, the firm chooses an item $f_{\hat{\theta}} = (\tau_{\hat{\theta}}, d_{\hat{\theta}}, \psi_{\hat{\theta}})$ from the menu. Faced with the tariff $\psi = \psi_{\hat{\theta}}$, the firm chooses a stopping time and a decision to maximizes its expected utility as given by (3.1).

Definition 19. *A direct liability mechanism M is incentive compatible if for each $\theta \in \Theta$ it is optimal to chooses the item f_θ from M and the strategy (τ_θ, d_θ) .*

In general, a direct liability mechanism may implement absurd policies: for example, the firm could get a very high reward (i.e., a negative penalty) if it launches the product when X_t is very high (and, hence, the product is very risky). We rule out such a possibility and focus on *admissible* mechanisms:

Definition 20. *An IC direct liability mechanism is admissible if each type θ 's strategy is characterized by thresholds $\underline{x}_\theta \leq \bar{x}_\theta$ such that θ launches the product if X_t drops below \underline{x}_θ and abandons it if X_t exceeds \bar{x}_θ .*

In practice, it may be difficult for a regulator to contract with the firm at the outset and agree on penalties that depend finely on a firm's private information before it launches

a product and, even earlier, before the firm decides how much due diligence to perform before deciding whether to launch its product. It is therefore valuable to determine when a direct liability mechanism can be implemented by a tariff that is independent of the firm's private information.

Definition 21. *A direct liability mechanism is a tariff mechanism if the tariffs $\{\psi_\theta\}_{\theta \in \Theta}$ are independent of θ .*

Theorem 4. *Any admissible direct liability mechanism is outcome-equivalent to a tariff mechanism.*

PROOF. Consider any admissible mechanism M and let $\underline{x}_\theta = \underline{x}_\theta^{\psi_\theta}$ and $\psi_\theta = \psi_\theta(\underline{x}_\theta)$ denote the firm's launch threshold and penalty in case of damage that are implemented under mechanism M when the firm has type θ .

We introduce a ceiling mechanism \tilde{M} as follows: for each θ , $\tilde{\psi}_\theta$ gives the maximal penalty l for all x except at \underline{x}_θ , where it gives ψ_θ . The ceiling mechanism \tilde{M} is IC and implements the same thresholds \underline{x}_θ , because under M the firm faces the penalty only when it launches the product and higher penalties at other levels can only reduce the incentive to deviate.

If M prescribes the same threshold \underline{x} to types $\theta \neq \theta'$, the penalties ψ_θ and $\psi_{\theta'}$ must be identical. Otherwise, one type would want to misreport its type and M would not be incentive compatible.

We define the tariff ψ as follows:

$$\psi(\underline{x}_\theta) = \psi_\theta$$

for all $\theta \in \Theta$ and

$$\psi(x) = l$$

otherwise.

This tariff is independent of the firm's private information. Moreover, it implements the same launch thresholds as M , as is easily checked. ■ □

Theorem 4 shows that any admissible liability mechanism can be implemented by a tariff. From now on, we invoke Theorem 4 and focus without loss of generality on admissible mechanisms that are implemented by tariffs, hereafter “admissible tariffs”.

Given any admissible tariff $\psi : x \mapsto \psi(x)$, each type θ faces a Markovian decision problem in which the state variable at time t is X_t . Therefore, there exist thresholds $\underline{x}_\theta^\psi \leq \bar{x}_\theta^\psi$ such that type θ stops acquiring information when the process X leaves the interval $(\underline{x}_\theta^\psi, \bar{x}_\theta^\psi)$, launches the product at $\underline{x}_\theta^\psi$ and abandons it at \bar{x}_θ^ψ .

Our next result establishes a single-crossing property for the firm.

Lemma 12. *Consider any admissible tariff ψ , level x , and type $\theta \in \Theta$. If θ prefers acquiring information at x to immediately launching the product at x , then so does any type $\theta' \geq \theta$.*

PROOF. We fix a tariff function ψ and a level x , and suppose that $X_t = x$ at some time t that we normalize to 0 for simplicity. Suppose that some type θ prefers the strategy that consists in launching the product at $\underline{x} < x$ and abandoning it at $\bar{x} > x$, and let $p = \Pr(y = 1|\theta)$.

If θ launches the product at x , it gets:

$$(3.4) \quad \pi - p\psi(x).$$

Let T^g, f^g denote the expected hitting time and the probability of hitting \underline{x} if $y = 0$ (the product is good), and T^b and f^b be defined similarly if $y = 1$ (the product is damaged). If θ continues until hitting \underline{x} or \bar{x} , its expected payoff is

$$(3.5) \quad p(f^b(\pi - \psi(\underline{x})) - cT^b) + (1 - p)(f^g \times \pi - cT^g).$$

Comparing (3.4) and (3.5), continuing is optimal if

$$(3.6) \quad p(f^b(\pi - \psi(\underline{x})) + \psi(x) - cT^b) + (1 - p)(f^g \pi - cT^g) \geq \pi.$$

The left-hand side is a convex combination of two terms: $a = f^b(\pi - \psi(\underline{x})) + \psi(x) - cT^b$ and $b = f^g \pi - cT^g$. The second term, b is less than π , because f^g is a probability. Therefore, (3.6) can hold only if the first term, a , is greater than π .

Rewriting (3.6), a firm that assigns probability p to $y = 1$ wishes to continue if

$$p(a - b) \geq \pi - b.$$

Since $a > b$, the coefficient of p is strictly positive. This implies that any type that assigns probability $p' > p$ to $y = 1$ also prefers the continuation strategy to launching the product immediately at x . ■ □

Lemma 12 has the following intuition: If a firm knew that the product were safe, it would optimally launch the product immediately. The return to acquiring more evidence is

negative in this case. Given any liability function, if a type wants to acquire more evidence it must be that doing so has a positive return conditional on the product being unsafe. The expected return from acquiring more evidence is thus increasing in the probability that the firm assigns to the product being faulty.

Lemma 12 immediately implies the following monotonicity result:

Proposition 20. *For any admissible tariff ψ , the launch thresholds $x^\psi(\theta)$ are decreasing in θ .*

This monotonicity result crucially hinges on the fact that the regulator can only charge the firm if it causes some damage. The following example² shows that if the regulator can charge the firm even when the product causes no damage, the launch thresholds increase with the type of the firm.

Example: Monotonicity Violation with Damage-Independent Fee

Suppose that the assumptions of our model are maintained with one exception: if the firm launches its product, the regulator charges the firm a fee $\eta(x) \leq l$ that depends on the evidence x demonstrated when the product is launched, independently of any damage subsequently caused by the product.

We assume that $l > \pi$, so that the regulator can deter the firm from launching the product at any given x by setting $\eta(x) = l$. An admissible direct revelation mechanism specifies, for each type, a launch threshold \underline{x}_θ and a fee $\eta_\theta = \eta(\underline{x}_\theta)$. Without loss of generality we assume that $\eta(x) = l$ for all $x \notin \{\underline{x}_\theta\}$.

²This example is partially inspired by the approval mechanisms in McClellan (2019) and Henry and Ottaviani (2019).

For the sake of this example, we assume for simplicity that the firm perfectly knows its product's riskiness, which means that there are two types of firms: *bad* firms with prior $\theta = 0$ and good firms with prior $\theta = 1$.

We construct an IC mechanism for which $\underline{x}_0 < \underline{x}_1$. We start by setting a launch threshold $\underline{x}_1 < 0$ for the bad firm and choose $\eta_1 = \eta(\underline{x}_1)$ low enough that (i) a bad firm forced to launch the product at \underline{x}_1 abandons the product at a threshold $\bar{x}_1 > 0$ such that $\bar{x}_1 > |\underline{x}_1|$ and (ii) this strategy yields a strictly positive expected payoff to the bad firm. Such a construction is always possible by choosing \underline{x}_1 close enough to 0.

Next, we fix some launch threshold $\underline{x}_0 \in (2\underline{x}_1, \underline{x}_1)$ for the good firm and choose $\eta_0 = \eta(\underline{x}_0)$ so that a good firm is indifferent between launching the product at \underline{x}_0 and at \underline{x}_1 . Such a construction is always possible by choosing \underline{x}_0 close enough to \underline{x}_1 and η_1 slightly lower than η_0 . By construction, a good firm is indifferent between the two items of menu $\{(\underline{x}_0, \eta_0) \text{ and } (\underline{x}_1, \eta_1)\}$.

To demonstrate incentive compatibility, there remains to show that a bad firm prefers the second item on this menu. Let \bar{x}^d denote the optimal abandonment threshold of the bad firm if it launches its product at \underline{x}_0 . Suppose first that $\bar{x}^d \leq 0$. This means that the firm prefers to abandon immediately, starting from $X_0 = 0$. This yields an expected payoff of zero and is dominated by the item $(\underline{x}_1, \eta_1)$. Now suppose that $\bar{x}^d > 0$. By construction, \underline{x}_1 is closer to \underline{x}_0 than it is to \bar{x}^d . Lemma 13, then implies that at state \underline{x}_1 , a good firm gets a strictly higher payoff than the bad firm by adopting the strategy $(\underline{x}_0, \bar{x}^d)$. The good firm is by construction indifferent between the strategy $(\underline{x}_0, \bar{x}_0)$ and stopping immediately at \underline{x}_1 , so the good firm weakly prefers to stop immediately to adopt the strategy $(\underline{x}_0, \bar{x}^d)$. Moreover, both types get exactly the same payoff if they stop at \underline{x}_1 ,

since both firms pay the $\eta(x_1)$. Therefore, the strategy $(\underline{x}_0, \bar{x}_d)$ must be strictly worse for the bad firm than stopping at \underline{x}_1 . This shows that the mechanism is incentive compatible for both types of firms, and that the risky firm ($\theta = 1$) launches the product with less evidence than the safe one ($\theta = 0$), since $\underline{x}_0 < \underline{x}_1$.

Lemma 13. *Consider the strategy that consists in launching the product at \underline{x} and abandoning it at $\bar{x} > \underline{x}$, and consider any $x \in (\underline{x}, (\underline{x} + \bar{x})/2)$. If $\eta(\underline{x}) < \pi$, the expected payoff from the strategy, starting from $X_0 = x$, is higher for the good firm than for the bad firm.*

PROOF. We follow the notation used in the proof of Lemma 12. Starting from $X_0 = x$, the Optional Sampling Theorem applied to the identity function $X_t \mapsto X_t$ and to type $\theta = 1$ implies that

$$E[X_\tau | x, \theta = 1] = x + E\left[\int_0^\tau -1 dt\right] = x - T^b.$$

Expressing the expectation on the left-hand side in terms of hitting probability f^b and rearranging yields:

$$(-1) \cdot T^b = (\underline{x} - x)f^b + (\bar{x} - x)(1 - f^b) = (\bar{x} - x) - f^b(\bar{x} - \underline{x})$$

Proceeding similarly for type $\theta = 0$, we get:

$$T^g = (\underline{x} - x)f^g + (\bar{x} - x)(1 - f^g) = (\bar{x} - x) - f^g(\bar{x} - \underline{x})$$

Summing the last two equations yields

$$(3.7) \quad T^g - T^b = 2(\bar{x} - x) - (f^g + f^b)(\bar{x} - \underline{x}).$$

We have $x \in (\underline{x}, (\underline{x} + \bar{x})/2)$, which implies that $2(\bar{x} - x) < \bar{x} - \underline{x}$ and that $f^g + f^b > 1$.³ Therefore (3.7) is negative, which shows that $T^b > T^g$. The difference of the good and bad types' expected payoffs is given by:

$$(f^g(\pi - \eta) - cT^g) - (f^b(\pi - \eta(\underline{x})) - cT^b) = \underbrace{[f^g - f^b]}_{>0} \underbrace{(\pi - \eta)}_{>0} + c \underbrace{[T^b - T^g]}_{>0}.$$

which is strictly positive. ■ □

3.5. Reducing Recklessness

Proposition 19, shows that the regulator would like to implement lower thresholds than the firm when the firm faces with a uniform penalty, regardless of the firm's private information. The next proposition shows that under these circumstances, it is without loss of generality to focus on tariffs that are nondecreasing functions of x , i.e., which impose a lower penalty, the more due diligence is demonstrated by the firm.

Proposition 21. *Suppose that Θ is finite and consider any thresholds $\{x_\theta\}_{\theta \in \Theta}$ that are (i) decreasing in θ and (ii) such that $x_\theta \leq \underline{x}_\theta^l$ for all $\theta \in \Theta$. Then, there exists a non-decreasing, piecewise-constant tariff ψ such that $\underline{x}_\theta^\psi = x_\theta$ for all $\theta \in \Theta$.*

³For the latter inequality, notice that the drifts of X_t are exact opposite for good and bad firms, so that $\frac{1}{2}(f^g + f^b)$ is the probability that the Brownian process X_t with drift either 1 or -1 with equal probability hits \underline{x} before \bar{x} when starting from x . Since x is closer to \underline{x} than it is to \bar{x} , this probability is greater than 1/2, which implies that $f^g + f^b > 1$.

PROOF. We index the elements of Θ from the smallest θ_1 to the largest $\theta_{|\Theta|}$ and construct the tariff ψ by moving from large values of x to lower ones. We start by setting $\psi(x) \equiv l$ for all $x \geq x_{\theta_1}$. At x_{θ_1} , we lower the tariff to a level ψ_1 that makes θ_1 exactly indifferent between launching the product at x_{θ_1} and at $\underline{x}_{\theta_1}^l$. We keep ψ constant at the level ψ_1 for $x \in (x_{\theta_2}, x_{\theta_1}]$. Since a firm's launch threshold when it faces a constant tariff \hat{l} is decreasing in \hat{l} , and since $\psi_1 < l$, we have

$$x_{\theta_1} < \underline{x}_{\theta_1}^l \leq \underline{x}_{\theta_1}^{\psi_1}$$

where $\underline{x}_{\theta_1}^{\psi_1}$ is the launch threshold used by type θ_1 when the tariff is constant and equal to ψ_1 . This implies that type θ_1 prefers threshold x_{θ_1} to any level $x \in (x_{\theta_2}, x_{\theta_1})$.

At x_{θ_2} , we lower the tariff ψ to a level ψ_2 that makes type θ_2 exactly indifferent between launching the product at x_{θ_2} and at its preferred level \hat{x}_2 among all $x > x_{\theta_2}$, given the tariff ψ constructed so far. By the single-crossing property established in Lemma 12, this implies that θ_1 prefers \hat{x}_2 to any x_{θ_2} and, combined with the previous paragraph, that θ_1 prefers x_{θ_1} to any $x \geq x_{\theta_2}$.

We set ψ equal to ψ_2 for all $x \in (x_{\theta_3}, x_{\theta_2}]$. Since $x_{\theta_2} \leq \underline{x}_{\theta_2}^l \leq \underline{x}_{\theta_2}^{\psi_2}$, type θ_2 prefers x_{θ_2} to any $x \in (x_{\theta_3}, x_{\theta_2})$. Another application of Lemma 12 guarantees that type θ_1 also prefers x_{θ_2} to any $x \in (x_{\theta_3}, x_{\theta_2})$.

Proceeding iteratively, we then lower ψ at x_{θ_3} to a level ψ_3 that makes type θ_3 exactly indifferent between launching the product at x_{θ_3} and at its preferred level $\hat{x}_3 > x_{\theta_3}$ given the tariff ψ constructed so far. Repeated applications of Lemma 12 guarantee that types θ_1, θ_2 prefer their respective thresholds $x_{\theta_1}, x_{\theta_2}$ to x_{θ_3} . We extend ψ by setting it constant, equal to ψ_3 for all $x \in (x_{\theta_4}, x_{\theta_3}]$. The proof is completed by induction. ■ □

3.6. Taxation Principle with Identifiable Information Acquisition

When an IC mechanism implements distinct thresholds for distinct types, the conclusion of Theorem 4 is a corollary of the Taxation Principle with Non-Contractible Events of our companion paper (Poggi and Strulovici (2020)).

According to that paper, a mechanism is *identifiable* if satisfies two conditions that we translate into the present setting. Let A denote the set of all possible strategies by the firm. Each element of A consists of a pair (τ, d) , where τ is a stopping time adapted to the filtration of X and d is measurable with respect to \mathcal{F}_τ^X . For any subset A' of A , let $X(A')$ denote the set of observable outcomes by the regulator if the firm chooses an action $a \in A'$ and causes some damage.

Definition 22. *An IC mechanism M is identifiable if there exists a partition $\mathcal{A} = \{A_k\}_{k=1}^K$ of A such that*

- (i) $X(A_k) \cap X(A_{k'}) = \emptyset$ for all $k \neq k'$.
- (ii) *All types θ who choose an action in A_k under the mechanism choose the same action of A_k .*

Proposition 22. *If M implements distinct launch thresholds for all types, then it is identifiable.*

PROOF. For each θ , let A_θ denote the set of firm strategies that use launch threshold \underline{x}_θ , and let $A_0 = A \setminus (\cup_{\theta \in \Theta} A_\theta)$. By assumption on M , $\underline{x}_\theta \neq \underline{x}_{\theta'}$ for all $\theta \neq \theta'$. Therefore, A_θ and $A_{\theta'}$ are disjoint for all $\theta \neq \theta'$ and $\mathcal{A} = \{A_0, A_\theta : \theta \in \Theta\}$ forms a partition of A . Condition (ii) is trivially satisfied since for each cell of \mathcal{A} there is at most one type taking action in that cell. Moreover Condition (i) is also satisfied by construction of the

partition: $X(A_\theta) = \{\underline{x}_\theta\}$ for all $\theta \in \Theta$ and, hence, $X(A_\theta) \cap X(A_{\theta'}) = \emptyset$ for all $\theta \neq \theta'$.

■

□

Corollary 6. *If an IC mechanism M implements distinct launch threshold for all types, it can be implemented by a tariff mechanism.*

PROOF. Proposition 22 implies that M is identifiable. The result then immediately follows from Theorem 1 in Poggi and Strulovici (2020)

■

□

Bibliography

- Angeletos, G.-M. and Werning, I. (2006). Crises and prices: Information aggregation, multiplicity, and volatility. *american economic review*, 96(5):1720–1736.
- Asquith, P., Beatty, A., and Weber, J. (2005). Performance pricing in bank debt contracts. *Journal of Accounting and Economics*, 40(1-3):101–128.
- Athey, S. (2001). Single crossing properties and the existence of pure strategy equilibria in games of incomplete information. *Econometrica*, 69(4):861–889.
- Baumol, W. J. (1965). *The stock market and economic efficiency*. Fordham University Press.
- Bernanke, B. S. and Woodford, M. (1997). Inflation forecasts and monetary policy. *Journal of Money, Credit, and Banking*, pages 653–684.
- Biagi, A. and Denicolò, V. (2014). Timing of discovery and the division of profit with complementary innovations. *Journal of Economics and Management Strategy*, 23(1):89–102.
- Boleslavsky, R., Kelly, D. L., and Taylor, C. R. (2017). Selloffs, bailouts, and feedback: Can asset markets inform policy? *Journal of Economic Theory*, 169:294–343.
- Bond, P., Edmans, A., and Goldstein, I. (2012). The real effects of financial markets. *Annu. Rev. Financ. Econ.*, 4(1):339–360.
- Bond, P. and Goldstein, I. (2015). Government intervention and information aggregation by prices. *The Journal of Finance*, 70(6):2777–2812.

- Bond, P., Goldstein, I., and Prescott, E. S. (2010). Market-based corrective actions. *The Review of Financial Studies*, 23(2):781–820.
- Börgers, T. (2017). (no) foundations of dominant-strategy mechanisms: a comment on Chung and Ely (2007). *Review of Economic Design*, 21(2):73–82.
- Breon-Drish, B. (2015). On existence and uniqueness of equilibrium in a class of noisy rational expectations models. *The Review of Economic Studies*, 82(3):868–921.
- Bryan, K. A. and Lemus, J. (2017). The direction of innovation. *Journal of Economic Theory*, 172:247–272.
- Bryan, K. A., Lemus, J., and Marshall, G. (2020). Crises and the Direction of Innovation.
- Che, Y.-K., Choe, C., and Rhee, K. (2018). Bailout stigma. *Available at SSRN 3208238*.
- Che, Y.-k. and Mierendorff, K. (2017). Optimal Sequential Decision with Limited Attention.
- Cohen, C., Abbas, S. A., Anthony, M., Best, T., Breuer, P., Miao, H., Myrvoda, A., Togo, E., et al. (2020). The role of state-contingent debt instruments in sovereign debt restructurings. Technical report, International Monetary Fund.
- Cordella, T., Ricci, L. A., and Ruiz-Arranz, M. (2010). Debt overhang or debt irrelevance? *IMF Staff Papers*, 57(1):1–24.
- Crowley, J. (2015). *Central and commercial bank balance sheet risk before, during, and after the global financial crisis*. International Monetary Fund.
- Dow, J. and Gorton, G. (1997). Stock market efficiency and economic efficiency: is there a connection? *The Journal of Finance*, 52(3):1087–1129.
- Duffie, D. and Dworczak, P. (2020). Robust benchmark design. *Journal of Financial Economics*, forthcoming.

- Dworczak, P. and Pavan, A. (2020). Preparing for the worst but hoping for the best: Robust (bayesian) persuasion. mimeo, Northwestern University.
- Fershtman, C. and Kamien, M. I. (1992). Cross licensing of complementary technologies. *International Journal of Industrial Organization*, 10(3):329–348.
- Flachsland, C., Pahle, M., Burtraw, D., Edenhofer, O., Elkerbout, M., Fischer, C., Tietjen, O., and Zetterberg, L. (2020). How to avoid history repeating itself: the case for an eu emissions trading system (eu ets) price floor revisited. *Climate Policy*, 20(1):133–142.
- Friehe, T. and Schulte, E. (2017). Uncertain product risk, information acquisition, and product liability. *Economics Letters*, 159:92–95.
- Gilbert, R. J. and Katz, M. L. (2011). Efficient division of profits from complementary innovations. *International Journal of Industrial Organization*, 29(4):443–454.
- Glasserman, P. and Nouri, B. (2016). Market-triggered changes in capital structure: Equilibrium price dynamics. *Econometrica*, 84(6):2113–2153.
- Goldstein, I. and Guembel, A. (2008). Manipulation and the allocational role of prices. *The Review of Economic Studies*, 75(1):133–164.
- Green, J. R. and Scotchmer, S. (1995). On the Division of Profit in Sequential Innovation. *RAND Journal of Economics*, 26(1):20–33.
- Greenspan, A. (2001). Harnessing market discipline: the region (september). *Federal Reserve Bank of Minneapolis, MN*, pages 10–11.
- Grochulski, B. and Wong, R. (2018). Contingent debt and performance pricing in an optimal capital structure model with financial distress and reorganization. *FRB Richmond Working Paper*.

- Grossman, S. J. and Stiglitz, J. E. (1980). On the impossibility of informationally efficient markets. *The American economic review*, 70(3):393–408.
- Hauk, E., Lanteri, A., and Marcet, A. (2020). Optimal policy with general signal extraction. *Economic Research Initiatives at Duke (ERID) Working Paper*.
- Hayek, F. A. (1945). The use of knowledge in society. *The American economic review*, 35(4):519–530.
- Hellwig, M. (1980). On the aggregation of information in competitive markets. *Journal of Economic Theory*, 22(3):477–498.
- Henry, E., Loseto, M., and Ottaviani, M. (2021). Regulation with Experimentation: Ex Ante Approval, Ex Post Withdrawal, and Liability.
- Henry, E. and Ottaviani, M. (2019). Research and the Approval Process. *American Economic Review*, 109(3):911–955.
- Jensen, R. T. and Miller, N. H. (2008). Giffen behavior and subsistence consumption. *American economic review*, 98(4):1553–77.
- Jordan, J. S. (1982). The generic existence of rational expectations equilibrium in the higher dimensional case. *Journal of Economic Theory*.
- Ke, T. T. and Villas-boas, J. M. (2019). Optimal learning before choice. *Journal of Economic Theory*, 180:383–437.
- Klabjan, D., Olszewski, W., and Wolinsky, A. (2014). Attributes. *Games and Economic Behavior*, 88:190–206.
- Kolstad, C. D., Ulen, T. S., and Johnson, G. V. (1990). Ex Post Liability for Harm vs Ex Ante Safety Regulation: Substitutes or Complements? *American Economic Review*, 80(4):888–901.

- Lee, Y.-H. A. (2019). A model of stock market-based rulemaking. *Available at SSRN 3440321*.
- Liang, A. and Mu, X. (2020). Complementary Information and Learning Traps. *The Quarterly Journal of Economics*, 135(1):389–448.
- Liang, A., Mu, X., and Syrgkanis, V. (2018). Optimal and myopic information acquisition. *ACM EC 2018 - Proceedings of the 2018 ACM Conference on Economics and Computation*, pages 45–46.
- Lucas, R. E. et al. (1976). Econometric policy evaluation: A critique. In *Carnegie-Rochester conference series on public policy*, volume 1, pages 19–46.
- Mayskaya, T. (2019). Dynamic Choice of Information Sources.
- McClellan, A. (2019). Experimentation and Approval Mechanisms.
- Ménière, Y. (2008). Patent law and complementary innovations. *European Economic Review*, 52(7):1125–1139.
- Moroni, S. (2019). Experimentation in Organizations.
- Nikandrova, A. and Pancs, R. (2018). Dynamic project selection. *Theoretical Economics*, 13(1):115–143.
- Ozdenoren, E. and Yuan, K. (2008). Feedback effects and asset prices. *The journal of finance*, 63(4):1939–1975.
- Pálvölgyi, D. and Venter, G. (2015). Multiple equilibria in noisy rational expectations economies. *Available at SSRN 2524105*.
- Poggi, F. and Strulovici, B. (2020). A Taxation Principle with Non-Contractible Events.
- Schmitz, P. W. (2000). On the joint use of liability and safety regulation. *International Review of Law and Economics*, 20(3):371–382.

- Scotchmer, S. and Green, J. (1990). Novelty and Disclosure in Patent Law. *RAND Journal of Economics*, 21(1):131–146.
- Shavell, S. (1984). Liability for Harm versus Regulation of Safety. *Journal of Legal Studies*, XIII:357–374.
- Shavell, S. (1992). Liability and the Incentive to Obtain Information about Risk. *Journal of Legal Studies*, XXI(June):259–270.
- Siemroth, C. (2019). The informational content of prices when policy makers react to financial markets. *Journal of Economic Theory*, 179:240–274.
- Wald, A. (1945). Sequential Tests of Statistical Hypotheses. *The Annals of mathematical Statistics*, 16(2):117–186.
- Warner, J. B., Watts, R. L., and Wruck, K. H. (1988). Stock prices and top management changes. *Journal of financial Economics*, 20:461–492.
- Wissenbach, I. and Taylor, E. (2020). Lufthansa soars after top shareholder backs bailout. *Reuters*.
- Woodford, M. (1994). Determinacy of equilibrium under alternative policy regimes. *Economic Theory*, 4(3):323–326.

APPENDIX A

Appendix I

A.1. Omitted proofs from Section 1.3

A.1.1. Proof Proposition 1

The proof is based on Lemma 14 that is stated for k projects. First, some preliminaries.

For any stationary strategy $x \in \mathcal{X}$ and initial state (S, X_0) , there is a *trajectory* $y_S : \mathbb{R}_+ \times \mathbb{R}^k \rightarrow \bar{\mathbb{R}}^k$ that is the (unique) solution to the differential equations

$$\begin{cases} \nabla y_S(t, X_0) = x(S, y_S(t, X_0)) \\ y_S(0, X_0) = X_0 \end{cases}$$

We will refer to $Y(S, X_0) = \lim_{t \rightarrow \infty} y_S(t, X_0)$ as the abandonment point of the strategy x given an initial state (S, X_0) .

Definition 23. *A strategy has increasing abandonment points if*

$$Y(S, X_0) \leq Y(\hat{S}, X_0) \text{ for all } S \subseteq \hat{S}.$$

Definition 24. *Two strategies x, \tilde{x} have the same abandonment points if for each initial state, the abandonment point is the same for both strategies.*

Lemma 14. *If two strategies x, \tilde{x} have the same abandonment points, and these abandonment points are increasing then the two strategies have the same expected payoff.*

PROOF. The proof works by induction. The Lemma trivially holds for $k = 1$. Assume that it holds for $k = 1, 2, \dots, m - 1$, we want to show that it holds for $k = m$.

Consider strategies x, \tilde{x} and a initial state (\emptyset, X_0) . Let $Y(\emptyset, X_0)$ be the associated abandonment point. For each set $S \neq \emptyset$, the continuation problem is analogous to one with less than n projects, so the lemma holds. Let $V(S, X)$ be the value of the two strategies at the state (S, x) for $S \neq \emptyset$.

Consider a strategy z with the same abandonment points than x and such that for any $S \neq \emptyset$, $z(S, X) = x(\emptyset, X)$ for all X with $x(\emptyset, X) \neq 0$. We can do this since $Y(S, X_0) \geq Y(\emptyset, X_0)$. For any τ , the new strategy has the same payoff than the original: either no project is successful and both abandon at the same point or the same project is succesful at the same point, and the continuation value is the same.

Similarly, we can construct a strategy \tilde{z} with the same abandonment points but such that for any $S \neq \emptyset$, $\tilde{z}(S, X) = \tilde{x}(\emptyset, X)$ for all X with $\tilde{x}(\emptyset, X) \neq 0$. \tilde{z} and \tilde{x} shield the same payoff. We end the proof by showing that z and \tilde{z} must also have the same payoff for any realization of the success times $\tau \in \mathbb{R}_+^k$.

Let $\bar{S} = \{i \in K : \tau_i < Y(\emptyset, X_0)\}$, that is the set of projects which completion time is below the abandonment point. Both z and \tilde{z} reach $Y(\emptyset, X_0)$ with probability one. The payoff is therefore

$$V(\bar{S}, Y(\emptyset, X_0)) - c \sum_{i \in \bar{S}} \tau_i - c \sum_{i \notin \bar{S}} Y_i(\emptyset, X_0)$$

Taking expectation over the realization of τ completes the proof. □

The intuition for Lemma 14 is the following: if the abandonment is increasing, then the current abandonment point is the least attention you are willing to put on the remaining projects by the end of the day. Since the attention it is going to be paid eventually, the order in which the agent does it is not gonna determine the outcome.

Proposition 1. Consider $k = 2$. If the projects are complements, then $V = \hat{V}$. If the projects are not complements, there exists a family of distribution $\{F_i\}$ such that $V \neq \hat{V}$.

(\Leftarrow). We want to show that q supermodular implies that any strategy that has the same abandonment points than an optimal strategy is also optimal.

Lemma 15. *For two complementary projects, any optimal strategy has increasing abandonment points.*

PROOF. By contradiction, assume that $Y_i(j, X_0) < Y_i(\emptyset, X_0)$. Then, \square

Lemma 16. *For two affiliated and complementary projects, any optimal strategy has increasing abandonment points.*

PROOF. We want to prove that for any optimal strategy x , $Y_i(j, X_0) \geq Y_i(\emptyset, X_0)$. By supermodularity of q , the marginal value of i is larger when j was completed than when it is not. If it is optimal to work on project i when it is not clear if j is going to be completed or not, it must be optimal to work on i when j was already completed.

Formally, by contradiction assume $Y_i(j, X_0) < Y_i(\emptyset, X_0)$. Then there is a time t such that $y_{\emptyset,i}(t, X_0) = Y_i(j, X_0)$. Let $X := y_{\emptyset}(t, X_0)$. Consider the trajectory starting at (\emptyset, X) associated with strategy x . If this strategy was copied starting on the state $(j, Y_i(j, X_0))$ with a dummy project j' that starts at x_j the expected continuation payoff must be negative (otherwise it is worth continuing at $Y_i(j, X_0)$).

$$\begin{aligned}
& q(j) + \int_{Y_i(j, X_0)}^{Y_i(\emptyset, X_0)} \frac{1 - F_i(\tau_i | X_0)}{1 - F_i(Y_i(j, X_0))} [h_i(\tau_i)(q(ij) - q(j)) - c] d\tau_i \geq \\
& q(j) + \int_0^T \frac{1 - F(Y(j, X_0) + y_\emptyset(\tau, X) - X)}{1 - F(Y(j, X_0))} \alpha_i(\tau_i) h_i(\tau_i) (q(ij) - q(j)) d\tau \\
& \quad - \int_{Y_i(j, X_0)}^{Y_i(\emptyset, X_0)} \frac{1 - F_i(\tau_i | X_0)}{1 - F_i(Y_i(j, X_0))} c d\tau \geq \\
& \int_0^T \frac{1 - F(Y(j, X_0) + y_\emptyset(\tau, X) - X)}{1 - F(Y(j, X_0))} [\alpha_i(\tau_i) h_i(\tau_i) (q(ij) - q(j)) + \alpha_j(\tau) h_j(\tau) q(j)] d\tau \\
& \quad - \int_{Y_i(j, X_0)}^{Y_i(\emptyset, X_0)} \frac{1 - F_i(\tau_i | X_0)}{1 - F_i(Y_i(j, X_0))} c d\tau \geq \\
& \int_0^T \frac{1 - F(Y(j, X_0) + y_\emptyset(\tau, X) - X)}{1 - F(Y(j, X_0))} [\alpha_i(\tau_i) h_i(\tau_i) q(i) + \alpha_j(\tau) h_j(\tau) q(j)] d\tau \\
& \quad - \int_{Y_i(j, X_0)}^{Y_i(\emptyset, X_0)} \frac{1 - F_i(\tau_i | X_0)}{1 - F_i(Y_i(j, X_0))} c d\tau \geq \\
& \quad V(\emptyset, Y(\emptyset, X)) \geq 0
\end{aligned}$$

But this strategy shields more than the continuation at X thus the project should stop X , so $X = Y(\emptyset, x_0)$ leading to a contraction. \square

Using Lemma 14, any strategy that has the same abandonment points than x is gonna get the same expected payoff and therefore be optimal.

(\Rightarrow). : The proof works by contrapositive. If q is *not* supermodular, there are cdfs $\{F_i, F_j\}$ such that $Y_i(j, X_0) < Y_i(\emptyset, X_0)$.

PROOF. Since q is not supermodular, $q(\{i\}) > q(\{i, j\}) - q(\{j\})$. Let $F_i = 1 - e^{-\lambda_i x}$ with λ_i such that

$$q(\{i\}) > \frac{c}{\lambda_i} > q(\{i, j\}) - q(\{j\})$$

and let j never succeed, i.e. $F_j = 0$. Rearranging we have that

$$\lambda q(\{i\}) - c > 0 > \lambda(q(\{i, j\}) - q(\{j\})) - c$$

What implies that for any X_0 , $Y_i(\emptyset, X_0) = \infty$ and $Y_i(\{j\}, X_0) = X_0$. □

A.1.2. Proof of Lemma 1

PROOF. h_i decreasing implies that $V(S, \cdot)$ is decreasing for all S .

If $X = X^*(\{A\}, X) = X^*(\{B\}, X)$, then $V(S, X) = q(S)$ and by complementarity of K , $V(\cdot, X)$ is supermodular. If $X \neq X^*(\{i\}, X)$ for some i ,

$$\frac{\partial V(K \setminus \{i\}, X)}{\partial X_i} = c - h_i(X_i) [V(K, X) - V(K \setminus \{i\}, X)] \leq 0$$

At any point $X \in X^*(\emptyset, \mathbb{R}_+^k)$ it must be that $c \geq h_i(x_i) \cdot V(\{i\}, X)$. So,

$$h_i(X_i) V(\{i\}, X) - h_i(X_i) [V(\{A, B\}, X) - V(K \setminus \{i\}, X)] \leq 0 \quad \forall X \in X^*(\emptyset, \mathbb{R}_+^k)$$

Rearranging,

$$V(\{B\}, X) + V(\{A\}, X) \leq V(\{A, B\}, X) \quad \forall X \in X^*(\emptyset, \mathbb{R}_+^k)$$

□

A.2. Omitted proofs from Section 1.4

A.2.1. Proof of Propositions 2 and 3

Some preliminaries. Let δ_i be $\lambda_i^H - \lambda_i^L$. Using Bayes' rule, the beliefs $p_i(X)$ evolve

$$p_i(X) = \frac{p_i e^{-\delta_i x_i}}{(1 - p_i) + p_i e^{-\delta_i x}}$$

As the agent becomes more pessimistic, the subjective hazard rate $h_i(t)$ becomes lower.

$$h_i(X) = \lambda_L^i + p_i(X)\delta_i$$

Notice that

$$g_i > 1 \quad \Leftrightarrow \quad \frac{\lambda_i^L \cdot \lambda_i^H}{\lambda_i^L + \lambda_i^H} > c$$

We prove the alternative:

Proposition 3':

1. If $g_i > 1$ for both projects, then it is optimal to work on them in sequence.
2. If $g_i < 1$ for both projects, then the greedy strategy is optimal.

The proof of the proposition is split in three lemmatas. Lemma 17 proves that g_i controls the monotonicity of project i 's hazard-to-value ratio. Lemma 18 shows that when both projects have an increasing hazard-to-value ratio, it is efficient to work on them in sequence. Lemma 19 shows that when both projects have a decreasing hazard-to-value ratio the greedy strategy is efficient.

Lemma 17. h_i/v_i is monotone. Moreover, $\text{sgn}((h_i/v_i)') = \text{sgn}(g_i - 1)$.

PROOF. First we show that the monotonicity of h/v depends on whether the value v is higher or lower than an expression R .

$$\begin{aligned}
 \text{sgn}((h_i/v_i)') &= \text{sgn}(h'_i v_i - h_i v'_i) \\
 &= \text{sgn}(h'_i v_i - h_i(c - h_i(1 - v_i))) \\
 &= \text{sgn}\left(\underbrace{\frac{h_i(h_i - c)}{h_i^2 + h'_i}}_{R(t)} - v_i\right)
 \end{aligned}$$

Change of variables. In belief space, the concavity of R is determined by whether g_i is larger or lower than one.

$$\begin{aligned}
 \hat{R}'(p) &= \frac{2\delta^2 \lambda_L \lambda_H (\lambda_L \lambda_H - c(\lambda_L + \lambda_H))}{(\lambda_L^2 + p\delta(\lambda_L + \lambda_H))^3} \\
 &= M(g - 1)
 \end{aligned}$$

Now we consider two cases: $\lambda_L < c$ and $\lambda_L \geq c$.

Case I: $\lambda_L \geq c$. In this case, the agent would never stop. The value is linear in the beliefs.

$$v(p) = 1 - p \frac{c}{\lambda_H} - (1 - p) \frac{c}{\lambda_L}$$

Since $v(0) = R(0)$ and $v(1) = R(1)$,

$$g > 1 \quad \Leftrightarrow \quad v(p) < R(p) \quad \forall p \in (0, 1)$$

Case II: $\lambda_L < c$. In this case, the agent abandons if sufficient time passes with no success. v is convex (information is valuable) and R is concave:

$$\lambda_L < c \quad \Rightarrow \quad \frac{\lambda_H}{\lambda_L + \lambda_H} \lambda_L < c \quad \Leftrightarrow \quad g_i > 1$$

Since $v(1) = R(1)$ and $v(\hat{p}) = R(\hat{p})$ where \hat{p} is the stopping belief.

$$v(p) < R(p) \quad \text{for any } p \in (\hat{p}, 1)$$

□

Lemma 18. *If h_i/v_i is strictly increasing for $i = A, B$, it is optimal to work on the projects in sequence.*

PROOF. By contradiction. Assume that $x := Y(\emptyset, x_0) > 0$. Let $r_i(t) := \frac{h_i(t)}{v_i(t)}$ and $g_i(t) := \frac{h'_i(t)}{v'_i(t) \cdot r_i(t)}$. Since x is an interior stopping point, it must be that $h_A(x_A)v_B(x_B) = h_B(x_B)v_A(x_A) = c$.

Abusing notation I write f_i instead of $f_i(x_i)$. First we show that $r'_A + r'_B > 0$ implies $\frac{h'_A v_A}{h_A v'_A} \cdot \frac{h'_B v_B}{h_B v'_B} < 1$.

$$\begin{aligned} r'_A + r'_B > 0 & \Leftrightarrow \frac{h'_A v_A - h_A v'_A}{v_A^2} + \frac{h'_B v_B - h_B v'_B}{v_B^2} > 0 \\ & \Leftrightarrow \frac{h_A v'_A}{v_A^2} \left(\frac{h'_A v_A}{h_A v'_A} - 1 \right) + \frac{h_B v'_B}{v_B^2} \left(\frac{h'_B v_B}{h_B v'_B} - 1 \right) > 0 \end{aligned}$$

For all (x_A, x_B) such that $h_A v_B = h_B v_A = c$,

$$\frac{h_A v'_A}{v_A^2} = \frac{h_B v'_A}{v_B v_A} = \frac{h_A v'_B}{v_A v_B} = \frac{h_B v'_B}{v_B^2}$$

Where the first and last equality use $h_A/v_A = h_B/v_B$ and the intermediate one uses that $h_B v'_A = h_B(c - h_A(1 - v_A)) = -h_B h_A(1 - v_A - v_B)$ (since $c = h_A v_B$) and equal to $h_A v'_B$ by symmetry. So,

$$\begin{aligned} r'_A + r'_B > 0 &\Leftrightarrow \frac{h_A v'_A}{v_A^2} \left[\left(\frac{h'_A v_A}{h_A v'_A} - 1 \right) + \left(\frac{h'_B v_B}{h_B v'_B} - 1 \right) \right] > 0 \\ &\Leftrightarrow \left[\frac{h'_A v_A}{h_A v'_A} + \frac{h'_B v_B}{h_B v'_B} \right] < 2 \\ &\Leftrightarrow \frac{h'_A v_A}{h_A v'_A} \cdot \frac{h'_B v_B}{h_B v'_B} < 1 \end{aligned}$$

Where the second implication uses that v_A is decreasing. And the last one is that the sum of two positive numbers being less than two implies that the product is less than one.

The determinant of the Hessian for the value function $V(\emptyset, x)$ is

$$\det(H) = (1 - F_A)(1 - F_B)[h'_A h'_B v_A v_B - h_A h_B v'_A v'_B]$$

So

$$\det(H) < 0 \quad \text{iff} \quad \frac{h'_A v_A}{h_A v'_A} \cdot \frac{h'_B v_B}{h_B v'_B} < 1$$

And $\det(H) < 0$ rules out the candidate as an optimum (saddle point). \square

Lemma 19. *If h_i/v_i is strictly decreasing for $i = A, B$, the greedy strategy is optimal.*

PROOF.

$$\begin{aligned}
 r_i \searrow & \Leftrightarrow h'_i v_i - h_i v'_i < 0 \\
 & \Leftrightarrow \frac{h'_i v_i}{h_i v'_i} > 1
 \end{aligned}$$

So,

$$r_A \searrow \text{ and } r_B \searrow \Rightarrow \frac{h'_A v_A}{h_A v'_A} \cdot \frac{h'_B v_B}{h_B v'_B} > 1$$

This implies that there is at most one interior candidate for solution ($h_A v_B = h_B v_A = c$), and that if it exist it is the actual solution. We consider the two cases.

Case I: there is an interior solution candidate. Then this is the actual solution. Since at the solution $r_A = r_B$ and the r_i are decreasing, by working always on the project with highest r_i , the point is eventually reached.

Case II: there in no interior solution candidate. Then it must be that $h_i v_j = c \Rightarrow h_j v_i > c$. Thus, the solution is to work in sequence starting with project j . Moreover, $h_j/v_j > h_i/v_i$ for all x such that $h_i v_j \geq c$, so the greedy strategy also works in sequence starting with j . \square

A.3. Equilibrium

There is a continuum of agents, $m \in [0, 1]$. Each agen decides, at each instant, what project to work on $\alpha_m(t) \in \{A, B\}$. Once all agents stop developing, the value of the joint development is split across the agents. The payoff of an agent will depend on what

innovations are successfully developed, who developed the innovations, and the timing of development.

We take a reduced-form approach to the problem, with focus on the first stage: let $W_i(X)$ be the expected payoff that is captured by the first agent to innovate when the innovation is i and the state was (\emptyset, X) .

Definition 25. A first-stage strategy for agent $m \in [0, 1]$ is a function $s_m : \mathbb{R}_+^2 \rightarrow \{A, B, \emptyset\}$.

We assume that the problem is non-trivial: if both projects are easy then it is efficient to work on them, and it is efficient to abandon immediately if both projects were known to be difficult.¹

A concern is that with complements, the competition in development stage will bias the allocation of resources toward the projects that can be developed faster, and that these will be developed first even if it is efficient to leave these projects for later.

The payoff for an agent m that was working on a project i at the moment in which this project was completed is proportional to the value of that project $V(\{i\}, X)$ and inversely proportional to the mass of individuals working on that project at that point. Let $\pi(s, Y, x)$ be the expected payoff of following first-stage strategy s when the evolution of the process is Y and the aggregate strategy is x .

¹In terms of the parameters the condition is

$$\frac{\lambda_A^L \cdot \lambda_B^L}{\lambda_A^L + \lambda_B^L} < \frac{c}{\gamma} < \frac{\lambda_A^H \cdot \lambda_B^H}{\lambda_A^H + \lambda_B^H}.$$

Definition 26. A first-stage stationary equilibrium consists of a first-stage strategy for each agent, an industry allocation $x : \mathbb{R}_+^2 \rightarrow [0, 1]^2$ and an evolution of cumulative resources $Y : \mathbb{R}_+ \times \mathbb{R}_+^2 \rightarrow \mathbb{R}_+^2$ such that

(1) Each agent maximizes his expected payoff given the evolution of resources.

$$s(X) \in \arg \max_{s' \in \mathcal{S}} \pi(s, Y(\cdot, X), x(X))$$

(2) The industry allocation aggregates all individual strategies.

$$x(X) = \int_0^1 s_m(X) \, dm$$

(3) The evolution of resources is consistent with the allocation of the industry.

$$\begin{cases} \nabla Y(t, X) = x(Y(t, X)) \\ Y(0, X) = X \end{cases}$$

Lemma 20. Any stationary equilibrium satisfies

$$x_i(\emptyset, X) > 0 \quad \Rightarrow \quad h_i(X)V(\{i\}, X) \geq \max\{c, h_j(X)V(\{j\}, X)\}$$

PROOF. in the Appendix A.3.1. □

Patent races might introduce distortions.

Definition 27. There is a race effect when there is no equilibrium whose allocation is efficient.

A.3.1. Proof of Lemma 20

PROOF. Individual profits are

$$\int_0^t \underbrace{\frac{1 - F_A(Y_X(\tau))}{1 - F_A(X_A)} \cdot \frac{1 - F_B(Y_X(\tau))}{1 - F_B(X_B)}}_{\Pr(\text{reach } \tau)} \left[\underbrace{x_{s(Y_X(\tau))}(X) \cdot h(Y_X(\tau))}_{\text{rate success at } \tau} \cdot \underbrace{\frac{V(s(Y_X(\tau)), Y_X(\tau))}{x_{s(Y_X(\tau))}(X)}}_{\text{expected payoff if successful}} - c \right] d\tau$$

Since the individuals do not have a marginal effect on the trajectory and take it as given, the way to optimize the individual profits is every instant.

$$s \in \arg \max_{s' \in \{A, B, \emptyset\}} \{h_{s'}(Y_X(\tau))V(s', Y_X(\tau)) - c\}$$

□

A.4. Extensions: proofs

A.4.1. Proof of Proposition 8

PROOF. Suppose that

$$V(S, x) = \hat{V}(S, x) = \max_{\hat{S} \in 2^{K \setminus S}} \left\{ q(S \cup \hat{S}) - c \sum_{i \in \hat{S}} \lambda_i^{-1} \right\}$$

□

A.4.2. Discrete time

In this appendix we will consider the discrete time case $T = \{1, 2, 3, \dots, \infty\}$. At any time before stopping the agent decides which project to work on $\alpha_t \in \{A, B, \emptyset\}$. Let F_i be

the distribution of successes for project i , and $h_i : T \rightarrow [0, 1]$ the respective hazard rate. Finally, let $v_i : T \rightarrow [0, 1]$ the value of the joint project when only project i is incomplete as a function of the time spent working on project i .

$$v_i(x_i) := q(j) + \max_{T \geq x_i} \left\{ \sum_{x=x_i+1}^T \frac{1-F(x)}{1-F(x_i)} [h(x)(1-q(j)) - c] \right\}$$

Proposition 23. *h_i/v_i decreasing for both projects implies that the greedy strategy is efficient.*

PROOF. Grab an optimal abandonment point $x^* := Y(\emptyset, 0)$ and a trajectory to it. The trajectory has to be greedy at the time before the abandonment point. Otherwise, the optimality of x^* is violated.

Consider now a greedy trajectory and the point (x_L, x_B^*) in that trajectory where crosses $x_B = x_B^*$ (the rightmost one). If the optimum is to the right of the path ($x_A^* > x_L$) then by optimality,

$$\frac{h_A}{v_A}(x_L) \geq \frac{h_i}{v_i}(x_i^* - 1) \geq \frac{h_j}{v_j}(x_j^*)$$

If $x_L = x_i^* - 1$ then the first inequality holds with equality and there is a greedy path to the optimum: the one we considered changing at the indifferent point (x_L, x_A^*) . If $x_L < x_i^* - 1$ then by strict monotonicity of h/v the inequality holds strictly, what would violate greediness of the strategy at (x_L, x_A^*) \square

Proposition 24. *h/v increasing for both tasks implies that the efficient allocation is in sequence.*

PROOF. Suppose that the optimal stopping point $x^* = Y(\emptyset, 0)$ is interior, i.e. $x^* > 0$. Since last period is myopically optimal for each trajectory,

$$\begin{aligned}\frac{h_A}{v_A}(x_A^* - 1) &\geq \frac{h_B}{v_B}(x_B^*) > \frac{h_B}{v_B}(x_B^* - 1) \\ \frac{h_B}{v_B}(x_B^* - 1) &\geq \frac{h_A}{v_A}(x_A^*) > \frac{h_A}{v_A}(x_A^* - 1)\end{aligned}$$

Where the strict inequalities come from the h/v being increasing for both projects. Thus, a contradiction. \square

A.4.3. One h/v increasing and one decreasing

Lemma 21. *If the horizontal sum of the two h/v is increasing, then it is optimal to develop the projects in sequence.*

PROOF. Consider $q(y) := (h_A/v_A)^{-1}(y) + (h_B/v_B)^{-1}(y)$ decreasing for all $y \in R := (h_A/v_A)([0, \bar{t})) \cap (h_B/v_B)([0, \bar{t}))$. Taking the derivative this implies that

$$\frac{1}{(h_A/v_A)'((h_A/v_A)^{-1}(y))} + \frac{1}{(h_B/v_B)'((h_B/v_B)^{-1}(y))} < 0 \quad \forall y \in \mathbb{R}$$

$$\frac{(h_A/v_A)'((h_A/v_A)^{-1}(y)) + (h_B/v_B)'((h_B/v_B)^{-1}(y))}{(h_A/v_A)'((h_A/v_A)^{-1}(y)) \cdot (h_B/v_B)'((h_B/v_B)^{-1}(y))} < 0 \quad \forall y \in \mathbb{R}$$

$$(h_A/v_A)'((h_A/v_A)^{-1}(y)) + (h_B/v_B)'((h_B/v_B)^{-1}(y)) > 0 \quad \forall y \in \mathbb{R}$$

Or, in other words: $r'_A(x_A) + r'_B(x_B) > 0$ for all points (x_A, x_B) with $h_A(x_A)/v_A(x_B) = h_B(x_B)/v_B(x_B)$. We can use the same logic used in the proof of Lemma 18 to rule out interior points. \square

A.4.4. Imperfect complements

$\lambda_L > c/(1-q)$ then the agent would never stop. The value is independent of q and linear.

The monotonicity of h/v is equivalent to the case where $q = 0$.

Consider now $\lambda_L \in (c, c/(1-q))$. There is a belief at which the agent stops.

$$\hat{p} = \frac{c/(1-q) - \lambda_L}{\delta}$$

If $R(\hat{p}) > v(\hat{p}) = q$ and R is concave, h/v is increasing.

$$R(\hat{p}) > q$$

$$\frac{c^2 q}{(1-q)[c(\lambda_L + \lambda_H) - (1-q)\lambda_L \lambda_H]} > q$$

Interesting case: $[c(\lambda_L + \lambda_H) - (1-q)\lambda_L \lambda_H] > 0$.

$$\left(\frac{c}{(1-q)} \right)^2 \geq \frac{c}{(1-q)} (\lambda_L + \lambda_H) - \lambda_L \lambda_H$$

$$\frac{c}{(1-q)} \left(\frac{c}{(1-q)} - \lambda_L \right) \geq \lambda_H \left(\frac{c}{(1-q)} - \lambda_L \right)$$

$$\frac{c}{(1-q)} \geq \lambda_H$$

But if this is the case, then the agent does not wish to work on the development even when sure that it is relatively easy.

A.5. Appendix II

A.6. Proofs of Section 2.3

A.6.1. Preliminary results

It will be useful to establish some properties of θ . Let \bar{P} be the set of p for which $\theta_M(p) \neq \emptyset$. Corollary 7 below implies that for continuous M , the set of p for which $\theta_M(p) = \emptyset$ is open.

Lemma 22. *$p \mapsto \theta_M(p)$ is compact-valued. If M is continuous at p' then $p \mapsto \theta_M(p)$ is upper hemicontinuous at p' .*

PROOF. Compact valued is easy: if $R(M(p), \theta) - p \neq 0$ then by continuity of R this holds for all θ' in a neighborhood of θ .

Now upper hemicontinuity. Let V be an open neighborhood of $\theta_M(p)$. Then $\Theta \setminus V$ is compact, so there exists $\kappa > 0$ such that $R(M(p), \theta) - p > \kappa$ for all $\theta \in \Theta \setminus V$. Then by continuity of R, M there exists an open neighborhood U of p such that $R(M(p'), \theta) - p' > \kappa$, and thus $\theta_M(p') \in V$, for all $p' \in U \cap \bar{P}$. Thus $p \mapsto \theta_M(p)$ is upper hemicontinuous. \square

Lemma 23. *If R is weakly increasing in θ then $\theta_M(p)$ is convex valued.*

PROOF. $\theta_M(p) = \{\theta \in \Theta : R(M(p), \theta) = p\}$. If $R(M(p), \cdot)$ is monotone, $R(M(p), \theta') = R(M(p), \theta'') = p$ implies $R(M(p), \theta) = p$ for all $\theta \in (\theta', \theta'')$. \square

A.6.2. Proof of Theorem 1

PROOF. We first prove the result for continuous M , and then extend it to all of \mathcal{M} . Assume without loss of generality that $R(a, \cdot)$ is increasing for all a . Lemmas 26 and 27 and Corollary 7 apply.

Let $\theta_1 < \theta_2 < \theta_3$ be interior, and suppose $P(\theta_1) > P(\theta_2)$ and $P(\theta_3) > P(\theta_2)$ (the other type of non-monotonicity is dealt with symmetrically). We first want to show that $[P(\theta_2), \max\{P(\theta_1), P(\theta_3)\}] \subseteq \bar{P}$. Suppose not, so there is some $p' \in [P(\theta_2), \max\{P(\theta_1), P(\theta_3)\}]$ such that $\theta(p) = \emptyset$. Then p' is either Type L or Type H. Assume it is Type H (symmetric argument using Lemma 26 if it is Type L). First, assume $p' \in [P(\theta_2), \min\{P(\theta_1), P(\theta_3)\}]$. By Lemma 27, part (i) there is a $p > p'$ such that $\theta_2 \in \theta(p)$. Thus there is multiplicity in state θ_2 . Then by continuity of M and R , and R weakly increasing in θ , there is multiplicity for all θ in $[\theta_2, \theta_2 + \varepsilon)$ and/or $[\theta_2, \theta_2 - \varepsilon)$ for some $\varepsilon > 0$, violating multiplicity. Thus $[P(\theta_2), \min\{P(\theta_1), P(\theta_3)\}] \subseteq \bar{P}$. Suppose instead that $p' \in [\min\{P(\theta_1), P(\theta_3)\}, \max\{P(\theta_1), P(\theta_3)\}]$. Then by Lemma 5, either part (i) or part (ii), there is multiplicity in one of θ_1, θ_3 . Then there is multiplicity on a positive measure set, since these are interior.

Assume that $P(\theta_3) \geq P(\theta_1)$ (symmetric argument for reverse inequality). Suppose there exists $\theta' \geq \theta_2$ such that $\theta' \in \theta(P(\theta_1))$. Note that R weakly increasing in θ implies that $\{\theta \in \Theta : R(a, \theta) = p\}$ is convex for all a , so $\theta_2 \in \theta(P(\theta_1))$. Thus if such a θ' exists there will be multiplicity in state θ_2 , and, by the same argument as above, there will be multiplicity for a positive measure of states.

It remains to show that the existence of such a θ' is implied by our assumptions. Suppose instead that $\theta(P(\theta_1)) \subseteq [\underline{\theta}, \theta_2)$. We will show that this implies that there exists

$p > P(\theta_1)$ such that $\theta_2 \in \theta(p')$, so there is multiplicity in θ_2 , and as before this will imply multiplicity for a positive measure of states. Suppose no such p' exists. Let $\tilde{p} = \sup\{p \in [P(\theta_1), P(\theta_3)] : \max \theta(p) < \theta_2\}$, which is well defined by Lemma 22. Since $\theta(\tilde{p})$ is convex, the assumption that no such p' exists implies that either $\max \theta(\tilde{p}) < \theta_2$ or $\min \theta(\tilde{p}) > \theta_2$. Then we have a violation of upper hemicontinuity at \tilde{p} . Thus there exists $p > P(\theta_1)$ such that $\theta_2 \in \theta(p')$, as desired.

Now, for the extreme states we want to see that the monotonicity is maintained. Let θ be interior, then $p(\bar{\theta}) = R(M(p(\bar{\theta})), \bar{\theta}) > R(M(p(\bar{\theta})), \theta)$ and $R(M(p(\bar{\theta})), \theta) > \underline{p} := \inf\{\bar{P}\}$. So there is an equilibrium price $\tilde{p} \in (p(\bar{\theta}), \underline{p})$. If $p(\bar{\theta}) > \bar{p} := \sup\{\bar{P}\}$, then the previous does not imply multiplicity, but if $p(\bar{\theta}) < \bar{p}$, it does for all θ with associated prices in $(p(\bar{\theta}), \bar{p})$. The same argument holds to prove that $p(\underline{\theta}) \leq \underline{p}$.

It remains to show that the result holds for discontinuous $M \in \mathcal{M}$. For P to be non-monotone without violating robustness to multiplicity it must be that P is discontinuous. Suppose P is discontinuous at θ , and assume without loss of generality that P is decreasing below θ and left-continuous at θ . Let $\theta' > \theta$ be such that $P(\theta') > P(\theta)$. Then it must be that M is discontinuous on $(P(\theta), P(\theta'))$, otherwise the argument above for continuous M would apply. Let $\underline{p} = \inf\{p \in (P(\theta), P(\theta')) : M \text{ is discontinuous at } p\}$. Since M is continuous on $(P(\theta), \underline{p})$, the argument for continuous M implies that $\theta_M(p) = \theta$ for all $p \in (P(\theta), \underline{p})$. But then, by the definition of \mathcal{M} , M must be continuous in a neighborhood of \underline{p} , which is a contradiction. \square

A.6.3. Theorem 2

PROOF. First, note that under the assumption that R is strictly increasing in θ , Theorem 1 implies that P must be *strictly* monotone; otherwise measurability would be violated. Given this, to show necessity of i we first show that continuity of Q is necessary for CUI.

Suppose first that Q is discontinuous at an interior state θ' . If $P(\theta) := R(Q(\theta), \theta)$ is not also discontinuous at θ' then there can be no $M \in \mathcal{M}$ that implements Q . Assume P is discontinuous at θ . Suppose without loss of generality that P is increasing. Under strict monotonicity of $\theta \mapsto R(a, \theta)$, we have $|\theta_M(p)| \leq 1$ for all p .

Assume first that M is continuous. Thus Lemma 24 implies that $(\lim_{\theta \nearrow \theta'} P(\theta), \lim_{\theta \searrow \theta'} P(\theta)) \subseteq \bar{P}$. If $\theta_M(p) \neq \theta'$ for some $p \in (\lim_{\theta \nearrow \theta'} P(\theta), \lim_{\theta \searrow \theta'} P(\theta))$ then there will be multiplicity in some state $\theta'' \neq \theta'$ by Lemma 28. But if $\theta_M(p) = \theta'$ for some $p \in (\lim_{\theta \nearrow \theta'} P(\theta), \lim_{\theta \searrow \theta'} P(\theta))$ then there is multiplicity in state θ' .

If M is discontinuous it must still be continuous in a neighborhood N of $P(\theta')$. Then the same argument implies that $\theta_M(p) = \theta'$ for all $p \in N \cap (\lim_{\theta \nearrow \theta'} P(\theta), \lim_{\theta \searrow \theta'} P(\theta))$.

Suppose that P is decreasing; we want to show that local upper monotonicity is satisfied at $(Q(\underline{\theta}, \underline{\theta}))$ and local lower monotonicity is satisfied at $(Q(\bar{\theta}), \bar{\theta})$. We prove the former here, the latter is symmetric. Suppose local upper monotonicity is not satisfied at $(Q(\underline{\theta}, \underline{\theta}))$. Since M must be continuous in a neighborhood of $P(\underline{\theta})$ (by definition of \tilde{M}) the violation of local upper monotonicity implies that for any $\varepsilon > 0$ there exists a $p \in (P(\underline{\theta}, \underline{\theta}) + \varepsilon)$ such that $R(M(p), \underline{\theta}) \leq p$. Since $\theta \mapsto R(a, \theta)$ is strictly increasing and R is continuous, for ε small enough we will also have $R(M(p), \bar{\theta}) > p$. But then by

continuity of R there exists θ such that $R(M(p), \theta) = p$. Since $p > P(\underline{\theta})$ this implies that there is multiplicity at θ .

Now for sufficiency of *ii*. The argument immediately preceding Theorem 2 implies that continuity of Q and strict monotonicity of P are sufficient to rule out multiplicity involving prices in $P(\Theta)$. If P is increasing then define $M(p) = M(P(\bar{\theta}))$ for all $p > P(\bar{\theta})$ and $M(p) = M(P(\underline{\theta}))$ for all $p < P(\underline{\theta})$. Then there can be no equilibria involving prices in $\mathcal{P} \setminus P(\Theta)$. When P is decreasing let m, ε satisfy the conditions of local upper monotonicity at $Q((\underline{\theta}), \theta)$. Then defining $M = m$ on $(P(\underline{\theta}), P(\underline{\theta}) + \varepsilon)$ guarantees that there is no equilibrium at any such prices. Therefore $M \in \mathcal{M}$ can be discontinuous at $P(\underline{\theta}) + \varepsilon$. Defining $M(p) = Q(\bar{\theta})$ for all $p > P(\underline{\theta}) + \varepsilon$ guarantees that there can be no equilibria involving prices above $P(\underline{\theta})$. A symmetric construction is used to guarantee that there are no equilibrium prices below $P(\bar{\theta})$. \square

A.6.4. Proposition 9

PROOF. First for necessity. Theorem 1 implies that P must be weakly monotone. If it is not strictly monotone then it will violate measurability, given that $R(a, \cdot)$ is strictly monotone. The necessity of the local upper/lower monotonicity conditions follows from the same argument as Theorem 2. This proves necessity of *i*.

To show necessity of *ii*., suppose Q has a discontinuity at an interior state θ' that is not bridgeable. Given that we have established *i*, assume without loss of generality that P is strictly increasing and left continuous. We first show that M must be continuous on $(P(\theta'), \lim_{\theta \searrow \theta'} P(\theta))$. Suppose not, and let $\underline{p} = \inf\{p \geq P(\theta') : M \text{ is discontinuous at } p\}$. By definition of \mathcal{M} , it must be that $\underline{p} > P(\theta')$. Under strict monotonicity of $\theta \mapsto R(a, \theta)$,

we have $|\theta_M(p)| \leq 1$. If $\theta_M(p) \neq \theta'$ for some $p \in (P(\theta'), \underline{p})$ then there will by multiplicity by Lemma 28. But then, by definition of \mathcal{M} , M must be continuous on a neighborhood of \underline{p} , which contradicts the definition of \underline{p} .

We have established that $(\lim_{\theta \nearrow \theta'} P(\theta), \lim_{\theta \searrow \theta'} P(\theta))$. If the discontinuity at θ' is not bridgeable then there is no continuous M such that the following three conditions hold: a) $\lim_{\theta \nearrow \theta'} M(P(\theta)) = \lim_{\theta \nearrow \theta'} Q(\theta)$, b) $\lim_{\theta \searrow \theta'} M(P(\theta)) = \lim_{\theta \searrow \theta'} Q(\theta)$ and, c) $\theta_M(p) = \theta'$ for all $p \in (\lim_{\theta \nearrow \theta'} P(\theta), \lim_{\theta \searrow \theta'} P(\theta))$. To see this, notice that any such M would constitute a monotone path from $\lim_{\theta \nearrow \theta'} Q(\theta)$ to $\lim_{\theta \searrow \theta'} Q(\theta)$.

Now for necessity of *iii* and *iv*. Note that this conditions will be satisfied if the discontinuities are bridgeable. If they are not bridgeable then the argument for necessity of the local upper/lower monotonicity conditions is the same as that given for Theorem 2.

Now for sufficiency. Assume without loss that P is strictly increasing. Define $M(p) = Q(P^{-1}(p))$, which is well defined on $P(\Theta)$ by *i*. Moreover M is continuous on $P(\Theta)$ under *ii*. Q is continuous at any interior state θ at which P is continuous. It remains to define M on $\mathcal{P} \setminus P(\Theta)$. This is done as in Theorem 2. \square

A.6.5. Corollary 4

PROOF. First necessity. If Q is not strictly monotone then for ε small enough there will be no ε -approximation that is strictly monotone. Thus by Theorem 1 there are no CWUI ε -approximations. Suppose *ii* is violated. Since P is strictly monotone it can have at most countably many discontinuities. Thus Q must have a positive measure δ of degenerate discontinuities. If $\varepsilon < \delta$ then for any ε -approximation of Q there will be a

degenerate discontinuity that is outside of the set of states for which $Q' \neq Q$. But then Q' has a degenerate discontinuity, and so is not CWUI.

Now for sufficiency. Given Corollary 3 we need only show that Q can be approximated around all degenerate discontinuities. This follows immediately from the definition of correctable. \square

A.6.6. Proposition 10

PROOF. Corollary 3 implies that P is strictly monotone and that whenever Q is discontinuous so is P . Thus P (and Q) can have at most countably many discontinuities. The proposition will follow if we can show that for any $\varepsilon > 0$ and any θ^* at which Q is discontinuous, we can continuously approximate Q around θ^* without changing Q outside of $(\theta^* - \varepsilon, \theta^* + \varepsilon)$.

Since R is continuous and P is discontinuous at θ^* , there exists $\delta < \varepsilon$ and $\theta' \in (\theta^*, \theta^* + \delta)$ such that $Q(\theta^*), Q(\theta')$ and δ satisfy the conditions of continuous bridgeability. Thus there exists a continuous Q' on $[\theta^*, \theta^* + \delta]$ such that $Q'(\theta^*) = \lim_{\theta \nearrow \theta^*} Q(\theta)$, $Q'(\theta^* + \delta) = Q(\theta^* + \delta)$, and $R(Q'(\theta), \theta)$ is strictly increasing on $[\theta^*, \theta^* + \delta]$. Since ε was arbitrary, this gives the desired approximation. \square

A.6.7. Lemma 24

Lemma 24. *Assume R is weakly increasing in θ . For any continuous M that is robust to multiplicity, let p_1, p_2 be prices such that $\theta_M(p_1)$ and $\theta_M(p_2)$ are contained in the interior of θ . Then*

$$[\min\{p_1, p_2\}, \max\{p_1, p_2\}] \in \bar{P}.$$

PROOF. By Theorem 1, the price function P is monotone, so without loss of generality assume that it is increasing, and let $p_2 > p_1$. Assume towards a contradiction that there exists $p \in (p_1, p_2)$ such that $\theta_M(p) = \emptyset$. By Lemma 25 p is either type H or type L. Suppose it is type L, i.e. $R(M(p), \underline{\theta}) - p > 0$. Since $\theta_M(p_1) \neq \emptyset$, it must be that $R(M(p_1), \underline{\theta}) - p_1 \leq 0$. Moreover, since $\underline{\theta} \notin \theta_M(p_1)$ by assumption, the inequality is strict: $R(M(p_1), \underline{\theta}) - p_1 < 0$. Then by continuity there exists $p' \in (p_1, p)$ such that $R(M(p'), \underline{\theta}) - p' = 0$. Let $\theta_1 = \min \theta_M(p_1)$, which exists by Lemma 22 (by assumption $\theta_1 > \underline{\theta}$). Since P is increasing, $p' > p_1 > P(\theta)$ for all $\theta \in [\underline{\theta}, \theta_1)$. Then by Lemma 28 there is multiplicity for all states in $\theta \in [\underline{\theta}, \theta_1)$, which is a contradiction. If p is type H then the proof is symmetric, using p_2 rather than p_1 . \square

A.6.8. Proof of Proposition 11

PROOF. This essentially follows from Proposition 9. The only modifications are the following. Condition *iv* is clearly necessary and sufficient for there to be no monotonicity involving actions in $Q(\Theta)$. The modification of *ii* from bridgeable to properly bridgeable is necessary and sufficient for there to be no multiplicity involving actions not in $Q(\Theta)$. There is no need to modify condition *iii* since it guarantees existence of that are all type H (at $\bar{\theta}$) or type L (at $\underline{\theta}$), and thus involve no multiplicity. \square

A.6.9. Implementable price functions

In some cases the principal may not care directly about the actions they take, only about the price that they induce. In this section we ask the following question: for which price

functions there exists an action function such that (Q, P) is CWUI. We call such a P CWUI.

Definition 28. *A price function $P : \Theta \rightarrow \mathbb{R}$ is in range if for each $\theta \in \Theta$, $P(\theta) \in R(\mathcal{A}, \theta)$.*

Proposition 25. *Under strict monotonicity of R in θ , a price function is implementable if and only if it is in range and an injection.*

PROOF. in the Appendix A.6.9.1 □

We will call intersection states the ones where there is at least two different actions a_1, a_2 with $R(a_1, \theta) = R(a_2, \theta)$. Let Θ_I be the set of such states. We will make the following extra assumptions on R :

Mixture continuity. For any θ , any a'', a' such that $R(a'', \theta) > R(a', \theta)$, and any $p \in [R(a', \theta), R(a'', \theta)]$, there is a unique α such that $R(\alpha a' + (1 - \alpha)a'', \theta) = p$.

Isolated intersections. For every $\theta \in \Theta_I$, there exists an $\epsilon > 0$ such that $B_\epsilon(\theta) \cap \Theta_I = \{\theta\}$.

Intersection smoothness. $R_2(a, \theta)$ exists for every intersection state θ and a that puts weight only on intersecting actions for that state.

Definition 29. *A price function $P : \Theta \rightarrow \mathbb{R}$ satisfies the **kink's condition** iff there exist C^1 functions \bar{P} and \underline{P} in range and such that $\bar{P}(\theta) \geq P(\theta) \geq \underline{P}(\theta)$.*

The kink's condition effectively means that every kink of P in the upper envelope of $R(a, \theta)$ is concave, and every kink in the lower envelope is convex. Moreover, the kink's

condition implies that if there is a θ such that $R(\mathcal{A}, \theta)$ is a singleton, the price function has to be differentiable at θ .

Proposition 26. *Under strict monotonicity of R in θ , mixture continuity, isolated intersections, and intersection smoothness, a price function is CWUI if and only if it is in range, strictly monotone, and satisfies the kink's condition.*

PROOF. in the ?? A.6.9.2. □

There are two primary components of the proof of Proposition 26. The first is that \bar{P} is convex. The second is to show that if $\theta(p)$ is non-monotone then there will be multiplicity. Identification is used to prove both parts of the proposition, but it is not necessary for either. One simple relaxation under which the result is preserved is to allow for actions with constant payoffs.

Weak identification. R is weakly increasing in θ . Moreover, if $R(a, \cdot)$ is not strictly increasing then it is constant.

Proposition 27. *Under weak identification, a pair (Q, P) is implementable if and only if $Q(P^{-1}(p))$ is a singleton for all $p \in P(\Theta)$.*

Proposition 28. *Under R weakly increasing in θ , a price function is CWUI if and only if it is in range and weakly monotone and whenever it is flat at a price p , it is so for the whole set $\theta_M(p)$.*

PROOF. in the ?? A.6.9.3. □

A.6.9.1. Proof of Proposition 25.

PROOF. (\Rightarrow): suppose not an injection. There are θ and θ' with $P(\theta) = P(\theta')$. By identification, $R(Q(\theta), \theta) \neq R(Q(\theta'), \theta')$, which by rational expectations means that $P(\theta) \neq P(\theta')$, a contradiction. If not in range, then there exist a $\theta \in \Theta$ such that $P(\theta) \notin R(\mathcal{A}, \theta)$, i.e. there is no $a \in \mathcal{A}$ such that $R(a, \theta) = P(\theta)$, so $R(Q(\theta), \theta) \neq P(\theta)$, violating rational expectations.

(\Leftarrow): Since $P(\theta)$ is in range, for each $\theta \in \Theta$ there exists a a with $R(a, \theta) = P(\theta)$. let's define $Q(\theta)$ by a selection in the rational expectations condition: $R(Q(\theta), \theta) = P(\theta)$. Measurability is satisfied trivially since $P(\theta) \neq P(\theta')$ for all $\theta \neq \theta'$. \square

A.6.9.2. Proof of Proposition 26.

PROOF. (\Rightarrow): Take M that implements P . For all $p \in R(\mathcal{A}, \Theta)$ there is at most a unique $\theta \in \Theta$ that satisfies $R(M(p), \theta) = p$. Otherwise identification would be violated. This defines a function $\theta(p)$.

Let \bar{P} be the set of all prices for which there is an interior solution. We want to show that \bar{P} is convex. Pick $p, p' \in \bar{P}$ and $\alpha \in (0, 1)$ we want to see that $p_\alpha := \alpha p + (1 - \alpha)p' \in \bar{P}$. Let θ and θ' the associated states of p and p' . Continuity of R plus identification imply strict monotonicity of R in θ and for all a . Assume without loss that $\theta' > \theta$.

We will prove that

$$(A.1) \quad R(M(p_\alpha), \theta) \leq p_\alpha \leq R(M(p_\alpha), \theta')$$

Consider a violation of the second inequality. If $p_\alpha > R(M(p_\alpha), \theta')$ notice that also, $p = R(M(p), \theta) < R(M(p), \theta')$. Therefore, we have

$$p_\alpha - R(M(p_\alpha), \theta') > 0 \quad \text{and} \quad p - R(M(p), \theta') < 0$$

By continuity and since $\theta' \in \Theta^\circ$, there exists an $\bar{\varepsilon} > 0$ such that for all $\tilde{\theta} \in B_{\bar{\varepsilon}}(\theta')$

$$p_\alpha - R(M(p_\alpha), \tilde{\theta}) > 0 \quad \text{and} \quad p - R(M(p), \tilde{\theta}) < 0$$

By continuity there is a $p_1 \in (p_\alpha, p)$ with $p_1 - R(M(p_1), \theta') = 0$. But for $0 < \epsilon < \bar{\varepsilon}$ we have $p' = R(M(p'), \theta') < R(M(p'), \theta' - \epsilon)$. There exists a $p_2 \in (p_\alpha, p')$ such that $R(M(p_2), \theta' - \epsilon) = p_2$. $p_1 \neq p_2$, so there is multiplicity in a set of states $[\theta', \theta' + \epsilon)$. With a similar logic we can rule out $p_\alpha \leq R(M(p_\alpha), \theta)$.

Finally, by continuity of π in θ and using Equation (A.1), there is a $\hat{\theta}$ in (θ, θ') such that $p_\alpha - R(M(p_\alpha), \hat{\theta}) = 0$, therefore $p_\alpha \in \bar{P}$.

The function $\theta(p)$ is continuous in the set \bar{P} . However we could have discontinuities for the two prices that are associated with the extreme states $\bar{\theta}$ and $\underline{\theta}$.

We show now that $\theta(p)$ is monotone in \bar{P} . Suppose that is not, i.e. there are prices $p_l < p_m < p_h$ such that either $\theta(p_m) < \min\{\theta(p_l), \theta(p_h)\}$ or $\theta(p_m) > \max\{\theta(p_l), \theta(p_h)\}$. Suppose the first (the symmetric argument holds for the other case). Then for all $\theta \in (\theta(p_m), \min\{\theta(p_l), \theta(p_h)\})$ and by continuity there are prices $p_\theta^1 \in (p_l, p_m)$ and $p_\theta^2 \in (p_m, p_h)$ with $\theta(p_\theta^1) = \theta(p_\theta^2)$. This violates multiplicity.

We can invert $\theta(p)$ in \bar{P} . The only problem is when $\theta(p)$ is flat, but any selection would give us that the inverse is strictly monotone.

Now, for the extreme states we want to see that the monotonicity is maintained. Let θ be interior, then $p(\bar{\theta}) = R(M(p(\bar{\theta})), \bar{\theta}) > R(M(p(\bar{\theta})), \theta)$ and $R(M(p(\bar{\theta})), \theta) > \underline{p} := \inf\{\bar{P}\}$. So there is an equilibrium price $\tilde{p} \in (p(\bar{\theta}), \underline{p})$. If $p(\bar{\theta}) > \bar{p} := \sup\{\bar{P}\}$, then the previous does not imply multiplicity, but if $p(\bar{\theta}) < \bar{p}$ It does for all θ with associated prices in $(p(\bar{\theta}), \bar{p})$. The same argument holds to prove that $p(\underline{\theta}) \leq \underline{p}$.

(\Leftarrow): P is strictly monotone and bounded so there is a countable number of discontinuities. Fill those to get a continuous and monotone $\theta(p) := \sup\{\theta : P(\theta) < p\}$.

Let $\bar{M} : P(\Theta) \rightrightarrows \mathcal{A}$ be the set of actions that give price p at the corresponding state i.e. $a \in \bar{M}(p)$ if and only if $R(a, \theta(p)) = p$.

If p is not an intersection price, then \bar{M} is LHC at p . Therefore in a ball around p there is a continuous selection. If p is a interior intersection, then we can consider the set of actions that are not involved in the intersection and select a continuous M . \square

A.6.9.3. Proof of Proposition 28.

PROOF. Assume without loss of generality that that $R(a, \cdot)$ is increasing for all a . Lemmas 26 and 27 and Corollary 7 apply.

Let $\theta_1 < \theta_2 < \theta_3$ be interior, and suppose $P(\theta_1) > P(\theta_2)$ and $P(\theta_3) > P(\theta_2)$ (the other type of non-monotonicity is dealt with symmetrically). We first want to show that $[P(\theta_2), \max\{P(\theta_1), P(\theta_3)\}] \subseteq \bar{P}$. Suppose not, so there is some $p' \in [P(\theta_2), \max\{P(\theta_1), P(\theta_3)\}]$ such that $\theta(p') = \emptyset$. Then p' is either Type L or Type H. Assume it is Type H (symmetric argument using Lemma 26 if it is Type L). First, assume $p' \in [P(\theta_2), \min\{P(\theta_1), P(\theta_3)\}]$. By Lemma 27, part (i) there is a $p > p'$ such that $\theta_2 \in \theta(p')$. Thus there is multiplicity in state θ_2 . Then by continuity of M and R , and R weakly increasing in θ ,

there is multiplicity for all θ in $[\theta_2, \theta_2 + \varepsilon)$ and/or $[\theta_2, \theta_2 - \varepsilon)$ for some $\varepsilon > 0$, violating multiplicity. Thus $[P(\theta_2), \min\{P(\theta_1), P(\theta_3)\}] \subseteq \bar{P}$. Suppose instead that $p' \in [\min\{P(\theta_1), P(\theta_3)\}, \max\{P(\theta_1), P(\theta_3)\}]$. Then by Lemma 5, either part (i) or part (ii), there is multiplicity in one of θ_1, θ_3 . Then there is multiplicity on a positive measure set, since these are interior.

Assume that $P(\theta_3) \geq P(\theta_1)$ (symmetric argument for reverse inequality). Suppose there exists $\theta' \geq \theta_2$ such that $\theta' \in \theta(P(\theta_1))$. Note that R weakly increasing in θ implies that $\{\theta \in \Theta : R(a, \theta) = p\}$ is convex for all a , so $\theta_2 \in \theta(P(\theta_1))$. Thus if such a θ' exists there will be multiplicity in state θ_2 , and, by the same argument as above, there will be multiplicity for a positive measure of states.

It remains to show that the existence of such a θ' is implied by our assumptions. Suppose instead that $\theta(P(\theta_1)) \subseteq [\underline{\theta}, \theta_2)$. We will show that this implies that there exists $p > P(\theta_1)$ such that $\theta_2 \in \theta(p')$, so there is multiplicity in θ_2 , and as before this will imply multiplicity for a positive measure of states. Suppose no such p' exists. Let $\tilde{p} = \sup\{p \in [P(\theta_1), P(\theta_3)] : \max \theta(p) < \theta_2\}$, which is well defined by Lemma 22. Since $\theta(\tilde{p})$ is convex, the assumption that no such p' exists implies that either $\max \theta(\tilde{p}) < \theta_2$ or $\min \theta(\tilde{p}) > \theta_2$. Then we have a violation of upper hemicontinuity at \tilde{p} . Thus there exists $p > P(\theta_1)$ such that $\theta_2 \in \theta(p')$, as desired.

Now, for the extreme states we want to see that the monotonicity is maintained. Let θ be interior, then $p(\bar{\theta}) = R(M(p(\bar{\theta})), \bar{\theta}) > R(M(p(\bar{\theta})), \theta)$ and $R(M(p(\bar{\theta})), \theta) > \underline{p} := \inf\{\bar{P}\}$. So there is an equilibrium price $\tilde{p} \in (p(\bar{\theta}), \underline{p})$. If $p(\bar{\theta}) > \bar{p} := \sup\{\bar{P}\}$, then the previous does not imply multiplicity, but if $p(\bar{\theta}) < \bar{p}$, it does for all θ with associated prices in $(p(\bar{\theta}), \bar{p})$. The same argument holds to prove that $p(\underline{\theta}) \leq \underline{p}$. \square

Strictly monotone price functions may have jump discontinuities. However such price functions can always be approximated arbitrarily well by continuous and strictly increasing functions.

A.6.10. Intermediate results

Lemma 25. *Fix a continuous M . Assume $R(a, \cdot)$ is (weakly) increasing for all a (the same holds if decreasing, with $\underline{\theta}$ and $\bar{\theta}$ switched). Then each p such that $\theta(p) = \emptyset$ is of one and only one of the following two types:*

- *Type L: $R(M(p), \underline{\theta}) > p$.*
- *Type H: $R(M(p), \bar{\theta}) < p$.*

PROOF. Since $\theta \mapsto R(M(p), \theta)$ is increasing p cannot be of both types. If p is of neither then by continuity there exists a $\theta \in [\underline{\theta}, \bar{\theta}]$ such that $R(M(p), \theta) = p$. But then $\theta(p)$ is not empty. \square

Corollary 7. *The set of prices $\{p : \theta(p) = \emptyset\}$ is open.*

Lemma 26. *Assume $R(a, \cdot)$ is (weakly) increasing for all a (the same holds if decreasing, with $\underline{\theta}$ and $\bar{\theta}$ switched) and M is continuous. Let p be Type L and $\theta'' > \theta'$.*

- i. *If there exists $p'' > p$ such that $\theta'' \in \theta(p'')$ then there exists $p' \in (p, p'']$ such that $\theta' \in \theta(p')$.*
- ii. *If there exists $p'' < p$ such that $\theta'' \in \theta(p'')$ then there exists $p' \in [p'', p)$ such that $\theta' \in \theta(p')$.*

PROOF. We will prove (i), the proof for (ii) is symmetric. $R(M(p), \underline{\theta}) > p$ since p is type L. Moreover, under monotonicity

$$p'' = R(M(p''), \theta'') \geq R(M(p''), \theta') \geq R(M(p''), \underline{\theta}).$$

Then by continuity of R and M , there exists $\underline{p} \in (p, p'']$ such that $R(M(\underline{p}), \underline{\theta}) = \underline{p}$. By monotonicity we have $R(M(\underline{p}), \theta') \geq R(M(\underline{p}), \underline{\theta}) = \underline{p}$ and $p'' = R(M(p''), \theta'') \geq R(M(p''), \theta')$. Then by continuity of R, M there exists $p' \in [\underline{p}, p'']$ such that $R(M(p'), \theta') = p'$, so $\theta' \in \theta(p')$ as desired. \square

Lemma 27. *Assume $R(a, \cdot)$ is (weakly) increasing for all a (the same holds if decreasing, with $\underline{\theta}$ and $\bar{\theta}$ switched) and M is continuous. Let p be type H and $\theta'' > \theta'$.*

- i. *If there exists $p' > p$ such that $\theta' \in \theta(p')$ then there exists $p'' \in (p, p']$ such that $\theta'' \in \theta(p'')$.*
- ii. *If there exists $p' < p$ such that $\theta' \in \theta(p')$ then there exists $p'' \in [p', p)$ such that $\theta'' \in \theta(p'')$.*

PROOF. Analogous to that of Lemma 26. \square

Lemma 28. *(Generalized intermediate value theorem). Let $F : [0, 1] \rightarrow [0, 1]$ be a compact and convex valued, upper hemicontinuous correspondence. Let $p_1 < p_2$. Let $y_1 \in F(p_1)$ and $y_2 \in F(p_2)$. Then for any $\tilde{y} \in (\min\{y_1, y_2\}, \max\{y_1, y_2\})$ there exists $p \in [p_1, p_2]$ such that $\tilde{y} \in F(p)$.*

PROOF. Define $p^* := \sup\{p \in [p_1, p_2) : \max F(p) < \tilde{y}\}$. If $p^* = p_1$ then $\max F(p) \geq \tilde{y}$ for all $p \in (p_1, p_2)$. Assume none of these hold with equality (otherwise we are done).

Then if $\min F(p) \leq \tilde{y}$ for some $p \in (p_1, p_2]$ then we are done, by convexity of F . So suppose $\min F(p) > \tilde{y}$ for all $p \in (p_1, p_2]$. Then $\tilde{y} \in F(p_1)$: otherwise, by convexity of $F(p_1)$, we have $\max F(p_1) < \tilde{y}$, which violates upper hemicontinuity. Thus we are done if $p^* = p_1$.

Suppose instead that $p^* = p_2$. If $\min F(p_2) \leq \tilde{y}$ then we are done, by convexity of $F(p)$. Suppose $\min F(p_2) > \tilde{y}$. Then by the definition of p^* , it must be that for any $\varepsilon > 0$ there exists $p \in (p_2 - \varepsilon, p_2)$ such that $\max F(p) < \tilde{y}$. But this violates upper hemicontinuity of F at p_2 . Thus we are done if $p^* = p_2$.

It only remains to address the case of $p^* \in (p_1, p_2)$. It must be that $\max F(p^*) \geq \tilde{y}$: if not then by upper hemicontinuity there exists $\varepsilon > 0$ such that $\max F(p) < \tilde{y}$ for all $p \in [p^*, p^* + \varepsilon)$, but this would contradict the definition of p^* . If $\min F(p^*) \leq \tilde{y}$ then we are done, by convexity. So suppose $\min F(p^*) > \tilde{y}$. Then by upper hemicontinuity there exists $\varepsilon > 0$ such that $\min F(p) > \tilde{y}$ for all $p \in (p^* - \varepsilon, p^*]$. But this contradicts the definition of p^* . \square

A.6.10.1. Proof of Proposition 16.

PROOF. *Claim 0.* For any $\theta' \in (\underline{\theta}, \bar{\theta})$ and p' be such that $\theta' \in \theta_M(p')$, there exist p'' such that $\theta_M(p'') \cap \{\underline{\theta}, \bar{\theta}\} \neq \emptyset$, $\theta_M(p) \neq \emptyset$ for all $p \in (\min\{p', p''\}, \max\{p', p''\})$ and M is continuous on $(\min\{p', p''\}, \max\{p', p''\})$ (when this interval is non-empty).

Let $\theta' \in (\underline{\theta}, \bar{\theta})$ be arbitrary, and let p' be such that $\theta' \in \theta_M(p')$. If $\{p \leq p' : \theta_M(p) = \emptyset\}$ is empty then $p'' = \arg \min_{a \in \mathcal{A}} R(a, \underline{\theta})$ satisfies the conditions of the claim. Similarly, if $\{p \geq p' : \theta_M(p) = \emptyset\}$ is empty then $p'' = \arg \max_{a \in \mathcal{A}} R(a, \bar{\theta})$ satisfies the conditions of the claim. Assume that $\{p \leq p' : \theta_M(p) = \emptyset\} \neq \emptyset$ and $\{p \geq p' : \theta_M(p) = \emptyset\} \neq \emptyset$. Let $\underline{p} = \sup\{p \leq p' : \theta_M(p) = \emptyset\}$ and $\bar{p} = \inf\{p \geq p' : \theta_M(p) = \emptyset\}$. Since $M \in \mathcal{M}$, we

have $\underline{p} < p' < \bar{p}$. Since M must be continuous on (\underline{p}, \bar{p}) , we have $\theta_M(\underline{p}) \cap \{\underline{\theta}, \bar{\theta}\} \neq \emptyset$ and $\theta_M(\bar{p}) \cap \{\underline{\theta}, \bar{\theta}\} \neq \emptyset$. This proves Claim 0.

Claim 1. Let $\theta' \in (\underline{\theta}, \bar{\theta})$ and p' such that $\theta' \in \theta_M(p')$. Let p'' be such that $\theta_M(p) \neq \emptyset$ for all $p \in (\min\{p', p''\}, \max\{p', p''\})$ and M is continuous on $(\min\{p', p''\}, \max\{p', p''\})$ (when this interval is non-empty). Then if $\underline{\theta} \in \theta_M(p'')$ and $p'' \leq p'$ ($p'' \geq p'$) there exists an equilibrium with a price function that is increasing (decreasing) on $[\underline{\theta}, \theta']$. Similarly, if $\bar{\theta} \in \theta_M(p'')$ and $p'' \geq p'$ ($p'' \leq p'$) there exists an equilibrium with a price function that is increasing (decreasing) on $[\theta', \bar{\theta}]$.

We will show the claim for $\bar{\theta} \in \theta_M(p'')$ and $p'' \geq p'$; all others cases are symmetric. For any θ , the set $\theta_M^{-1}(\theta)$ is compact: if $R(M(p), \theta) \neq p$ then this holds for all \tilde{p} in a neighborhood p , since $M \in \mathcal{M}$ is continuous around equilibrium prices. If $p' = p''$ then we are done: convexity of $\theta_M(p)$ (Lemma 23) implies that there is a constant, and thus monotone, equilibrium price function on $[\theta', \bar{\theta}]$. Assume instead that $p'' > p'$. If there exists $\theta^* \in (\theta', \bar{\theta})$ such that $p^* > p''$ for any $p^* \in \theta_M^{-1}(\theta^*)$ then there exists $\tilde{\theta} \in (\theta', \bar{\theta})$ such that $p'' \in \theta_M^{-1}(\tilde{\theta})$, by continuity of M on (p', p'') and Lemma 28. Then convexity of $\theta_M(p'')$ implies that we can construct a flat price function above $\tilde{\theta}$. Therefore assume no such θ^* exists. By a symmetric argument, we can assume that $\theta_M^{-1}(p) \cap [p', p''] \neq \emptyset$ for all $\theta \in [\theta', \bar{\theta}]$.

We want to construct an increasing equilibrium price function on $[\theta', \bar{\theta}]$. Consider an arbitrary price function \tilde{P} such that $\tilde{P}(\theta) \in \theta_M^{-1}(\theta) \cap [p', p'']$ for all $\theta \in [\theta', \bar{\theta}]$, $\tilde{P}(\underline{\theta}) = p'$, and $\tilde{P}(\bar{\theta}) = p''$. We will show that any violations of monotonicity can be ironed without leading to further violations.

Claim 1.2. Suppose $\tilde{P}(\theta_2) < \tilde{P}(\theta_1) < \tilde{P}(\theta_3)$ for $\bar{\theta} > \theta_3 > \theta_2 > \theta_1 >$. Then there exists $p \in \theta_M^{-1}(\theta_2) \cap [\tilde{P}(\theta_1), \tilde{P}(\theta_3)]$.

Claim 1.2 follows immediately from Lemma 28. This in turn shows that Claim 1 holds for $\bar{\theta} \in \theta_M(p'')$ and $p'' \geq p'$, which is what we wished to show.

Claim 0 and Claim 1 together imply the existence of a monotone price function. \square

A.6.10.2. Proof of Proposition 17.

PROOF. By Proposition 16, M admits an equilibrium with a monotone price function P . Let Q be the associated action function. For any state θ such that $r(Q(\theta), P(\theta))$ is non-degenerate, let $\hat{Q}(\theta') = R(Q(\theta), \theta)$ for all $\theta' \in r(Q(\theta), P(\theta))$. Clearly $\hat{P}(\theta) := R(\hat{Q}(\theta), \theta)$ will also be monotone, and (\hat{Q}, \hat{P}) is also implemented by M . It remains to show that M can be modified on $\mathcal{P} \setminus \hat{P}(\Theta)$ in order to implement (\hat{Q}, \hat{P}) uniquely. This follows from Proposition 11. Note that \hat{Q} will have no degenerate discontinuities since M is continuous on $\hat{P}(\Theta)$. \square

A.6.11. Proof of Lemma 10

PROOF. An equilibrium exists for any increasing M by Tarski's fixed point theorem. That the price function will be increasing follows from the fact that $a \mapsto R(a, \theta)$ is decreasing and $\theta \mapsto R(a, \theta)$ is increasing. If P is increasing and M is increasing, there will be no equilibrium involving prices above $P(\bar{\theta})$ or below $P(\underline{\theta})$.

We show that M can have no discontinuities on $[P(\underline{\theta}), P(\bar{\theta})]$, which implies that P is continuous. Suppose, towards a contradiction that there is a non-empty set D of discontinuities in this region, and let $p' = \inf D$. By definition of \mathcal{M} , $p' \in (P(\underline{\theta}), P(\bar{\theta}))$. Let $a' = \lim_{p \nearrow p'} M(p)$. For any $p \in (P(\underline{\theta}), p')$ and any $a \in (M(P(\underline{\theta}), a'))$ there exists

$\theta \in (\underline{\theta}, \bar{\theta})$ such that $R(a, \theta) = p$. This follows from the fact that $a \mapsto R(a, \theta)$ is decreasing. Then for any $p \in (P(\underline{\theta}), p')$ there exists θ such that $R(M(p), \theta) = p$, since M is increasing and continuous on $(P(\underline{\theta}), p')$. This contradicts the definition of p' . \square

A.6.12. Proof of Lemma 11

PROOF. Condition *i* is immediate. For *ii*, first note that for $p \in (R(\underline{a}, \underline{\theta}), R(\underline{a}, \bar{\theta}))$ it must be that $M(p) > \underline{a}$; if not then $R(M(p), \theta) = p$ for some $\theta \in (\underline{\theta}, \bar{\theta})$. Suppose there is no discontinuity on $(P(\underline{\theta}), R(\underline{a}, \underline{\theta})]$. Then M must be decreasing over this domain to prevent multiplicity, and $\lim_{p \searrow R(\underline{a}, \underline{\theta})} M(p) = \underline{a}$. But for $p \in (R(\underline{a}, \underline{\theta}), R(\underline{a}, \bar{\theta}))$ it must be that $M(p) > \underline{a}$, so there must be a discontinuity. A symmetric argument applies to $(R(\bar{a}, \bar{\theta}), P(\bar{a})]$

Conditions *iii* and *iv* follow from a similar argument. Define \bar{p} by $\bar{p} = \sup\{p : M \text{ is decreasing on } (P(\bar{\theta}), \bar{p})\}$. The argument above implies that $\bar{p} \leq R(\underline{a}, \underline{\theta})$. This implies *iii*. A symmetric argument implies *iv*. \square

A.6.13. Proof of Section 2.5.1

Lemma 29. *Given a continuous function $F : X \times (0, 1) \rightarrow X$ on a compact, convex subset X of a Euclidean space, define the function*

$$G(t) = \{x \in X : F(x, t) = x\}.$$

Then $G(t)$ is non-empty for all t (Brouwer's fixed point theorem). Moreover, if $G(t)$ is single valued then G is upper and lower hemicontinuous at t .

PROOF. Since $G(t)$ is single valued upper hemicontinuity implies lower hemicontinuity. We want to show that for any open neighborhood V of $G(t)$ there exists a neighborhood U of t such that $G(t') \subseteq V$ for all $t' \in U$.

Claim 1. For any open neighborhood V of $G(t)$ there exists a $\kappa > 0$ such that

$$|F(x, t) - x| > \kappa \quad \forall \quad x \in X \setminus V.$$

The proof of claim 1 is as follows. $X \setminus V$ is a closed subset of a compact set, and thus compact. The function $x \mapsto F(x, t) - x$ is continuous, so it attains its minimum on $X \setminus V$. Since $G(t)$ is unique and $G(t) \notin X \setminus V$, this minimum is strictly greater than zero, so the desired κ exists.

To complete the proof of Lemma 29, we need to show that there exists an open neighborhood U of t such

$$|F(x, t') - x| > \kappa \quad \forall \quad x \in X \setminus V, \quad t' \in U.$$

By continuity of $t' \mapsto F(x, t') - x$, for each x there exists a ε_x such that $|t' - t| < \varepsilon_x$ implies $|F(x, t') - x| > \kappa$. For each x , define $\ell(x, \varepsilon) = \min\{|F(x, t') - x| : |t' - t| \leq \varepsilon/2\}$, which exists by continuity of F and compactness of $|t' - t| \leq \varepsilon/2$. Define

$$B(x) = \{x' \in X : \ell(x', \varepsilon_x) > \kappa\}.$$

By continuity of $x \mapsto F(x, t') - x$, $B(x)$ contains an open neighborhood of x (Berge's maximum theorem). Let $\tilde{B}(x)$ be this open neighborhood. The set $\cup_{x \in X \setminus V} \tilde{B}(x)$ covers $X \setminus V$. Then by compactness of $X \setminus V$ there exists a finite sub-cover. Let u be the smallest

ε_x corresponding to an x such that $\tilde{B}(x)$ is in the finite sub-cover. Then $U = \{t' \in (0, 1) : |t' - t| < u\}$. \square

Proposition 29. *Given a continuous function $F : X \times \Theta \times (0, 1) \rightarrow X$ on a compact, convex subset X of a Euclidean space, define the function*

$$G(t, \theta) = \{x \in X : F(x, \theta, t) = x\}.$$

Then $G(t, \theta)$ is non-empty for all t, θ (Brouwer's fixed point theorem). Moreover, let S be any compact subset of Θ such that $G(t, \theta)$ is single valued for all $\theta \in S$. Then $t \mapsto G(t, \theta)$ is upper and lower hemicontinuous at t , uniformly over S .

PROOF. Since $G(t, \theta)$ is single valued on S it suffices to show upper hemicontinuity. Let $V(\theta)$ be an open neighborhood of $\theta \mapsto G(t, \theta)$ on S . Without loss of generality (since Θ is compact and $G(t, \theta)$ single valued on S), let $V(\theta) = \{x \in X : |G(t, \theta) - x| < \delta\}$ for some $\delta > 0$, or equivalently, $V(\theta) = \cup_{x \in G(t, \theta)} N_\delta(x)$. We want to show that there exists an open neighborhood U of t such that $t' \in U$ implies $G(t', \theta) \subseteq V(\theta)$ for all $\theta \in S$.

Claim 1. $X \setminus V(\theta)$ is upper and lower hemicontinuous on S .

The proof of Claim 1 is as follows. Since $G(t, \theta)$ is single valued,

$$X \setminus V(\theta) = X \setminus N_\delta(G(t, \theta))$$

where $N_\delta(x)$ is the open ball around x with radius δ . We first show upper hemicontinuity.

Let W be an open set containing $X \setminus V(\theta)$. Without loss of generality, let

$$W = X \setminus \bar{N}_{\delta-\rho}(G(t, \theta))$$

for some $\rho \in (0, \delta)$ where $\bar{N}_{\delta-\rho}(x)$ is the closed ball around x with radius $\delta - \rho$.² By Lemma 29, we know that $\theta \mapsto G(t, \theta)$ is upper and lower hemicontinuous at all $\theta \in S$. By upper hemicontinuity of $\theta \mapsto G(t, \theta)$ at θ , there exists an open neighborhood B of θ such that $\theta' \in B$ implies $|x - G(\theta, t)| < (\delta - \rho)/2$ for all $x \in G(\theta', t)$. Then $\bar{N}_{\delta-\rho}(G(t, \theta)) \subset \cup_{x \in G(t, \theta')} N_\delta(x) = V(\theta')$ for all $\theta' \in B$. Thus $V(\theta') \subset W$ for all $\theta' \in B$, which shows upper hemicontinuity.

For lower hemicontinuity, let $W \subset X$ be an open set intersecting $X \setminus V(\theta)$. This holds if and only if there exists $x' \in W$ such that $|x' - G(t, \theta)| > \delta$. By upper hemicontinuity of $\theta \mapsto G(t, \theta)$ at θ , there exists an open neighborhood B of θ such that $\theta' \in B$ implies $|x - G(\theta, t)| < (|x' - G(t, \theta)| - \delta)/2$ for all $x \in G(\theta', t)$. Then $\theta' \in B$ implies $|x' - x| > \delta$ for all $x \in G(t, \theta')$. Thus $x' \notin \cup_{x \in G(t, \theta')} N_\delta(x) = V(\theta')$, so $W \cap X \setminus V(\theta') \neq \emptyset$ for all $\theta' \in B$, which shows lower hemicontinuity. This completes the proof of Claim 1.

We know from Lemma 29 that for each $\theta \in S$ there exists $\varepsilon_\theta, \kappa_\theta > 0$ such that

$$(A.2) \quad |t' - t| < \varepsilon_\theta \implies |F(x, \theta, t') - x| > \kappa_\theta \quad \forall x \in X \setminus V(\theta).$$

Claim 2. For each $\theta \in S$ there exists an open neighborhood $B(\theta)$ of θ such that

$$\theta' \in B(\theta) \text{ and } |t' - t| < \varepsilon_\theta \implies |F(x, \theta, t') - x| > \kappa_\theta \quad \forall x \in X \setminus V(\theta'),$$

where $\varepsilon_\theta, \kappa_\theta$ satisfy (A.2).

The proof of this claim is as follows. Define

$$z(\theta, \varepsilon) := \min\{|F(x, \theta, t') - x| : |t' - t| \leq \varepsilon/2, x \in X \setminus V(\theta)\},$$

² W so defined is open in X , but not in the space of which X is a subset.

which is well defined by compactness of $X \setminus V(\theta)$. By Berge's maximum theorem and Claim 1, $\theta \mapsto z(\theta, \varepsilon)$ is continuous at any $\theta \in S$. By (A.2) we know that $z(\theta, \varepsilon_\theta) > \kappa_\theta$ for all $\theta \in S$. Then for any $\theta \in S$ there exists an open neighborhood $B(\theta)$ of θ such that $\theta' \in B(\theta)$ implies $z(\theta', \varepsilon_\theta) > \kappa_\theta$. This proves Claim 2.

To complete the proof of Proposition 29, note that $\cup_{\theta \in S} B(\theta)$ is an open cover of S . By compactness of S there exists a finite sub-cover. Let I be the set of $\theta \in S$ that index this sub-cover. Let $\varepsilon = \min\{\varepsilon_\theta : \theta \in I\}/2$. Then

$$|t' - t| < \varepsilon \implies |F(x, \theta, t') - x| > 0 \quad \forall x \in X \setminus V(\theta) \text{ and } \theta \in S.$$

Since $G(t', \theta)$ is non-empty for all t', θ we have that $|t' - t| < \varepsilon$ implies that for all θ , $G(t', \theta) \subseteq V(\theta)$, which shows upper hemicontinuity as desired. \square

A.6.13.1. Proof of Theorem 3.

PROOF. We first show that the result holds for continuous M . Since any discontinuity in $M \in \mathcal{M}$ must be bounded away from the set of equilibrium prices, this implies that the result also holds for all $M \in \mathcal{M}$.

Let M be continuous. Let $F(a, \theta, t) = M(R(a, \theta, t))$, where t continuously parameterizes the function R . Then F is continuous since M is continuous. Moreover, $G(t, \theta) = \tilde{Q}(\theta, t)$ will be single valued on all but a zero-measure set of states when M is weakly robust to multiplicity, and single valued everywhere when M is robust to multiplicity. Therefore for any $\varepsilon > 0$ we can find a compact set S such that $G(t, \theta)$ is single valued for all $\theta \in S$. When M is robust to multiplicity let $S = \Theta$. Then Proposition 29 applies, which gives the result. \square

A.6.13.2. Proof of Proposition 15.

PROOF. First, note that $|\mathcal{U}(R)| = 1$. It is non-empty since Lemma 29 and robustness to multiplicity imply that $\theta \mapsto \tilde{Q}(\theta, R)$ is a continuous function on all but a zero measure set of states, and is thus $\theta \mapsto u(\theta, Q(\theta))$ integrable for all $Q \in \tilde{Q}(\cdot, R)$. It is single valued since all $Q \in \tilde{Q}(\cdot, R)$ are the same on all but a zero measure set of states.

Since $|\mathcal{U}(R)| = 1$, upper hemicontinuity implies lower hemicontinuity, so it suffices to show the former. Thus we want to show that for any $\delta > 0$ there exists an open neighborhood $B \subseteq \mathcal{C}$ of R such that $R' \in B$ implies $\mathcal{U}(R') \subseteq (\mathcal{U}(R) - \delta, \mathcal{U}(R) + \delta)$. Since the set of states at which $\tilde{Q}(\theta, R)$ is not single valued has zero measure, for any $\varepsilon > 0$ there exists a compact set S such that $\{\theta \in \Theta : |\tilde{Q}(\theta, R)| \neq 1\} \subset \Theta \setminus S$ and $\lambda(\Theta \setminus S) < \varepsilon$ (where λ is Lebesgue measure). For any such S , there exists a neighborhood $B_S \subset \mathcal{C}$ of R such that $R' \in B_S$ implies

$$\left| \int_S u(\theta, Q(\theta)) dH(\theta) - \mathcal{U}(R) \right| < \delta/2$$

Taking ε small enough gives implies $\mathcal{U}(R') \subseteq (\mathcal{U}(R) - \delta, \mathcal{U}(R) + \delta)$ as desired. \square

A.6.13.3. Proof of Lemma 9.

PROOF. Suppose M is discontinuous at p' , and let $\theta' \in \theta_M(p'|R)$. First, suppose that $p \mapsto R(M(p), \theta')$ is continuous at p' . Since M is discontinuous, there exists an open neighborhood U of $M(p')$ such that for any $\varepsilon > 0$ there exists $p'' \in N_\varepsilon(p')$ with $M(p) \notin U$. Since $p \mapsto R(M(p), \theta')$ is continuous at p' , for any $\delta > 0$ we can choose ε small to guarantee $|R(M(p''), \theta') - R(M(p'), \theta')| < \delta$. But then let \hat{R} be a continuous function

in a δ -neighborhood of R such that $\hat{R}(M(p''), \theta') = p'$, so $M(p'') \in \tilde{Q}_{\hat{R}}(\theta'|M)$. Therefore we cannot have upper hemicontinuity of $R \mapsto \tilde{Q}_R(\theta'|M)$ at R .

Now, suppose $p \mapsto R(M(p), \theta')$ is discontinuous at p' . Assume M is left-continuous at p' (symmetric argument for right-continuous, and similar for removable discontinuity). Then there exists $\varepsilon > 0$ such that either $R(M(p), \theta') < p$ for all $p \in [p' - \varepsilon, p')$ or $R(M(p), \theta') > p$ for all $p \in [p' - \varepsilon, p')$. Assume without loss of generality that the former holds. Then let \hat{R} be a continuous function such that $\hat{R}(M(p), \theta') > R(M(p), \theta')$ for all $p \in [p' - \varepsilon, p')$. For \hat{R} close to R there will be a neighborhood U of p' such that $\hat{R}(M(p), \theta') \neq p$ for all $p \in U$. This is because M is discontinuous at p' . Then $R \mapsto \tilde{Q}_R(\theta'|M)$ cannot be lower hemicontinuous at R . \square

A.6.13.4. Proof of Lemma 4.

PROOF. If the function $m := Q \circ P^{-1}$ on $P(\Theta)$ is continuous then we are done; it does not matter how M is defined outside of $P(\Theta)$. Suppose m has a discontinuity at p' in the interior of $P(\Theta)$, and let $\theta' = P^{-1}(p')$. First, suppose P is continuous at θ' . For a given $\varepsilon > 0$, the single crossing assumption implies that there is at most one state $\theta^* \in [\theta' - \varepsilon, \theta' + \varepsilon]$ such that $R(Q(\theta' - \varepsilon), \theta^*) = R(Q(\theta' + \varepsilon), \theta^*)$. Let $\hat{Q} = Q$ on $\Theta \setminus [\theta' - \varepsilon, \theta' + \varepsilon]$, and for any $\tilde{\theta} \in (\theta' - \varepsilon, \theta' + \varepsilon) \setminus \theta^*$, let $\hat{Q} = Q(\theta' - \varepsilon)$ for $\theta \in [\theta' - \varepsilon, \tilde{\theta}]$, and $\hat{Q} = Q(\tilde{\theta} + \varepsilon)$ for $\theta \in (\tilde{\theta}, \theta' + \varepsilon]$. Let \hat{P} be the price function associated with \hat{Q} . Then $\hat{m} := \hat{Q} \circ \hat{P}^{-1}$ is continuous on $\hat{P}(\Theta) \cap (p' - \delta, p' + \delta)$ for some $\delta > 0$.

A similar construction applies to the case where P is discontinuous at θ' . In this case assume WLOG P is left-continuous and let θ'' be the right limit of $\theta_M(p)$ at p' . Then we modify \hat{Q} as above, except that \hat{Q} is unchanged on (θ', θ'') . \square

A.6.13.5. Proof of Proposition 12.

PROOF. Sufficiency is obvious. Conditions *iii* and *iv* are necessary, as discussed in the paragraph preceding Proposition 12. To show necessity of *i* and *ii*, restrict attention to a one-dimensional strictly ordered chain in Θ (e.g. the diagonal). For the restriction of Q to this chain, necessity of *i* and continuity for interior states then follow from the same arguments as in the uni-dimensional case. Under *iii* and *iv*, this implies that *i* holds; if there is a non-monotonicity on some chain then there will be a non-monotonicity on every chain. Similarly Q must be continuous on the interior. \square

A.6.13.6. Proof of Lemma 5.

PROOF. First note that $s_i \mapsto x_i(p|a, \mathcal{I}, s_i)$ is Lipschitz continuous since Ω is bounded and $s_i = \omega + \varepsilon_i$ for a normally distributed ε_i . Increasing ω by δ has the same effect on aggregate demand as increasing s_i by δ for all i . Then $\omega \mapsto X(p|a, \mathcal{I}, \omega)$ is Lipschitz continuous since σ_i and τ_i are bounded in the population. \square

A.6.13.7. Proof of Lemma 6.

PROOF. First, note if (ω'', z'') and (ω', z') are elements of \mathcal{I} , with $\omega'' > \omega'$ then it must be that $z'' > z'$. This follows from the fact that aggregate demand is strictly increasing in ω and strictly decreasing in p .

The function $\omega \mapsto X(p|a, \mathcal{I}, \omega)$ is Lipschitz continuous by Lemma 5. So for any $\kappa > 0$ there exists $\delta > 0$ such that for any $(\omega'', z''), (\omega', z') \in \mathcal{I}$, we have $|\omega'' - \omega'| < \delta$ implies $|z'' - z'| < \kappa$. In other words, there is uniform bound on the “slope” of \mathcal{I} in $\Omega \times \mathcal{Z}$ space. Since the prior distribution on $\Omega \times \mathcal{Z}$ is absolutely continuous, this implies the desired result. \square

A.6.13.8. Proof of Lemma 7.

PROOF. Define the random variable $\tilde{V}^a := \pi(a, \theta) = \beta_0^a + \beta_1^a \theta$. Then define $\tilde{S}_i^a := \beta_1^a s_i + \beta_0^a = \tilde{V}^a + \beta_1^a \varepsilon_i$. Thus conditional on knowing the principal's action, investor i 's observation of s_i is equivalent to observing a signal \tilde{S}_i^a which is equal to the true dividend \tilde{V}^a plus normal random noise, where the variance of the noise term depends on a ; it is given by $\sigma_{ai}^2 = (\beta_1^a)^2 \sigma_i^2$. The results then follows from the proof of Proposition 2.2 in Breon-Drish (2015) (Online Appendix). The proposition in Breon-Drish (2015) pertains to the information sets revealed by equilibrium price functions which are continuous and satisfy a differentiability assumption. However for the relevant direction of the proof, these conditions are only needed to guarantee that the distribution of \tilde{V}^a conditional on \mathcal{I} has a density, which is implied here by Lemma 6. \square

A.6.13.9. Proof of Proposition 13.

PROOF. Given Lemma 7, we just need to identify what the coefficients on the linear statistic are.

Fix M , and let $L_M : \Omega \times \mathcal{Z} \times \mathcal{A} \rightarrow \mathbb{R}$ be the equilibrium statistic in a generalized linear equilibrium in which the price reveals exactly a hyperplane. Define the random variable $\tilde{V}^a := \pi(a, \omega) = \beta_0^a + \beta_1^a \omega$. Then define $\tilde{S}_i^a := \beta_1^a s_i + \beta_0^a = \tilde{V}^a + \beta_1^a \varepsilon_i$. Thus conditional on knowing the principal's action, investor i 's observation of s_i is equivalent to observing a signal \tilde{S}_i^a which is equal to the true dividend \tilde{V}^a plus normal random noise, where the variance of the noise term depends on a ; it is given by $\sigma_{ai}^2 = (\beta_1^a)^2 \sigma_i^2$. Let \tilde{L}^a be the random variable $L_M(\omega, z, a)$.

We first fix the principal's action at a , and generalize Breon-Drish (2015) Proposition 2.1 to allow for supply shocks with a truncated normal distribution. We will therefore suppress dependence of $\tilde{S}_i^a, \tilde{V}^a, \tilde{L}^a$ on the action a for the time being. Abusing notation, write the statistic L in terms of v , rather than ω ; that is, $L(v, z|a) = \alpha v - z$, suppressing the dependence on M .³ For fixed a , the truncation is the only difference between the current setting and that of Breon-Drish (2015) Proposition 2.1. By the same steps as the proof for Proposition 2.1 in Breon-Drish (2015) Online Appendix, we can show that the conditional distribution of \tilde{V}^a conditional on $\tilde{S}_i^a = s_i$ and $\tilde{L}^a = \ell$ is given by

$$(A.3) \quad dF_{\tilde{V}|\tilde{S},\tilde{L}}(v|s_i, \ell) = \frac{\mathbb{1}[\ell - \alpha v \in (-b, b)] \exp \left\{ \left(\frac{1}{\sigma_{ai}^2} s_i + \frac{\alpha}{\sigma_Z^2} \ell \right) v - \frac{1}{2} \left(\frac{1}{\sigma_{ai}^2} + \frac{\alpha^2}{\sigma_Z^2} \right) v^2 \right\} dF_{\tilde{V}}(v)}{\int_{\frac{\ell-b}{\alpha}}^{\frac{\ell+b}{\alpha}} \exp \left\{ \left(\frac{1}{\sigma_{ai}^2} s_i + \frac{\alpha}{\sigma_Z^2} \ell \right) x - \frac{1}{2} \left(\frac{1}{\sigma_{ai}^2} + \frac{\alpha^2}{\sigma_Z^2} \right) x^2 \right\} dF_{\tilde{V}}(x)},$$

where $\mathbb{1}[\cdot]$ is the indicator function. This is not in the *exponential family* of distributions, as defined in Breon-Drish (2015) Assumption 10. Nonetheless, it will have similar properties.

We can write the conditional distribution in (A.3) as

$$\mathbb{1}[\ell - \alpha v \in (-b, b)] \exp \left\{ \hat{L}(s_i, \ell) v - g \left(\hat{L}(s_i, \ell); \alpha, \ell \right) \right\} dH(v; \alpha),$$

³This abuse of notation is done to match the notation of Breon-Drish (2015). Note that in that paper “ a ” is used in place of α to denote the slope of the equilibrium statistic. The reader examining Breon-Drish (2015) should not confuse this with the notation for the principal action used in the current paper.

where

$$\begin{aligned}\hat{L}(s, \ell) &= \left(\frac{1}{\sigma_{ai}^2} s_i + \frac{\alpha}{\sigma_Z^2} \ell \right) \\ g_i(\hat{L}; \alpha, \ell) &= \log \left(\int_{\frac{\ell-b}{\alpha}}^{\frac{\ell+b}{\alpha}} \exp \left\{ \left(\frac{1}{\sigma_{ai}^2} s_i + \frac{\alpha}{\sigma_Z^2} \ell \right) x - \frac{1}{2} \left(\frac{1}{\sigma_{ai}^2} + \frac{\alpha^2}{\sigma_Z^2} \right) x^2 \right\} dF_{\tilde{V}}(x) \right) \\ dH_i(v; \alpha) &= \exp \left\{ -\frac{1}{2} \left(\frac{1}{\sigma_{ai}^2} + \frac{\alpha^2}{\sigma_Z^2} \right) v^2 \right\} dF_{\tilde{V}}(v)\end{aligned}$$

This has the following important implication (essentially the same as Lemma A6 in Breon-Drish (2015)). Since the conditional distribution must integrate to 1, i.e.

$$\int_{\frac{\ell-b}{\alpha}}^{\frac{\ell+b}{\alpha}} \exp \left\{ \hat{L}(s_i, \ell) v - g \left(\hat{L}(s_i, \ell); \alpha, \ell \right) \right\} dH(v; \alpha) = 1$$

we have that

$$\int_{\frac{\ell-b}{\alpha}}^{\frac{\ell+b}{\alpha}} \exp \left\{ \hat{L}(s_i, \ell) v \right\} dH(v; \alpha) = \exp \left\{ g \left(\hat{L}(s_i, \ell); \alpha, \ell \right) \right\}.$$

As a result, for any $t \in \mathbb{R}$ we have

$$\mathbb{E} \left[\exp \{ t \tilde{V} \} | s, \ell \right] = \exp \left\{ g \left(t + \hat{L}(s_i, \ell); \alpha, \ell \right) - g \left(\hat{L}(s_i, \ell); \alpha, \ell \right) \right\}.$$

The remainder of the proof for the fixed-action case proceeds as in Breon-Drish (2015) Proposition 2.1. In particular, this shows that in any generalized linear equilibrium with fixed action a ,

$$\alpha = \int_i \frac{\tau_i}{\sigma_{ai}^2} di.$$

Since $v = \beta_0^a + \beta_1^a \omega$ and $\sigma_{ai}^2 = (\beta_1^a)^2 \sigma_i^2$ we have

$$L^*(\omega, z|a) = \beta_0^a \int_i \frac{\tau_i}{\sigma_{ai}^2} di + \left(\frac{1}{\beta_1^a} \int_i \frac{\tau_i}{\sigma_{ai}^2} di \right) \cdot \omega - z$$

Since only the level sets of L^* matter, we can ignore the first term.

We now show that the result holds under feedback as well. Given M , the investor knows which action the principal will take conditional on the price. In a generalized linear equilibrium, the investor's demand is therefore determined by maximizing utility given that the price is p , the action is $M(p)$, the observed signal is \tilde{S}_i^a , and the extended state is in $\{(\omega, z) : L_M(\omega, a|a) = \ell\}$ for the value of ℓ corresponding to price level p . The remaining question is which $L_M(\cdot|a)$ could constitute equilibrium statistics given action a and decision rule M . The first part of the proof shows that if the principal's action is fixed at a then there is a unique equilibrium statistic $L^*(\omega, z|a)$. Since all investors know the principal's action once they observe the price, this L^* must be the equilibrium statistic, regardless of M . \square

A.6.13.10. Proof of Proposition 14.

PROOF. Proposition 13 says that the equilibrium price must reveal at least the linear statistic. We want to show that the price can reveal no more than this. For $p \in \tilde{P}(\Omega, \mathcal{Z})$ let $l^*(p)$ be the linear statistic revealed by p . Suppose that $\mathcal{I}(p) := \{(\omega, z) : \tilde{P}(\omega, z) = p\} \neq l^*(p)$, so that the price reveals more than the linear statistic. We show that in this case there will be multiplicity. This follows from the fact that the set of states $\{(\omega, z) : X(p|M(p), \mathcal{I}(p), \omega) = z\}$ is the entire linear statistic $l^*(p)$. This follows from the

proof of Lemma 6 and Proposition 2.2 in Breon-Drish (2015) (Online appendix), which shows that individual demands will be linear in signals for any price. \square

A.7. Bridgeability

This section discusses bridgeability further. We provide sufficient conditions for the various notions of bridgeability, and show that they are satisfied in common settings.

Let (\mathcal{A}, \succ) be a partially ordered set. Say (\mathcal{A}, \succ) is *upward directed* if for any two $a'', a' \in \mathcal{A}$ there exists $c \in \mathcal{A}$ such that $c \succ a''$ and $c \succ a'$. Downward directed is defined analogously.⁴ We use the notation $a''_\alpha a' \equiv \alpha a'' + (1 - \alpha)a'$. Say that \succ is preserved by mixtures if for any $a'' \succ a'$ and $\alpha \in (0, 1)$, $a'' \succ a''_\alpha a' \succ a'$. Finally, say that $a \mapsto R(a, \theta)$ is *strongly monotone with respect to \succ* if $a'' \succ a'$ and $a'' \neq a'$ implies $R(a'', \theta) > R(a', \theta)$. We use the notation $a''_\alpha a' \equiv \alpha a'' + (1 - \alpha)a'$. The following proposition gives sufficient conditions for full bridgeability, but it is also useful because the proof of the existence of a monotone path is constructive. This construction could potentially be useful in applications.

Proposition 30. *Let (\mathcal{A}, \succ) be a partially ordered set that is both upward and downward directed, and such that \succ is preserved by mixtures. If $R(\cdot, \theta)$ is strongly monotone with respect to \succ then there is a monotone path between a' and a'' at θ iff $R(a'', \theta) \neq R(a', \theta)$*

PROOF. The condition $R(a', \theta) \neq R(a'', \theta)$ is obviously necessary. It remains to show that it is sufficient. That is, we want to show that there exists a monotone path between any $a'', a' \in \mathcal{A}$ such that $R(a', \theta) \neq R(a'', \theta)$. Assume without loss that $R(a'', \theta) > R(a', \theta)$.

⁴A lattice is an upward and downward directed set, but the converse is not true.

If $a'' \succ a'$ then the ray from a'' to a' is a monotone path. This follows since \succ is preserved by mixtures and $R(\cdot, \theta)$ is strongly monotone.

Suppose a' and a'' are not ordered. Let \bar{a} be an upper bound for a'', a' , i.e. $\bar{a} \succ a''$ and $\bar{a} \succ a'$, and let \underline{a} be a lower bound. Both exist since (\mathcal{A}, \succ) is upward and downward directed. By continuity of R , there exists $\bar{\lambda} \in (0, 1)$ such that $R(\bar{a}_{\bar{\lambda}} a', \theta) = R(a'', \theta)$. Similarly there exists $\underline{\lambda} \in (0, 1)$ such that $R(a''_{\underline{\lambda}} \underline{a}, \theta) = R(a', \theta)$.

We will now construct one half of the monotone path from a' to a'' . Let $t : [0, 1] \rightarrow [\bar{\lambda}, 1] \times [0, 1]$ be a continuous and strictly monotone function, and let $t_i(x)$ be the i^{th} coordinate of $t(x)$. For each $x \in (0, 1)$, we have $R(\bar{a}_{t_1(x)} a', \theta) > R(a'', \theta)$, $R(\underline{a}_{t_2(x)} a', \theta) < R(a', \theta)$, and $\bar{a}_{t_1(x)} a' \succ \bar{a}_{t_1(x)} a'$. These properties follow from strong monotonicity of R and the fact that \succ is preserved under mixtures.

For each $x \in (0, 1)$, define $f(x)$ by $R((\bar{a}_{t_1(x)} a')_{f(x)} (\underline{a}_{t_2(x)} a'), \theta) = xR(a'', \theta) + (1 - x)R(a', \theta)$. We claim that $x \mapsto (\bar{a}_{t_1(x)} a')_{f(x)} (\underline{a}_{t_2(x)} a')$ is a continuous function. It is a well defined function by strong monotonicity of R . It is continuous since R and t are continuous. Moreover, by construction $x \mapsto R((\bar{a}_{t_1(x)} a')_{f(x)} (\underline{a}_{t_2(x)} a'), \theta)$ is strictly increasing, and $(\bar{a}_{t_1(0)} a')_{f(0)} (\underline{a}_{t_2(0)} a') = a'$. Therefore $x \mapsto (\bar{a}_{t_1(x)} a')_{f(x)} (\underline{a}_{t_2(x)} a')$ forms one half of a monotone path from a' to a'' . The other half of the monotone path is defined analogously, using a'' and $\underline{\lambda}$ in place of a' and $\bar{\lambda}$. \square

Proposition 30 makes it easy to identify when a discontinuity will be bridgeable. For example, it implies that when \mathcal{A} is a chain a gap between a' and a'' will be bridgeable at θ iff $R(\cdot, \theta)$ is strictly monotone on (a', a'') .

More importantly, Proposition 30 implies that every discontinuity will be bridgeable when $\mathcal{A} = \Delta(Z)$, i.e. the set of distributions on some set Z , under mild assumptions on R .

Let $\pi(z, \theta)$ be a real valued function, with $\theta \mapsto \pi(z, \theta)$ continuous for all z . For example, $\pi(a, \theta)$ could represent a company's cash flow as a function of the state and government intervention $z \in Z$. In state θ , any $a \in \mathcal{A}$ induces a distribution $F(a, \theta)$ on \mathbb{R} via $\pi(\cdot, \theta)$. Let \succ_{FOSD} be the first-order stochastic dominance order. This partial order on $\Delta(\mathbb{R})$ induces a preorder \succeq on \mathcal{A} . Define $a'' \succ a'$ by $a'' \succeq a'$ and $\neg(a' \succeq a'')$ if $a'' \neq a'$, and $a' \succ a'$ for all a' . If $\pi(z', \theta) \neq \pi(z'', \theta)$ for all $z'' \neq z'$ then $\succeq = \succ$. Then $a \mapsto R(a, \theta)$ is strongly monotone if $F(a'', \theta) \succ F(a', \theta)$ implies $R(a'', \theta) \succ R(a', \theta)$. The partially ordered set (\mathcal{A}, \succ) satisfies the conditions of Proposition 30 (when $\pi(z', \theta) \neq \pi(z'', \theta)$ for all $z' \neq z''$ it is in fact a lattice).

Corollary 8. *If $\mathcal{A} = \Delta(Z)$ and for all θ $a \mapsto R(a, \theta)$ is strongly monotone with respect to the order induced by first-order stochastic dominance, then the environment is fully bridgeable.*

It will also be useful to establish a related notion of bridgeability. Say that there exists a monotone path from (a', θ') to (a'', θ'') if there exists a continuous function $\gamma : [0, 1] \rightarrow \mathcal{A} \times [\theta', \theta'']$ such that $\gamma(0) = (a', \theta')$, $\gamma(1) = (a'', \theta'')$, $x \mapsto \gamma_1(x)$ is weakly increasing and $R(\gamma_1(x), \gamma_2(x))$ is strictly increasing. The path is *strongly monotone* if moreover $\gamma_1(x)$ is strictly increasing.

Recall that the environment is *continuously bridgeable* if for any $\theta^* \in \Theta$ there exists $\varepsilon > 0$ such that if a', a'' is bridgeable at θ^* and $R(a'', \theta) \neq R(a', \theta)$ for all $\theta \in [\theta^*, \theta^* + \varepsilon]$ then there exists a sup-norm continuous function $\sigma(\cdot | a', a'') : [\theta^*, \theta^* + \varepsilon] \rightarrow \mathcal{A}^{[0,1]}$ such that $\sigma(\theta | a', a'')$ is a monotone path from a' to a'' for all $\theta \in [\theta^*, \theta^* + \varepsilon]$. Say that the environment is *continuously fully bridgeable* if it is full bridgeable and continuously bridgeable.

Lemma 30. *Assume $\theta \mapsto R(a, \theta)$ is strictly monotone for all $a \in \mathcal{A}$ and the environment is continuously fully bridgeable. Then the environment is correctable if for all θ such that $R(a', \theta) = R(a'', \theta)$ for all $a \in \mathcal{A}$ the following holds: there exists $a_1 \in \mathcal{A}$ and $\delta > 0$ such that*

- i. $R(a_1, \theta') > R(a, \theta')$ for all $a \neq a_1$ and $\theta' \in (\theta - \delta, \theta)$.*
- ii. $R(a, \theta') > R(a_1, \theta')$ for all $a \neq a_1$ and $\theta' \in (\theta, \theta + \delta)$.*

PROOF. If P is decreasing then the existence of an approximating Q' around any degenerate discontinuity follows immediately from continuous bridgeability. Assume therefore that P is increasing.

First suppose that there is a degenerate discontinuity at some θ such that there exist a', a'' with $R(a', \theta) \neq R(a'', \theta)$. Assume there exists \bar{a} such that $R(\bar{a}, \theta) > R(Q(a), \theta)$ (the argument for the reverse inequality is symmetric). Since P is increasing and continuous at θ , for any $\varepsilon > 0$ there exists $\theta'' \in (\theta, \theta + \varepsilon)$ and $\theta' \in (\theta - \varepsilon, \theta)$ such that $R(\bar{a}, \theta) > R(Q(\theta''), \theta)$ for all $\theta \in (\theta', \theta'')$. Then since the environment is continuously bridgeable (in particular between $Q(\theta'')$ and \bar{a}) there exists a continuous Q' on $[\theta', \theta'']$, with the corresponding P' strictly increasing, such that $R(Q'(\theta'), \theta') \in (R(Q(\theta'), \theta'), R(Q(\theta''), \theta''))$ and $Q'(\theta'') = Q(\theta'')$. Thus the environment is correctable.

Now suppose there is a degenerate discontinuity at some θ such that $R(a'', \theta) = R(a', \theta)$ for all $a', a'' \in \mathcal{A}$. Then under conditions *i* and *ii* the following Q' satisfies the conditions for correcting the degenerate discontinuity: for any $\varepsilon < \delta$, $Q' = a_1$ on $(\theta - \varepsilon, \theta + \varepsilon)$ and equals Q elsewhere. \square

Note that Lemma 30 implies that the environment is correctable if for all θ there exist a', a'' such that $R(a', \theta) \neq R(a'', \theta)$. Except for unusual cases, the environment will be continuously fully bridgeable when it is fully bridgeable. For example, the environment of Corollary 8 is continuously fully bridgeable when Z is finite and $\theta \mapsto \pi(z, \theta)$ is differentiable for all z .

Lemma 31. *Assume $\mathcal{A} = \Delta(Z)$ for some finite Z , $\pi(z, \theta)$ is differentiable for all z , and for all θ , $a \mapsto R(a, \theta)$ is strongly monotone with respect to the first-order stochastic dominance induced order. Then the environment is continuously fully bridgeable.*

PROOF. First suppose $\min_{z'', z' \in Z} |\pi(z'', \theta^*) - \pi(z', \theta^*)| > 0$. Then by continuity of $\theta \mapsto \pi(z, \theta)$, there exists $\varepsilon > 0$ such that $\pi(z'', \theta) > \pi(z', \theta) \Leftrightarrow \pi(z'', \theta^*) > \pi(z', \theta^*)$ for all $\theta \in [\theta^*, \theta^* + \varepsilon]$ and z', z'' . Thus the partial order on \mathcal{A} induced by first-order stochastic dominance is the same for all $\theta \in (\theta^* - \varepsilon, \theta^* + \varepsilon)$. This implies that the join and meet are the same for any a', a'' , and so the construction used in the proof of Proposition 30 can make use of the same join and meet. Then the conditions of continuous bridgeability are implied by continuity of R .

Now suppose $\pi(z'', \theta^*) = \pi(z', \theta^*)$ for all $z'', z' \in B \subset Z$. Suppose that for any $\delta > 0$ there exists $\theta \in [\theta^*, \theta^* + \delta]$ and $z'', z' \in B$ such that $\pi(z'', \theta) > \pi(z', \theta)$. Then by differentiability of π in θ , there exists a set $C \subset B$ and $\varepsilon > 0$ such that such that *i)* $\pi(z'', \theta) = \pi(z', \theta)$ for all $\theta \in [\theta^*, \theta^* + \delta]$ and all $z', z'' \in C$, and *ii)* $\pi(z'', \theta) > \pi(z', \theta) \Leftrightarrow \pi(z'', \theta') > \pi(z', \theta')$ for all $\theta, \theta' \in (\theta^*, \theta^* + \varepsilon]$ and all $z', z'' \in Z \setminus C$. Then the FOSD-induced order on \mathcal{A} is the same for any $\theta', \theta'' \in [\theta^*, \theta^* + \delta]$. Moreover, this order is a superset of the FOSD-induced order at θ^* : if a'' first-order stochastically dominates a' at $\theta' \in (\theta^*, \theta^* + \delta]$

then it will also do so at θ^* . Thus for any a', a'' we can use the join and meet for the FOSD order induced by $\theta \in (\theta^*, \theta^* + \delta]$ to construct the monotone path θ^* as well. Then the conditions of continuous bridgeability are implied by continuity of R . \square

The following are weaker notions of bridgeability, which it will only be necessary to define on the extreme states $\underline{\theta}, \bar{\theta}$.

Definition 30. A discontinuity in Q at $\bar{\theta}$ is **upper-bridgeable** if there exists a path γ from $\lim_{\theta \nearrow \bar{\theta}} Q(\theta)$ to $Q(\bar{\theta})$ such that $R(\gamma(x), \bar{\theta}) \leq \max\{\lim_{\theta \nearrow \bar{\theta}} R(Q(\theta), \theta), R(Q(\bar{\theta}), \bar{\theta})\}$ for all $x \in [0, 1]$, with equality iff

$$\gamma(x) = \arg \max_{a \in \{\lim_{\theta \nearrow \bar{\theta}} Q(\theta), Q(\bar{\theta})\}} R(a, \bar{\theta}).$$

Definition 31. A discontinuity in Q at $\underline{\theta}$ is **lower-bridgeable** if there exists a path γ from $\lim_{\theta \searrow \underline{\theta}} Q(\theta)$ to $Q(\underline{\theta})$ such that $R(\gamma(x), \underline{\theta}) \geq \min\{\lim_{\theta \searrow \underline{\theta}} R(Q(\theta), \theta), R(Q(\underline{\theta}), \underline{\theta})\}$ for all $x \in [0, 1]$, with equality iff

$$\gamma(x) = \arg \min_{a \in \{\lim_{\theta \searrow \underline{\theta}} Q(\theta), Q(\underline{\theta})\}} R(a, \underline{\theta}).$$

A.8. Deriving R

We first present general conditions under which the market admits a reduced-form representation, and then show when these conditions are satisfied in specific settings.

The reduced form is easily derived in models where agents in the market do not learn from the price. If this is the case, the the reduced form will obtain provided the state is properly defined. The state used in the reduced form representation must identify all uncertainty in the market. This may or may not be the same as the payoff-relevant state

from the perspective of the principal. For example, in the noisy REE model studied in Section 2.4.2, the state includes both the payoff relevant state and the supply shock.

However, the derivation of a reduced-form representation is more challenging when agents in the market learn from the price, as in rational expectations models. This is because the decision rule will shape the information revealed by the price. As a result there may be different equilibrium prices consistent with a given equilibrium action in the same state, when there is different information revealed by the price. We therefore focus here on identifying conditions under which the market admits a reduced form representation in rational expectations models.

Consider the following general rational expectations model, which nests many of the examples studied in this paper. For a fixed principal action $a \in \mathcal{A}$, market clearing is defined by the condition $\chi(\theta, p, \Lambda|a) = 0$, where p is the price and $\Lambda \subseteq \Theta$ is the public information, which will be revealed by the price. For example, χ could represent aggregate demand in a model in which investors submit limit orders as a function of the private information and the public information revealed by the price. The assumption that market clearing is measurable with respect to the state θ requires that the state be properly defined, as discussed above.

Still fixing the principal's action at a , a necessary condition for the equilibrium price of p in state θ to reveal the event Λ is that

$$(A.4) \quad \theta \in \Lambda \subseteq \{\theta' \in \Theta : \chi(\theta', p, \Lambda|a) = 0\}.$$

Refinements imposed as part of the solution concept may limit the set of public information sets that can be revealed by the price in equilibrium. In other words, in may

not be possible for every Λ satisfying Equation (A.4) to be revealed in equilibrium. The following condition is sufficient to guarantee that the market admits a reduced-form representation, taking into account the fact that the solution concept may impose restrictions on the information revealed by the price.

B1. $\forall a \in \mathcal{A}$, and $\theta \in \Theta$, \exists a unique p such that

$$\theta \in \Lambda \subseteq \{\theta' \in \Omega \times \mathcal{Z} : \chi(\omega', p, \Lambda|a) - z' = 0\}$$

for some $\Lambda \subseteq \Theta$.

In words, this assumption says that there is a unique price that could arise in equilibrium at state θ when the action a is taken for all states in the public information set to which θ belongs. This differs from the stronger assumption that there is also a unique information set to which θ can belong in this case, stated below.

B1'. $\forall a \in \mathcal{A}$, and $\theta \in \Theta$, \exists a unique p and unique $\Lambda \subseteq \Theta$ such that

$$\theta \in \Lambda \subseteq \{\theta' \in \Omega \times \mathcal{Z} : \chi(\omega', p, \Lambda|a) - z' = 0\}.$$

While B1 is sufficient for the market to admit a reduced form, we will see problems in which B1' can be established.

Proposition 31. *Under B1, the market admits a reduced-form representation.*

PROOF. Given M , let F_M be equilibrium price function. Fix $\theta = (\omega, z)$. Let $\Lambda = \{\theta' \in \Theta : F_M(\theta') = F_M(\theta)\}$. Let a be the equilibrium action at state θ . Since the action is measurable with respect to the price, the same action is take at all $\theta' \in \Lambda$. A necessary

condition for equilibrium:

$$\Lambda \subseteq \{\theta' \in \Omega \times \mathcal{Z} : \chi(\omega, F_M(\theta), \Lambda|a) = z\}$$

Under B1, this condition uniquely determines $p = F_M(\theta)$. All we assumed about M is that eq. action in state θ is a . Thus the market admits a reduced form. \square

The key insight in the proof of Proposition 31 is that while the information revealed by the price in a given state depends on global features of the equilibrium price function, the action is measurable with respect to the price. Thus the action will be the same in all states in a given public information set. We can therefore use local properties of equilibrium to identify the equilibrium price.

A.8.1. Asset market

We show here that summarizing the market through the function R is consistent with a model of information aggregation. Suppose there is a unit mass of traders. Traders receive conditionally independent signals σ_i about the state, with conditional distribution $h(\cdot|\theta)$. Assume that $h(\cdot|\theta) \neq h(\cdot|\theta')$ for all $\theta \neq \theta'$. Traders are expected utility maximizers. The payoff to trader i who purchases a quantity x of the asset when the principal takes action a , the state is θ , and the asset price is p is given by $V_i(a, \theta, x, p)$, which is assumed to be strictly decreasing in p .⁵ For a fixed action a the demand of trader i who observes signal σ and knows that the state is in $\mathcal{I} \subseteq \Theta$ is given by

$$x_i(p|a, \sigma_i, \mathcal{I}) = \max_x E[V_i(a, \theta, x, p)|\sigma, \mathcal{I}].$$

⁵For example, each trader has a strictly increasing Bernoulli utility function u_i and wealth w_i , and $V_i(a, \theta, x, p) \equiv u_i(x(\pi(a, \theta) - p) + w_i)$.

Assume $p \mapsto x_i$ is strictly decreasing for all i (which is implied by assuming, for example, that $(x, p) \mapsto v_i(a, \theta, x, p)$ satisfies strict single crossing). Trader heterogeneity, both of utilities and beliefs, is allowed for, but for simplicity assume that there are finitely many trader types, meaning finitely many distinct demand functions in the population. Normalizing the aggregate supply of the asset to zero, the market clearing condition is

$$\int_0^1 x_i(p|a, \sigma_i, \mathcal{I}) di = 0.$$

Since there is a continuum of traders and a finite number trader types aggregate demand is deterministic, conditional on the state and the principal action a . Thus we can write market clearing in state θ as

$$X(p|a, \mathcal{I}, \theta) = 0.$$

Let $P^*(a, \mathcal{I}, \theta)$ be the unique price that clears the market.

Given any price function $\tilde{P} : \Theta \rightarrow \mathbb{R}$, let $\mathcal{I}_{\tilde{P}} : \Theta \rightarrow 2^\Theta$ be the coarsest partition with respect to which \tilde{P} is measurable. We say that \tilde{P} *induces* partition $\mathcal{I}_{\tilde{P}}$.

A *rational expectations equilibrium* (REE) given decision rule M consists of a price function \tilde{P} such that $X(\tilde{P}(\theta)|M(\tilde{P}(\theta)), \mathcal{I}_{\tilde{P}}(\theta), \theta) = 0$ for all θ . Let \mathcal{M} be the set of decision rules for which there exists a REE. For any decision rule $M \in \mathcal{M}$, let \tilde{P}_M be the associated REE price function.

As defined, a rational expectations equilibrium only determines the beliefs of traders on path, that is for prices in $\tilde{P}_M(\Theta)$. However as we are concerned about equilibrium multiplicity, we also need to specify the inferences made off-path. We therefore augment the definition of rational expectations equilibrium by assuming that there is some map

$\lambda : \mathcal{P} \setminus \tilde{P}_M(\Theta) \rightarrow 2^\Theta$ specifying the inferences drawn from the price for off-path beliefs.

We require only that λ be consistent with market clearing, that is

$$\lambda(p) \subseteq \{\theta \in \Theta : X(p|M(p), \lambda(p), \theta) = 0\}$$

The population distribution of signals is different for any distinct $\theta, \theta' \in \mathcal{I}$. It is therefore natural to assume that, unless all states in \mathcal{I} are payoff equivalent, there will exist some pair of states $\theta, \theta' \in \mathcal{I}$ such that $P^*(a, \mathcal{I}, \theta) \neq P^*(a, \mathcal{I}, \theta')$. The following assumption is sufficient for B1 in this setting.

A1. For any $a \in \mathcal{A}$ and $\mathcal{I} \subseteq \Theta$, if $P^*(a, \mathcal{I}, \theta) = P^*(a, \mathcal{I}, \theta')$ for all $\theta, \theta' \in \mathcal{I}$ then $P^*(a, \mathcal{I}, \theta) = P^*(a, \theta, \theta)$ for all $\theta \in \mathcal{I}$.

This assumption is discussed further following the statement of the proposition.

We want to show the equivalence between implementable mechanisms and rational expectations equilibria.

Proposition 32. *Under A1, there exists a function $R : \mathcal{A} \times \Theta \rightarrow \mathbb{R}$ such that for any decision rule M there exists a rational expectations equilibrium with price function \tilde{P} if and only if M implements \tilde{P} given market clearing function R .*

PROOF. First, we want to show that there exists an R such that for any decision rule M , if there exists a REE given M , with price function \tilde{P} , then M implements \tilde{P} given market clearing function R . Suppose that for decision rules M_1, M_2 there exist REE, with price functions \tilde{P}_1 and \tilde{P}_2 respectively. Let $\mathcal{I}_{\tilde{P}_1}$ and $\mathcal{I}_{\tilde{P}_2}$ be the partitions of Θ induced by \tilde{P}_1 and \tilde{P}_2 respectively.

Define $R(a, \theta) = \{\tilde{P}_M(\theta) : M \in \mathcal{M}, M(\tilde{P}_M(\theta)) = a\}$. That is $R(a, \theta)$ is the set of prices that can be supported as part of a REE for which the equilibrium action in state θ is a .

We want to show that R as defined above is a function. In other words, we want to show that if for some state θ , the equilibrium mixed is a under both M_1 and M_2 (that is, $M_1(\tilde{P}_1(\theta)) = M_2(\tilde{P}_2(\theta)) = a$), then $\tilde{P}_1(\theta) = \tilde{P}_2(\theta)$. Since \tilde{P}_j induces $\mathcal{I}_{\tilde{P}_j}$, it must be that $P^*(a, \mathcal{I}_{\tilde{P}_j}, \theta') = \tilde{P}_j(\theta)$ for all $\theta' \in \mathcal{I}_{\tilde{P}_j}$ for $j \in \{1, 2\}$. Then A1 implies that $P^*(a, \mathcal{I}_{\tilde{P}_j}, \theta) = P^*(a, \theta, \theta)$ for $j \in \{1, 2\}$, so $\tilde{P}_1(\theta) = \tilde{P}_2(\theta)$ as desired. The same argument applies to the off-path information sets specified by λ .

The other direction is straightforward. By the definition of implementation, if M implements \tilde{P} given market clearing function R then \tilde{P} is a REE price function given decision rule M . \square

A1 is an assumption on the payoff structure and the information structure. It is satisfied in typical models of the asset market. For example, A1 will hold if the function $\theta \mapsto v_i(a, \theta, x, p)$ is strictly monotone for all $a, p, x > 0$, and all i ; and the distribution of posteriors induced by $h(\cdot|\theta)$ is monotone (in an appropriate sense) in θ .⁶

The following are sufficient conditions for A1, along with a concrete example that satisfies these conditions. For the example, let $\mathcal{A} = [0, 1]$, $v_i(a, \theta, x, p) = u(x \cdot (\pi(a, \theta) - p) + w_i)$ and assume that $\pi(a, \theta)$ is weakly increasing in θ .

⁶A sufficient condition for the monotonicity of $\theta \mapsto v_i(a, \theta, x, p)$ is co-monotonicity of $\theta \mapsto v_i(a, \theta, x, p)$ for all a (when $v_i(a, \theta, x, p) \equiv u_i(x \cdot (\pi(a, \theta) - p) + w_i)$ this is equivalent to co-monotonicity of $\pi(a, \cdot)$). Posterior monotonicity will hold, for example, if $\sigma = \theta + \delta$ for some continuously distributed zero mean random variable δ .

- (1) *Ordered signals.* Assume that $h(\cdot|\theta'') >_{MLR} h(\cdot|\theta')$ for all $\theta'' > \theta'$. This implies that the posteriors induced by signals are also ordered by MLR; higher signals induce MLR higher posteriors over Θ .

Example: $\sigma = \theta + \varepsilon$, where ε is zero-mean noise.

- (2) *Single-crossing between x, θ .* We want individuals to demand more of the asset when they get a high signal. Assume therefore that $V_i(a, \theta, x, p)$ satisfies single crossing between x and θ . Monotonicity of demand is implied by standard MCS results (see Athey (2001)).

Example: $u(x(\pi(a, \theta) - p) + w_i)$ satisfies single crossing in x and θ when $\theta \mapsto \pi(a, \theta)$ is increasing.

- (3) *Payoff equivalence.* For any \mathcal{I} , we want demand to be strictly increasing in θ unless $V_i(a, \theta, x, p) = V_i(a, \theta', x, p)$ for all $\theta, \theta' \in \mathcal{I}$.

Example: This holds given the assumptions made thus far (in particular, monotonicity of π).

A.8.2. Forecasts and macro aggregates

Many policy decisions are made with reference to macroeconomic outcomes. For example, the government may decide to increase the amount of unemployment benefits or fund worker-retention programs depending on initial jobless claims or the unemployment rate. Many such problems also have a dynamic component. For example, businesses deciding whether or not to fire employees may care about the future unemployment rate both as a signal of demand and as a determinant of government worker-retention policies. In such

settings, forward looking agents often make use of expert forecasts of the relevant macro variables, such as the unemployment rate.

A.8.2.1. Policy decision up-front. The policy maker may prioritize timeliness over accuracy when making certain policy decisions. In such cases it will be necessary for the policy maker to take an action before the relevant aggregate outcome has been realized. The policy maker will therefore make use of expert forecasts. For example, consider the problem of the government choosing the level of unemployment benefits. The policy maker may wish to act before relevant data, such as the unemployment rate in the coming month, has been collected. It must therefore rely on forecasts of the relevant variables. For simplicity, assume that the government conditions its benefits policy exclusively on expert forecasts of the unemployment rate for the coming month (it is straightforward to incorporate other sources of information).⁷

Forecasters wish to provide accurate estimates of the market outcome (We will refer to this simply as the outcome from now on). If there are many forecasters, each individual expects their prediction to have only a small effect on overall expectations.⁸ However they recognize that overall expectations will be used by the policy maker to take an action. These two factors imply that forecasters' private information will shape their expectations of policy decisions, which in turn will affect their forecasts.

This situation is easiest to model if we assume that forecasters observe each others' forecasts, and can make revisions based on what others say. The consistency condition is

⁷Another example in which expert forecasts may be used to inform policy is monetary policy: Bernanke and Woodford (1997) discuss targeting expert forecasts of inflation, rather than realized inflation.

⁸Bloomberg surveys around 80 economists for predictions on the monthly unemployment rate.

that each forecaster doesn't want to change their forecast given those of the others, and the announced policy rule.

Formally, this model is very similar to the market price model. Assume there is a continuum of forecasters \mathcal{F} . Each forecaster $i \in \mathcal{F}$ receives a signal σ_i about the state. Signals are conditionally independent across forecasters. Forecasters make predictions about the value of some variable v , which will not be realized until after the principal has taken an action. Forecasts may have different models of the world, i.e. ways to map their information to a prediction, but assume for simplicity that there are only finitely many models in the population.

The principal bases their decision on some real-valued function of the profile of forecasts, the forecast aggregate. Forecasters iteratively revise their predictions based on their observations of the forecast aggregate. We do not explicitly model the iterative procedure. Rather, we look for a rational expectations equilibrium conditional on the principal's announced decision rule M . In this context, assumption A1 can be restated as follows.

Fix any principal action a , $\mathcal{I} \subseteq \Theta$ and $\theta \in \Theta$, and value of the forecast aggregate f . Assume all forecasters know that the principal will take action a , that the state is in \mathcal{I} , and that the value of the forecast aggregate is f (in addition to their private signals). Let the $X(f|a, \mathcal{I}, \theta)$ be the new value of the forecast aggregate after forecasters have a chance to revise their predictions. This is a deterministic function since there are a continuum of forecasters with i.i.d. signals. Then forecasts reach a fixed point when

$$X(f|a, \mathcal{I}, \theta) = f.$$

Assume that there is unique fixed point for any a, \mathcal{I}, θ (which will be the case, for example, when individual forecasts, as well as the aggregator, are monotone in f), and denote this fixed point by $F^*(a, \mathcal{I}, \theta)$.

A1'. For any $a \in \mathcal{A}$ and $\mathcal{I} \subseteq \Theta$, if $F^*(a, \mathcal{I}, \theta) = F^*(a, \mathcal{I}, \theta')$ for all $\theta, \theta' \in \mathcal{I}$ then $F^*(a, \mathcal{I}, \theta) = F^*(a, \theta, \theta)$ for all $\theta \in \mathcal{I}$.

Assumption A1' is satisfied, for example, when the distribution of beliefs induced in the population is monotone (in an FOSD sense, with an appropriate order on beliefs) in the state, individuals forecasts are monotone in their beliefs, and the forecast aggregate is monotone in individual forecasts (in an FOSD sense).

The existence of R in this setting follows from Proposition 32

A.8.2.2. Policy decision ex-post. Some decision makers may condition their actions on realized outcomes, rather than expectations. Nonetheless, expert forecasts may play a role in shaping behavior. For example, Congress extended the time frame for spending PPP funds after observing that companies had difficulty re-hiring employees. Congress also approved a second tranche of PPP funds after the first was exhausted. Companies condition their payroll decisions or loan applications on expectations of future aggregate outcomes. The base their expectations on expert forecasts. Forecasters make predictions knowing that *i*) expectations will shape business decisions, and *ii*) business decisions will shape the policy response. Again, assume forecasters observe each others' forecasts. Then we need a fixed point that takes into account the feedback of forecasts on policy through business decisions.

Formally this case is similar to that discussed above. There are two periods. Some variable v will realize in the second period, and the principal will take an action in the second period conditional on v . Assume that the principal commits to a rule M mapping v to an action.

A unit mass of economic agents, call them individuals, care about the principal's future action, as well as some underlying state θ . In order to predict what the principal's action will be, individuals rely on the predictions of a set \mathcal{F} of forecasters. Individuals are fairly simplistic: they aggregate forecaster predictions in some way, for example by taking the mean, and assume that this forecast aggregate f will be the true value. They choose their actions based on the action implied by the principal's decision rule, as well as their own private information. Assume that individuals do not infer anything about the state from the forecasters' predictions.⁹ The actions of all individuals, along with the state, jointly determine the outcome v . When all individuals expect the principal to take action $a \in \mathcal{A}$ and the state is θ , the aggregate outcome in the second period will be given by $J(a, \theta)$.

As before, there is a unit mass of forecasters, each of whom receives a private signal about the state. Forecasters observe the current value of the forecast aggregate and revise their decisions. A fixed point is reached when $X(f|a, \mathcal{I}, \theta) = f$. This function incorporates the fact that a affects the aggregate outcome through $J(a, \theta)$. Assuming $A1'$ holds, we have an R function by Proposition 32.

The interesting part of the ex-post decision model is that principal is not intending to use the equilibrium variable, in this case the forecast, to make a decision. The principal may not even be able to commit to a mapping M from the aggregate outcome to an

⁹This type of inference can be added without too much complication, with a suitable version of assumption A1.

action. It could just be that agents anticipate the principal to behave in a certain way ex-post. Nonetheless, the forecast will be determined as a fixed point, and this will impact the aggregate outcome, and thus the principal's decision.

A.8.2.3. Alternative model. Suppose that there is a single forecaster who gets a signal σ . The forecaster is aware of the effect that their prediction will have on individual behavior. The forecaster simply reports their expectation of the outcome v , when this is well defined. This will be well defined iff there is a fixed point to the function $f \mapsto \mathbb{E}[J(M(p), \theta) | \sigma]$. Let $R(a, \sigma) = \mathbb{E}[J(M(p), \theta) | \sigma]$. The analysis of the paper applies, with σ replacing θ .

A.8.2.4. Adding forecast uncertainty. The fact that forecasters receive conditionally independent signals may seem unrealistic. It is straightforward to generalize to a situation in which signals are correlated. Assume that the state consists of a pair (κ, θ) . This is a special case of the multi-dimensional model in Section 2.4. As before, θ is the payoff-relevant state. κ simply determines the joint distribution of signals. Signals are conditionally independent given (κ, θ) .

Let Σ be the space of population signal profiles $\{\sigma_i\}_{i \in \mathcal{F}}$. Assume that there is a complete order on the space of signal profiles, which can be represented by a bijection $b : \Sigma \rightarrow [0, 1]$.¹⁰ Since b is a bijection, $b(\{\sigma_i\}_{i \in \mathcal{F}})$ contains the same information as $\{\sigma_i\}_{i \in \mathcal{F}}$. Then when all forecasters expect the principal to take action a and know that the current forecast aggregate is f , and know that $b(\{\sigma_i\}_{i \in \mathcal{F}}) \in \mathcal{I} \subseteq [0, 1]$, then the updated forecast aggregate will be $X(f | a, \mathcal{I}, b)$. Then the analysis proceeds as before, except that b replaces

¹⁰For example, $\sigma_i = \varepsilon_i + \kappa + \theta$, where ε_i are i.i.d. random variables with a common bounded-support distribution. In this case the order on the set of signal profiles is given by the population mean $\int_{i \in \mathcal{F}} \sigma_i di$.

θ . The principal will have to account for the residual uncertainty when choosing a decision rule.

The discussion in this section applies whether the set \mathcal{F} of forecasters is finite or infinite. However the assumption of a continuum of forecasters remains convenient for two reasons. First, the assumption that such a bijection b exists makes more sense when there is a continuum of forecasters (see the example in the footnotes). Second, when there are finitely many forecasters they will behave strategically. For example, there is in general no reason to expect that forecasts should reach a fixed point when forecasters take into the effect that their forecasts have on the forecast aggregate. For a single forecaster trying to minimize the expected difference between their prediction and the actual outcome, unless they perfectly observe the state, may find it optimal to make a forecast that they know cannot be correct.

A.9. Extension: no commitment

We have assumed throughout that the principal is able to commit to a decision rule. In this section we briefly analyze the situation in which the principal cannot commit.

We assume that all market participants understand the principal's preferences, and can thus predict what the principal will do as a function of the principal's information set. In any REE, the price function will reveal some information to the principal, as a function of which the principal will take their preferred action. Thus any equilibrium price function P will induce a map $m(\cdot; P) : \mathcal{P} \rightarrow \mathcal{A}$, where $m(p; P)$ is the principal's optimal action given the information revealed by $P(\theta) = p$ (or some mixture over these

in the case of indifference). A rational expectations equilibrium without commitment consists of a price function P and decision rule m such that

- i. $P(\theta) = R(m(P(\theta)), \theta)$ for all θ . *(rational expectations)*
- ii. For all p , $m(p)$ is an optimal action for the principal conditional on $\{\theta : P(\theta) = p\}$.
(principal optimality)

The principal optimality condition replaces the commitment condition in the definition of REE used under commitment.

Let $Q^* : \Theta \rightarrow \mathcal{A}$ be the principal's first-best action function. That is, Q^* specifies the principal's optimal action in each state. Assume for simplicity that $\theta \mapsto R(a, \theta)$ is strictly increasing for all a . Then given any m , there is at most a single state θ such that $p = R(m(p), \theta)$. Thus any REE price function must be fully revealing. Given this observation, we have the following equivalent definition of a REE without commitment

Lemma 32. *Assume $\theta \mapsto R(a, \theta)$ is strictly increasing for all a . Then (P, m) constitute an REE without commitment if and only if (P, Q^*) are implementable under commitment (as defined in Definition 8).*

This observation has the following immediate corollary.

Corollary 9. *If Q^* is not implementable then there does not exist an REE without commitment.*

Moreover, we can use the characterization results under commitment to understand equilibrium behavior without commitment. For example, Theorem 1 has the following implication.

Corollary 10. *If $\theta \mapsto R(Q^*(\theta), \theta)$ is non-monotone then either there will be discontinuities at some equilibrium prices or there will be multiple equilibria.*

In other words, the equilibrium will either be vulnerable to manipulation (and not be robust to structural uncertainty), or it will suffer from non-fundamental volatility. Note that there can only be multiple equilibria if there are states for which the principal has multiple optimal actions.