Boston University Theses & Dissertations

2022

Essays in economic theory

BOSTON UNIVERSITY

GRADUATE SCHOOL OF ARTS AND SCIENCES

Dissertation

ESSAYS IN ECONOMIC THEORY

by

ENRICO DE MAGISTRIS

B.Sc., Bocconi University, 2013
M.S., Bocconi University, 2015
M.A., Boston University, 2018

Submitted in partial fulfillment of the

requirements for the degree of

Doctor of Philosophy

2022

Approved by

First Reader

Barton L. Lipman, PhD
Professor of Economics

Second Reader

Juan Ortner, PhD
Associate Professor of Economics

Third Reader

Chiara Margaria, PhD
Assistant Professor of Economics

Acknowledgments

This PhD has been a long journey, and I owe my greatest gratitude to many people that have helped me arrive to the finish line.

First, I want to thank my main advisor, Bart Lipman. He has been an example of integrity, wisdom, wit, empathy, and eveything an advisor should be. I have learned an incredible amount from my interactions with him. His calm when contradicting something I was saying, his ability to see things deeply, his willingness to change his mind when I was right, his support and encouragement have been invaluable for this journey.

I want to thank Chiara Margaria and Juan Ortner for advising me through the PhD. They spent a lot of time listening to me, discussing proofs, details, suggesting new venues for my ideas. Their support, friendliness and calm have been a big reason why I was able to arrive where I did.

I also want to thank Andy Newman, Albert Ma, Dilip Mookherjee, Michael Manove, and Jawwad Noor for many deep and insightful conversations and suggestions that greatly improved my papers and my understanding of Economic Theory.

I want to thank my friends Vittoria, Pablo Ignacio Guzmán Pinto, Thomas, Stefano, Alden, Hanna, and Lalit for being the best teammates I could hope for during this journey. I also want to thank Fernando, Lisa and Sergio for their friendship, their example as older students, and their support.

I want to thank the participants to the Micro Theory Reading Group and the Game Theory Reading group for the most intriguing discussions and conversations of my PhD.

I thank my roommates Andrea, Maaisha and Martina for accompanying me through rough patches and better times, and being an awesome little community during the pandemic.

I will now switch to Italian.

Grazie alla comunitá italiana di Boston, Michele, Nicola, Giacomo, Francesca, Valeria, Matteo, Roberto per avermi dato un senso di appartenenza e di radici profondo, nonostante fossimo dall'altro capo dell'Atlantico.

Grazie a Marco e Valeria per avermi sempre dato rifugio a Providence, e avermi fatto sentire che avevo una casa a poche miglia di distanza.

Grazie ai miei amici, Claudio, Davide, Pietro, Nicola, Giuseppe, Michele, Gianluca, Gianmarco, Enrica, Giulia, Edoardo, Gabriele per avermi sempre accolto a casa come se non fossi mai partito.

Grazie a Ludovica, per aver reso gli ultimi tre anni del mio dottorato così belli, nonostante le avversità. Grazie per il tuo affetto, per il tuo supporto, per la tua forza, per il tuo incoraggiamento, per aver letto ogni draft di ogni mio paper, per il tuo esempio. Grazie di aver sempre creduto in me.

Grazie ai miei genitori per avermi sempre incoraggiato a perseguire i miei desideri, anche quando questo voleva dire essere a 6 ore di fuso orario di distanza. Grazie per il supporto continuo, l'affetto, la cura che mi mostrate ogni giorno.

Grazie a Anna, Carla, Benedetta, Francesco, Alberto, Tommaso e Carlina per aver reso così gioiosi i miei ritorni a casa.

Thank you for being part of this journey. I do not know if I could make it on my own. But I do know that it has been much happier with you on my side.

Enrico De Magistris

ESSAYS IN ECONOMIC THEORY

ENRICO DE MAGISTRIS

Boston University, Graduate School of Arts and Sciences, 2022

Major Professor: Barton L. Lipman, PhD
Professor of Economics

ABSTRACT

This dissertation consists of two essays in Game Theory with Incomplete Information. The common theme of the essays pertains to the interpretation and origin of incomplete information in strategic situations.

In the first essay, I study bargaining between a buyer and a seller when the buyer can invest in generating outside options at a cost, and the seller cannot observe his investment choice. I model the negotiation phase as an incentive compatible and ex-post individually rational direct mechanism, which maximizes a weighted average of the buyer's and seller's expected payoff. When the weight on the buyer is larger, trade happens with certainty, the price equals the seller's cost, and the buyer does not invest in generating outside options. When the weight on the seller is larger, the optimal mechanism is a posted price mechanism at the worst outside option that the buyer could have, provided that the cost function of the buyer is decreasing in first order stochastic dominance. The probability of trade is strictly below 1, even if it would be socially efficient to trade.

In the second essay, I propose a notion of Rationalizability, called Incomplete Preference Rationalizability, for games with incomplete preferences. Under an appropriate topological condition, the incomplete preference rationalizable set is non-empty and compact. I argue that incomplete orderings can be used to model incom-

plete information in strategic settings. Drawing on this connection, I show that in games with private values the sets of incomplete preference rationalizable actions, of belief-free rationalizable actions (Battigalli et al., 2011; Bergemann and Morris, 2017), and of interim correlated rationalizable actions (Dekel et al., 2007) of the universal type space coincide.

Contents

List of Abbreviations

$\mathcal{B}(X)$	....	Set of Borel measurable subsets of X
BFR	....	Belief-free Rationalizability
$\mathfrak{C}(X)$	....	Set of all Cartesian subsets of X
$\mathrm{Compl}(\succsim)$	....	Set of all completions of $\succsim$
$\mathrm{conv}\, X$	....	Convex hull of the set X
$\overline{\mathrm{conv}}\, X$	....	Closed convex hull of the set X
CDF	....	Cumulative Distribution Function
CS	....	Condorelli and Szentes (2020)
$C(X)$	....	Set of continuous real functions on the set X
$\Delta(X)$	....	Set of Borel probability measures on the set X
$\mathcal{D}[0,1]$	....	Set of CDFs with support contained in $[0,1]$
epIR	....	ex-post Individually Rational
FSD	....	First Order Stochastic Dominance
$\mathrm{Gr}(f)$	....	Graph of the function (or correspondence) f
IC	....	Incentive Compatible
ICR	....	Interim Correlated Rationalizability
IPR	....	Incomplete Preference Rationalizability
$\mathcal{M}(F)$	....	Set of F-feasible mechanisms
MPS	....	Mean Preserving Spread
$\mathrm{supp}\, F$	....	Support of F
PPM	....	Posted Price Mechanism

Chapter 1

Investing in Outside Options in Bargaining

1.1 Introduction

The availability of an outside option in a bargaining situation can improve on the terms of trade for the party with such alternative opportunity. For example, a seller may be willing to offer a discount to a buyer with an alternative price offer. The seller could convince the buyer to trade by slightly undercutting the other offer, if she knew its size. However, if the seller only knew that an outside option exists, the seller may offer a discount larger than that needed to trade.

Traditionally, economic theory has modelled asymmetric information about outside options using an exogenous prior distribution over the possible alternatives that the buyer could have. This approach is well suited to model a situation in which how the buyer obtained his outside option does not impact the outcome of the negotiation (Fudenberg et al., 1987). In some applications, it is more realistic to take a different approach, one in which outside options arise because of some action taken by the buyer. In other words, the buyer may have to exert effort to find a good alternative price, rather than being born with one at his disposal. In addition, the seller may not know either the outside opportunity that the buyer found, or how hard he looked for it.

As an example, consider an employer (the *buyer*, he) negotiating with a job applicant (the *seller*, she). The worker typically does not know whether the employer has interviewed other applicants before negotiating with her, and if he has found an applicant willing to take the job at a low wage. The employer decides which candidates to interview, and how many of them, taking into account two things. First, interviewing and negotiating with many candidates is costly, and there is no guarantee of finding an applicant willing to accept a low wage. Second, the employer has to anticipate what effect will this outside option have on the negotiations with the worker. Specifically, having an alternative offer at hand can offset its cost by benefitting the buyer in two ways during bargaining. On the one hand, the outside option provides him with a worst-case scenario payoff: if the worker adamantly demands a high wage, the employer knows he can resort to the alternative applicant. On the other hand, it creates the opportunity of extracting *information rents*, since the seller does not observe either his effort or the resulting alternative price: the worker may end up with a lower wage than what the employer was ready to agree on. Both effects can improve on the employer's bargaining position, by providing him with a credible threat of interrupting negotiations and allowing him to extract rents from the worker.

This paper studies the effect that the endogeneity of outside options has on the outcome of a negotiation and the distribution of bargaining power in the relation. To this end, I analyse bargaining between a buyer and a seller assuming that, prior to negotiating, the buyer can invest in generating outside options. I assume that the buyer's willingness to pay is larger than the seller's opportunity cost, so that trade is socially efficient. As a result, the buyer's investment in outside options is a form of ex-post opportunistic behavior (Klein et al., 1978; Morita and Servátka, 2018) to extract quasi-rents from the seller. I model investment as choosing a cumulative distribution function (CDF) on $[0, 1]$, where 0 is the opportunity cost of the seller, and 1 is the willingness to pay of the buyer. In the employer-worker example, this modelling choice corresponds to the idea that the employer may use several search strategies to find alternative candidates to interview, such as advertising the position to different populations, attending work fares and so on. More generally, one can think of a distribution over outside option as a (stochastic) technology, an action available to the buyer that generates a random alternative opportunity.

We typically model bargaining between the buyer and the seller as an extensive form game with incomplete information. Such game specifies the whole *bargaining protocol*: which party makes an offer at a given point, any potential exchange of information (via cheap talk), and so on. Sometimes, there are good reasons to study a specific bargaining protocol. For example, a monopolist facing many potential customers has all the bargaining power in a negotiation, which can be modelled by assuming she makes all the offers (as in Gul et al., 1986). If we believe the parties have equal bargaining power, perhaps, it is natural to assume that they alternate in making an offer (as in Rubinstein, 1982).

The choice of a specific bargaining protocol to model the negotiations inevitably adds an assumption about the distribution of bargaining power in the relationship between buyer and seller. It is therefore risky to study the impact of endogenous outside options on bargaining power within a specific protocol, as the conclusion of the analysis might depend on the protocol chosen. Therefore, rather than arbitrarily choosing a game form describing the negotiation process, I use Mechanism Design to represent its equilibrium outcome in reduced form. More precisely, by the *Revelation Principle*, for any equilibrium of any bargaining protocol, there exists a *direct revelation mechanism*, implementing the same allocation as a Bayes-Nash equilibrium in truthful reporting.

Thus, I model the negotiation stage as a direct revelation mechanism, which associates with each outside option that the buyer could report an *outcome*, specifying the probability and the price at which the buyer and the seller trade. In the main model, I assume that the bargaining mechanism is *efficient* in the following sense. Given the distribution chosen by the buyer, neither party can obtain a larger expected payoff in an alternative mechanism without the other party getting a lower payoff. I will refer to this notion of efficiency as *ex-ante Pareto efficiency* for conciseness.[1] This assumption can be interpreted as a re-negotiation proof property: if the bargaining mechanism were inefficient, the buyer and the seller would have an incentive to conduct the negotiations following a different bargaining protocol. Still, in Section 1.6.2 I show that the conclusions of the model are qualitatively the same without this assumption.

An equilibrium of the game is represented by a pair, consisting of the distribution over outside options and a direct revelation mechanism, with the following property. First, the distribution is optimal for the buyer, given the mechanism governing trade at the bargaining stage. Second, the mechanism is ex-ante Pareto optimal given the distribution chosen by the buyer. To find the equilibria of the model, it is useful to consider a (fictitious) game between the buyer and a (fictitious) *mechanism designer*. In this game, the buyer invests in a CDF over outside options, and the designer selects a mechanism to maximize a weighted average of the buyer's and seller's expected payoffs. I emphasize here that this is just a metaphor, albeit a useful one: the true game is between the buyer and the seller, and the mechanism is just a reduced form representation of its equilibrium. Building on the metaphor, I will use expressions like "the designer *favors* the buyer", if the buyer's weight is larger than the seller's, or that a Pareto efficient mechanism for distribution F is a "*best-reply* of the designer to F".

[1] Some care should be taken when interpreting this property in a global sense, since the outside options are interpreted as deals that the buyer is able to reach with an unmodelled third party. If the payoffs of these parties is taken into account in the model, it is possible that there exists a mechanism that is ex-ante Pareto superior, in which the buyer trades with the alternative vendor with whom the trade surplus is the largest.

Perhaps surprisingly, in equilibrium the buyer is unable to extract any information rent from the seller. If the buyer's Pareto weight is larger, in equilibrium he obtains the whole surplus, without having to invest in generating outside options (Theorem 2). If that is the case, there is no efficiency loss, and the social welfare reaches its maximum level. Since this result is robust to all variations of the main model I consider (see Section 1.6), in the rest of the introduction I will only discuss the case in which the seller's Pareto weight is larger. If the seller is favored, then the buyer invests in a random outside option with a two-point support, taking either value 0, or 1. The buyer and the seller only trade when the buyer's alternative price is 1, so that the seller matches the high outside option. When the alternative price is 0, the buyer does not trade with the seller, but rather exercises his outside option. Consequently, the buyer is unable to obtain any information rents from the seller: either the seller extracts his full surplus, or he has to resort to a (costly) outside option (Theorem 3).

The crucial assumption behind this result is that the investment cost in decreasing in First Order Stochastic Dominance (FSD). Using the employer-worker example, this assumption corresponds to the intuitive idea that it is easy for the employer to find someone willing to work for a high wage, whereas finding someone who would accept a low wage may require interviewing many candidates. While this assumption is not enough to fully characterize the equilibrium, it implies (Lemma 1) that the seller extracts all the buyer's surplus whenever there is trade. To characterize the equilibrium, I also assume that the cost function is decreasing in Mean Preserving Spread (MPS), which implies that the buyer invests in a distribution supported on the extreme points of the interval, 0 and 1. Monotonicity in MPS means that a distribution with some variability around its mean should cost more than one which is more concentrated around the same mean. The idea behind it is that the latter distribution is riskier than the former, and so it should cost more.

One may wonder if the absence of information rents is determined by the unobservability of investment. In a similar setting, Condorelli and Szentes (2020) (CS hereafter) show how the buyer generates information rents if he invests in valuations for the good and the seller observes his investment choice before making a take it or leave it offer to the buyer. When the investment generates outside options, instead, the buyer in equilibrium does not obtain any information rents. Rather, he invests in the same two-point distribution he would use under unobservable investment.

Thus, investing in outside options does not improve on the buyer's bargaining position in equilibrium, regardless of the observability of investment. Indeed, if the buyer starts with more bargaining power, which corresponds to the buyer's Pareto weight being larger than the seller's in the model, then in equilibrium the buyer extracts the entire trade surplus regardless of his investment choice. On the other hand, if he starts with less bargaining power, then investing in outside options does little to improve on his position. Neither it generates any information rent, since the seller extracts the entire surplus when the buyer's alternative price is high, nor it strengthens his (credible) threat of walking away from the negotiations when it is low, since in that case the seller is happy to not trade at all. I highlight here that investment is *not* useless for the buyer, since, depending on the realization of the random outside option, the buyer sometimes obtains a larger outside payoff than the inside payoff determined by trade.

Implicit in the analysis is the assumption that the buyer's investment decision is not contractible, corresponding to the idea that the buyer assesses his alternatives before trading with the seller. Non-contractibility generates an inefficiency akin to that arising in the hold-up problem literature. In a typical hold-up problem model, one party – the buyer, say– makes an investment that increases the trade surplus prior to negotiate on how the surplus is split. However, since the investment cost is sunk at the bargaining phase, the seller holds the buyer up, and appropriates (part of) the returns from investment. Anticipating this, the buyer's investment level is below the optimal one, because he does not fully internalize the social benefits from investment.

In contrast, in this paper the buyer invests in generating value outside of the relationship with the seller, resulting in an investment level above the optimal one. Indeed, since trade is socially efficient, any investment in outside options decreases social welfare, that is, the buyer should not exert any effort. This inefficiency arises whether or not the seller observes the buyer's action. To determine the size of this inefficiency in Section 1.5.2 I analyse a variation of the model in which the investment decision of the buyer is contractible. Specifically, the buyer and the seller sign a binding contract prior to the buyer's investment decision, which the seller observes. The contract that the parties sign specifies the rules governing trade that would follow each CDF that the buyer could choose. In the optimal contract, the buyer is rewarded with a low price if he does not invest in outside options, and punished with a price of 1 otherwise. Thus, the buyer has to decide whether to invest in generating outside options, at the cost of not trading with the seller, and not investing and receiving some surplus within the trade relationship. In equilibrium, the buyer is offered a price which makes him indifferent between accepting that price without looking for an outside option, versus investing at the cost of this relationship. The size of inefficiency is determined by the difference between the social welfare under the optimal contract, and the equilibrium social welfare under non-contractible investment, and is proportional to the equilibrium investment cost.

The paper is organized as follows. Section 1.2 reviews the literature. Section 1.3 introduces the main model, with unobservable investment, and the notation. Section 1.4 studies its equilibria. Section 1.5 analyses two variations of the model in which the investment choice is observed by the seller. In particular, 1.5.1 studies what happens if investment is still not contractible, while 1.5.2 analyses the best possible contract assuming that the buyer and the seller can agree on a price before the buyer invests in generating outside options. Section 1.6 studies the robustness of the conclusion of the main model to weakening of the assumptions on the cost function (Section 1.6.1), and the ex-ante Pareto optimality of the mechanism (Section 1.6.2). Section 1.7 concludes.

1.2 Related Literature

This paper contributes to, and combines approaches from, three main strands of the literature: the *hold-up problem*, *bargaining with outside options*, and *mechanism design with endogenous private information*.

Hold-up problem

The hold-up problem has long been recognized as a source of efficiency loss in bargaining and the organization of the firm (Klein et al., 1978; Grossman and Hart, 1986). In its basic version, a party invests less than the efficient level because another party is the residual claimant of the returns from investment. Many hold-up models posit that a buyer invests in his valuation of the good prior to negotiating with the seller. Different papers make different assumptions about how investment translates into valuations. Some assume a deterministic link (Grossman and Hart, 1986), some that the valuation gets determined stochastically (Hermalin and Katz, 2009), while in Condorelli and Szentes (2020)'s model the buyer chooses the (stochastic) technology that generates valuations. In all these examples, the buyer invests less than the socially optimal level. Indeed, he anticipates the fact that the seller will appropriate (part of) the returns from investment, since its cost is sunk at the bargaining stage. Several solutions for the hold-up problem have been proposed over the years: allocation of property rights (Hart and Moore, 1990), contractual solutions, depending on informational asymmetries (Rogerson, 1992), repeated investment through a long relationship (Che and Sákovics, 2004), unobservability of investment together with frequently repeated trade (Gul, 2001), just to name a few. Some papers also consider what happens if the buyer's investment impacts his outside option. In De Meza and Lockwood (1998)'s model, investment impacts both the trade surplus that can be achieved (which determines the buyer's *inside* option) as well as what he could obtain by not trading with the seller (his *outside* option). Indeed, as my paper shows, investment in outside options can be seen as a form of rent-seeking behavior, and the equilibrium investment level is above the socially optimum level.

Bargaining with outside options

While outside options have long been recognized as a source of bargaining power, incorporating them in formal models presents difficulties and ambiguities.

For instance, a seller may have an outside option either because she could trade with an alternative partner, or because she could consume the good herself (Fudenberg et al., 1987; Shaked and Sutton, 1984). Despite these ambiguities, outside opportunities have been used in a variety of settings. For example, they neutralize the effects of crazy types *à la* Abreu and Gul (2000) in Compte and Jehiel (2002), and imply the failure of the Coase conjecture (Coase, 1972) in Board and Pycia (2014). While all these papers assume that the outside option is constant over time, this needs not be the case. For example, traders in Wolinsky (1987) can interrupt the negotiation to look for an alternative partner; the buyer in Hwang and Li (2017) receives the outside option at a random time during the bargaining process; and the outside option itself changes stochastically over time in McClellan (2021). I contribute to this literature by showing how the investment in outside options impacts the bargaining outcome when the bargaining process is ex-ante efficient.

Mechanism design with endogenous private information

The Harsanyi type space approach models incomplete information as an *exogenous* component of the model, and has been used in most of the classical papers in mechanism design (e.g. Myerson, 1981; Myerson and Satterthwaite, 1983). When treating information as an endogenous component, at least two approaches emerge.

Under the first approach, the state of the world is an exogenous random variable, and agents can choose a signal containing information about it. In the early literature this choice was constrained in several possible ways. For instance, in Persico (2004), there is a unique signal that the agent can acquire, or not acquire. In Shi (2012), the signals that can be chosen are ordered by the monotone likelihood ratio property.[2] The recent literature has mostly focused on *fully flexible* information acquisition, assuming that the agent can choose *any* signal correlated with the state of the world. For example, in Roesler and Szentes (2017) and Ravid et al. (forthcoming) the state of the world is the buyer's valuation of the good, and the buyer can choose how precisely to learn about it before trading with the seller. Mensch (2020) analyses how a designer can influence the learning strategy of the agent by announcing a mechanism prior to his choice of a signal.

Under the second approach, the state of the world is endogenously determined by the agents' actions. Private information arises either because the principal does

[2] Bergemann and Välimäki (2006) surveys the early literature on information acquisition in mechanism design.

not observe the agent's action, as in the auction model of Gershkov et al. (2021), or because the mapping from actions to states is stochastic, and only the agent observes the realized state (as in Condorelli and Szentes, 2020). This is also the approach of the current paper.

I contribute to this literature by studying how *endogenous* private information about outside options differs from endogenous private information about valuations in a bargaining model.

1.3 The Model

1.3.1 Mathematical preliminaries and notation

A *cumulative distribution function* (CDF) on $\mathbb{R}$ is a weakly increasing and right continuous function such that $\lim_{x\to-\infty} F(x) = 0$, $\lim_{x\to\infty} F(x) = 1$. The *support* of F is the set (see Chung, 2001)

$$\operatorname{supp} F = \{x \in \mathbb{R} : \forall \epsilon > 0, F(x+\epsilon) - F(x-\epsilon) > 0\}.$$

The set of all CDFs with support contained in $[0,1]$ is denoted $\mathcal{D}[0,1]$, and is endowed with the L_1 distance:

$$d_1(F,G) = \int_0^1 |F(x) - G(x)| \mathrm{d}x.$$

Classical results by Machina (1982) show that d_1 metrizes the weak convergence topology of probability measures, so that $(\mathcal{D}[0,1], d_1)$ is a compact and convex subset of $L_1[0,1]$, the vector space of all real valued functions on $[0,1]$ whose absolute value is Lebesgue integrable.

Given two CDFs $F, G \in \mathcal{D}[0,1]$, say that F *first order stochastically dominates* G, written $F \succsim_{FSD} G$ if for all $x \in [0,1]$, $F(x) \leq G(x)$. Equivalently (Quirk and

Saposnik, 1962), $F \succsim_{FSD} G$ if and only if for all (weakly) increasing functions $u : [0,1] \to \mathbb{R}$,

$$\int_0^1 u(x)\mathrm{d}F(x) \geq \int_0^1 u(x)\mathrm{d}G(x).$$

Say that F is a mean preserving spread of G, written $F \succsim_{MPS} G$ if for all $x \in [0,1]$,

$$\int_0^x F(y) - G(y)\mathrm{d}y \geq 0,$$

with equality at $x = 1$. Equivalently (Rothschild and Stiglitz, 1970), $F \succsim_{MPS} G$ if and only if $\mathbb{E}_F b = \mathbb{E}_G b$ and, for all (weakly) increasing and (weakly) concave $u : [0,1] \to \mathbb{R}$,

$$\int_0^1 u(x)\mathrm{d}F(x) \leq \int_0^1 u(x)\mathrm{d}G(x).$$

Finally, I use $H_b \in \mathcal{D}[0,1]$ to denote the CDF concentrated at $b \in [0,1]$:

$$H_b(x) = \begin{cases} 0 & x < b, \\ 1 & x \geq b. \end{cases}$$

It is easy to see that H_b is a mean preserving *contraction* of any distribution with average b. On the opposite extreme of the MPS order there is the distribution supported on $\{0,1\}$, $(1-b)H_0 + bH_1$, which is a mean preserving spread of any other CDF with mean b. These two CDFs play an important role in the analysis of Section 1.4.

1.3.2 Model

A buyer (he) and a seller (she) negotiate for the ownership of an indivisible good. The buyer's valuation, equal to 1 and the seller's marginal cost, equal to 0, are common knowledge. Prior to the negotiation, the buyer can exert effort in order to obtain an *outside option*, which is interpreted as a price offer from an alternative vendor.[3] The seller does not observe the effort exerted by the buyer or the value of the outside option, when the negotiation starts.

The game is divided in two parts: an investment phase and a bargaining phase. In the *investment phase*, the buyer chooses a costly cumulative distribution function (CDF) over outside options. The buyer can choose any distribution function F from some collection of CDFs, $\mathcal{F} \subseteq \mathcal{D}[0,1]$. A cost function $C : \mathcal{D}[0,1] \to \mathbb{R}_+$ associates with each distribution F the cost of investing in it, $C(F)$. Assuming that C be defined on the whole space $\mathcal{D}[0,1]$ is without loss of generality, as one can simply assign an arbitrarily large value to $C(G)$ whenever $G \notin \mathcal{F}$. A *technology* is a pair composed of the cost function C, and the set of distributions available to the buyer, $\mathcal{F}$.

For an illustration, consider again the employer-job applicant example of the Introduction. Assume that there is a large mass of candidates that the employer could interview, and a corresponding distribution of wages, $\Phi \in \mathcal{D}[0,1]$, that they would accept. For simplicity, the employer decides the number n of candidates to interview, so that his outside option is the lowest wage that another candidate would accept. Thus, the buyer could generate the distribution of the minimum of n random draws from Φ, which I denote by Φ_n. The cost of distribution Φ_n is the opportunity cost of the employer of making n interviews.

Throughout the paper, I will maintain three minimal assumptions on the technology $(C,\mathcal{F})$. First, I assume that $\mathcal{F}$ is closed and C, restricted to $\mathcal{F}$, is continuous. Continuity of $C_{|\mathcal{F}}$ corresponds to the idea that similar distributions should have similar cost. Notice that, since $\mathcal{D}[0,1]$ is compact, so is $\mathcal{F}$. Second, I assume that $\mathcal{F}$ is a convex set, and $\mathcal{C}_{|\mathcal{F}}$ a convex function. This is just a *reduction of compounded lot-*

[3] It is easy to specify the model alternatively, by assuming that b is the outside option payoff that the buyer receives, rather than an alternative price.

teries assumption: if F and G are both elements of $\mathcal{F}$, then the buyer can generate the distribution $\gamma F + (1 - \gamma)G$ by playing a mixed action, choosing F with probability γ and G with probability $(1 - \gamma)$, so that assuming that $\gamma F + (1 - \gamma)G \in \mathcal{F}$ is without loss of generality. It is also without loss to assume that $C_{|\mathcal{F}}$ is convex, since otherwise the buyer could generate the distribution F by playing the mixed action with the least expected cost that generates F. Finally, I assume that the buyer can always generate an outside option equal to his valuation, at 0 cost. This is modelled by assuming that the distribution concentrated at 1, denoted H_1, is an element of $\mathcal{F}$ and $C(H_1) = 0$.

Once distribution F has been chosen, the buyer observes the realization of a random variable distributed according to F, which represents his outside option. The seller observes neither the investment choice of the buyer, nor the realized outside option. At this point, the game moves to the next stage: the *negotiation phase*. As mentioned in the introduction, in principle, the negotiation phase could be modelled as an extensive form game with incomplete information, which would specify the whole protocol. Instead, I use a reduced form approach, and model bargaining as a *direct revelation mechanism*, which satisfies three conditions. First, the mechanism is ex-ante Pareto optimal. The second requirement is *Incentive Compatibility*: the buyer should prefer reporting his outside option b truthfully rather than lying and reporting another value b'. The third and final requirement is *Ex-post Individual Rationality*: the seller should prefer trading to not trading, the buyer should prefer trading to exercising his outside option, whenever trade has positive probability. This condition corresponds to the idea that trade is always voluntary, and the buyer could interrupt the negotiations to exercise his outside option, if that benefits him. To find the equilibria of the game between the buyer and the seller, I study an auxiliary game between the buyer, who chooses a distribution over outside options, and a (fictitious) mechanism designer (she), who chooses the mechanism governing trade. The payoff function of the designer is a weighted average between the payoffs of the buyer and the seller. The equilibria of this auxiliary game correspond to the (Pareto efficient) equilibria of the investment-then-bargaining game between the buyer and the seller.

The set of actions of the designer in the auxiliary game is the set of direct revelation mechanisms that are incentive compatible and ex-post individually rational. A *revelation mechanism* is composed of two elements: a closed set X of possible reports by the buyer, and a pair of functions $(q, p) : X \to [0, 1]^2$. The value $q(x)$ specifies the probability of trade if the buyer reports x, whereas $p(x)$ is the price that the buyer pays to the seller if trade happens.[4] In a *direct* revelation mechanisms the set of reports X is the subset of $[0, 1]$ of all types deemed possible by the designer's belief.[5] By reporting $x \in X$, a buyer with outside option b obtains an expected payoff of

$$q(x)(1 - p(x)) + (1 - q(x))(1 - b) = 1 - b + q(x)(b - p(x)),$$

whereas the seller's payoff is

$$q(x)p(x).$$

A direct revelation mechanism defined on $X \subseteq [0, 1]$ is *incentive compatible* (IC) if truthful reporting is optimal for the buyer: for each $b, b' \in X$,

$$1 - b + q(b)(b - p(b)) \geq 1 - b + q(b')(b - p(b')).$$

It is *ex-post individually rational* (epIR) if players prefer trading to not trading when-

[4]Remember that the mechanism has to be ex-post individually rational, which is why $p(x) \in [0, 1]$.

[5]More precisely: X is a closed subset of $[0, 1]$ which coincides with the support of the belief of the designer.

ever the probability of trade is positive: for each $b \in X$,

$$q(b) > 0 \implies 0 \leq p(b) \leq b.$$

The designer selects an incentive compatible and ex-post individually rational mechanism to maximize a weighted average of the traders' payoff, given her belief about the distribution of the buyer's outside options. Let $\tilde{F}$ be the CDF that the designer *thinks* that the seller has selected, and let $X = \operatorname{supp} \tilde{F}$. The designer's expected payoff from selecting an incentive compatible and ex-post individually rational mechanism $(q, p) : X \to [0, 1]^2$ is

$$W_\alpha(q,p) = \alpha \underbrace{\int_0^1 q(b)p(b)\mathrm{d}\tilde{F}(b)}_{\text{Seller's payoff}} + (1-\alpha) \underbrace{\int_0^1 \Big(1 - b + q(b)\big(p(b) - b\big)\Big)\mathrm{d}\tilde{F}(b)}_{\text{Buyer's payoff}},$$

where $\alpha \in [0, 1]$ is a parameter capturing how much the designer "cares" about the seller. The *best reply* of the designer to the buyer choosing distribution $\tilde{F}$ is the solution of the following program:

$$\max_{(q,p):X\to[0,1]^2} \quad \alpha \int_0^1 q(b)p(b)\mathrm{d}\tilde{F}(b) + (1-\alpha) \int_0^1 \Big(1 - b + q(b)\big(p(b) - b\big)\Big)\mathrm{d}\tilde{F}(b),$$

$$\text{s.t. } \forall b, b' \in X, \quad 1 - b + q(b)(b - p(b)) \geq 1 - b + q(b')(b - p(b')),$$

$$\forall b \in X, \quad q(b) > 0 \implies 0 \leq p(b) \leq b.$$

The buyer has to take a decision at three points in the game. First, he has to select a distribution $F \in \mathcal{F}$. Second, once he has observed his outside option and he has been offered a mechanism, he has to decide whether or not to participate. Finally, provided he takes part in the mechanism, the buyer chooses which report to communicate to the designer.

Consider the buyer's problem, starting from this last decision. Suppose the buyer's outside option is b, and that he's participating in the mechanism $(q, p) : X \to [0, 1]^2$, where X is a closed subset of $[0, 1]$. The buyer chooses report $x \in X$ to solve

$$\max_{x \in X} 1 - b + q(x)(b - p(x)).$$

If the buyer's outside option, b is deemed possible by the mechanism (that is, $b \in X$), and if the mechanism is incentive compatible, then b solves this maximization problem. However, since the designer does not observe the distribution chosen by the buyer, it is possible that $b \notin X$; that is, b may not be one of the types that the designer thought possible when choosing the mechanism. In that case, the solution to the maximization problem will be some element $b' \in X$. It is also possible that participating in the mechanism is suboptimal for the buyer as no report gives an expected payoff larger than $1 - b$. That is, if

$$1 - b > \max_{x \in X}\{1 - b + q(x)(b - p(x))\},$$

then the buyer is better off not taking part in the mechanism.

At the time of the first decision, the buyer does not know what mechanism will be offered by the designer. So suppose he believes that mechanism $(q, p) : X \to [0, 1]^2$ will be the designer's choice. Then the largest expected payoff the buyer can obtain after playing distribution $F \in \mathcal{F}$ is

$$\int_0^1 \max\Big\{\max_{x \in X}\big\{1 - b + q(x)\big(b - p(x)\big)\big\}, 1 - b\Big\} \mathrm{d}F(b) - C(F).$$

The *best reply* of the buyer to mechanism $(q, p) : X \to [0, 1]^2$ is the solution to the following program:

$$\max_{F\in\mathcal{F}} \int_0^1 \max\left\{ \max_{x\in X} \left\{1-b+q(x)\big(b-p(x)\big)\right\}, 1-b\right\} \mathrm{d}F(b) - C(F).$$

Definition 1. An *Equilibrium* is a pair $(F^*, (q^*, p^*))$ such that

- $(q^*, p^*) : \operatorname{supp} F^* \to [0,1]^2$ is incentive compatible and ex-post individually rational,
- (q^*, p^*) is a best reply to F^* for the designer, and
- F^* is a best reply to (q^*, p^*) for the buyer.

The first formal result of the paper shows that, without loss of optimality, one can restrict attention to the class of *posted price mechanisms*. More precisely, the designer's best reply to CDF F always contains at least one posted price mechanism. This is an extension of classical results in the literature to allow for arbitrary cumulative distribution functions.

Definition 2. Let X be a closed subset of $[0,1]$. A posted price mechanism (on X) is a mechanism $(q,p) : X \to [0,1]^2$ such that there exists $\tilde{p} \in [0,1]$, and $\tilde{x} \in X$ with

$$q(x) = \begin{cases} 1 & x \geq \tilde{x} \\ 0 & \text{else,} \end{cases} \qquad \text{and} \qquad p(x) = \begin{cases} \tilde{p} & x \geq \tilde{x} \\ 0 & \text{else.} \end{cases}$$

If either $\tilde{p} = \tilde{x}$, or $\tilde{p} = 0$ and $\tilde{x} = \min\{x \in X\}$, then the mechanism is called a *posted price at* $\tilde{p}$.

Theorem 1. *Fix* $F \in \mathcal{D}[0,1]$. *Then there exists* $p \in \operatorname{supp} F$ *such that the posted price mechanism at* p *is a best reply to* F.

Proof. See the Appendix. □

So by Theorem 1, it is without loss to assume that the Designer posts a price, and the buyer decides whether or not to trade at that price. Indeed, suppose that the designer plays a posted price mechanism at $p \in X$, and the buyer's outside option is b. Then the buyer is willing to trade at price p if and only $1-b \leq 1-p$, or, equivalently, $p \leq b$. This is true whether b is an element of X, the set of outside options that the Designer deems possible, or not. Indeed, if $p \leq b$, then the buyer reports either b, if $b \in X$, or any type in X which is allowed to trade, if $b \notin X$. If $p > b$, then the buyer can either report b (if $b \in X$), or report a type in X who doesn't trade, or even refuse to participate in the mechanism altogether. Therefore, from now I will simplify the terminology, and say that the Designer selects a price $p \in [0,1]$.

When facing a posted price of p, the buyer's expected payoff from selecting distribution F is

$$\int_0^1 \max\{1-p, 1-b\} \mathrm{d}F(b) - C(F) = \int_0^p (1-b)\mathrm{d}F(b) + \int_p^1 (1-p)\mathrm{d}F(b) - C(F),$$

which equals

$$1 - p + \int_0^p F(b)\mathrm{d}b - C(F),$$

by integration by parts. On the other hand, the Designer's expected payoff from posting price p when the buyer chooses distribution F is

$$\alpha p(1 - \lim_{b' \uparrow p} F(b)) + (1-\alpha)\Big(1 - p + \int_0^p F(b)\mathrm{d}b\Big).$$

Thus, a pair (F^*, p^*) is an *equilibrium* if

1. Distribution F^* is optimal for the buyer, given that the Designer posts price p^*. That is, F^* is a solution to

$$\max_{G \in \mathcal{F}} 1 - p^* + \int_0^{p^*} G(b)\mathrm{d}b - C(G).$$

2. The posted price mechanism at p^* is optimal for the Designer, if the distribution of outside options of the buyer is F^*. That is, p^* solves

$$\max_{p\in[0,1]} \alpha p(1-\lim_{b\uparrow p} F^*(b)) + (1-\alpha)\Big(1-p+\int_0^p F^*(b)\mathrm{d}b\Big).$$

Before moving to the analysis of the general model in the next Section, consider again the employer-job applicant example. The employer decides the number of candidates to interview before negotiating with the job applicant. Suppose for simplicity that interviewing n candidates has an associated cost kn, where $k \in \mathbb{R}$ is a constant. If the employer interviews n candidates, and the i-th one asks for wage b_i, then the employer's outside option is $b = \min_{i=1,\dots,n} b_i$. Therefore, this process generates the distribution

$$\Phi_n(b) = \Pr{}_\Phi(\min_{i=1,n} b_i \leq b) = 1-(1-\Phi(b))^n,$$

where Φ is the underlying distribution of wages asked by the alternative candidates. If the buyer invests in a sample of size 0, his outside option is 1 with probability 1. So after accounting for the possibility of mixed strategies of the buyer, the set of distributions that the buyer can choose from is

$$\mathcal{F} = \overline{\mathrm{conv}}\left(\{1-(1-\Phi(\cdot))^n\}_{n\in\mathbb{N}} \cup \{H_1\}\right),$$

where $\overline{\mathrm{conv}}(S)$ is the closed convex hull of the set S. A (pure) equilibrium is then constituted by a pair (Φ_{n^*}, p^*) where n^* solves the buyer's problem when the Designer posts price p^*:

$$n^* \in \mathrm{argmax}_{n\in\mathbb{N}}\, 1 - \int_0^{p^*} (1-\Phi(b))^n \mathrm{d}b - kn,$$

and p^* solves the Designer's problem when the buyer's sample contains n^* prices:

$$p^* \in \operatorname{argmax}_p \alpha p(\Phi^{n^*}(p)) + (1-\alpha)\Big(1 - \int_0^p (1-\Phi(b))^{n^*} \mathrm{d}b\Big).$$

As an illustration, consider the case where the underlying distribution is a binomial distribution on $\{0,1\}$, with $\varphi = \Pr_\Phi(1)$. That is,

$$\Phi(b) = \begin{cases} 0 & b < 0, \\ 1-\varphi & 0 \leq b < 1, \\ 1 & b \geq 1. \end{cases}$$

Suppose that the Designer cares more about the seller than the buyer ($\alpha < \frac{1}{2}$). Since, regardless of the sample dimension of the buyer, his outside option will be either 0 or 1, in any equilibrium the Designer chooses price $p^* = 1$. This means that the equilibrium sample dimension, n^*, must solve

$$\max_n 1 - \int_0^1 \big(1 - (1-\varphi)^n\big) \mathrm{d}b - kn = \max_n 1 - \varphi^n - kn.$$

The function $f(x) = 1 - \varphi^x - kx$ is concave on $\mathbb{R}$, and is maximized at

$$x^* = \log\Big(\frac{k}{-\log\varphi}\Big)\frac{1}{\log\varphi}.$$

Therefore, the solution to the buyer's problem will either be $n = \lfloor x^* \rfloor$ or $n = \lceil x^* \rceil$, where $\lfloor x \rfloor$ is the largest integer below x, and $\lceil x \rceil$ is the smallest integer above x.

1.4 Analysis and Equilibrium

In order to fully characterize the equilibrium of the model, I impose three assumptions. First, I assume that the buyer can choose *any* distribution supported on

$[0, 1]$.

Assumption 1. $\mathcal{F} = \mathcal{D}[0, 1]$.

Second, I assume that the cost function C is decreasing in first order stochastic dominance. This corresponds to the idea that it is easy to find someone willing to sell an object for a high price, and hard to find a low price. Therefore, a distribution over alternative prices which assigns most mass to large outside options should cost less than one assigning most mass to low outside options.

Assumption 2. C is strictly decreasing in First Order Stochastic Dominance: if $F \succsim_{FSD} G$ then $C(F) \leq C(G)$, and the inequality is strict whenever $F \neq G$.

Third, the cost function is decreasing in Mean Preserving Spread. Since a decrease in mean preserving spread is associated with a decrease in the risk of the investment, the assumption simply says that a riskier investment choice should cost less than a safe one, *ceteris paribus*.

Assumption 3. C is strictly decreasing in Mean Preserving Spread: if F is a mean preserving spread of G, then $C(F) \leq C(G)$, and the inequality is strict whenever $F \neq G$.

A consequence of Assumption 2 is that H_1, the distribution concentrated on 1, is the only distribution which costs 0. The following three simple examples show that the requirements of Assumptions 1-3 can be satisfied by simple functional forms.

Example 1. Let $\tilde{\kappa} : [0, 1] \to \mathbb{R}_+$ be strictly decreasing and strictly concave, with $\tilde{\kappa}(1) = 0$. Then the function $\tilde{C} : \mathcal{D}[0, 1] \to \mathbb{R}_+$ defined by

$$\tilde{C}(F) = \int_0^1 \tilde{\kappa}(b)\mathrm{d}F(b)$$

is continuous (by Portmanteau theorem), and satisfies Assumptions 2 and 3. By linearity of the Riemann-Stieltjes with respect to the integrator, the cost function C is affine (hence, convex).

Example 2. Let $h : \mathbb{R}_+ \to \mathbb{R}_+$ be strictly increasing and strictly convex, with $h(0) = 0$, and define

$$\hat{C}(F) = h\Bigg(\int_0^1 \hat{\kappa}(b)\mathrm{d}F(b) \Bigg),$$

where $\hat{\kappa}$ is strictly decreasing, strictly concave and satisfies $\hat{\kappa}(1) = 0$ as in Example 1. Then $\hat{C}$ satisfies Assumptions 2 and 3, and is strictly convex:

$$\begin{aligned} \hat{C}(\lambda F + (1-\lambda)G) &= h\Bigg(\int_0^1 \hat{\kappa}\mathrm{d}\big(\lambda F + (1-\lambda)G\big) \Bigg), \\ &= h\Bigg(\lambda \int_0^1 \hat{\kappa}\mathrm{d}F + (1-\lambda) \int_0^1 \hat{\kappa}\mathrm{d}G \Bigg), \\ &< \lambda \hat{C}(F) + (1-\lambda)\hat{C}(G). \end{aligned}$$

Example 3. Let $\ell : [0,1] \to \mathbb{R}_+$ be continuous, convex, strictly decreasing and such that $\ell(1) = 0$. Then the function

$$\mathcal{L}(F) = \ell\Big(\int_0^1 b\mathrm{d}F(b) \Big)$$

satisfies Assumption 2, and is constant in Mean Preserving Spread: if $F \succsim_{MPS} G$, then $\mathcal{L}(F) = \mathcal{L}(G)$. The cost function $\mathcal{L}$ is said to be *mean based*.

In the rest of this Section, I characterize the equilibrium of the game as a function of the parameter α, capturing how much the Designer cares about the seller, relative to the buyer.

1.4.1 Case I: the Designer Cares More About the Buyer ($\alpha < \frac{1}{2}$)

Since the buyer's valuation is larger than the seller's cost, trade is socially optimal in this model. So, if the designer posts a price that ensures that trade happens with probability 1, then the social optimum would be achieved. If the buyer's outside option is distributed according to F, then probability 1 trade can be achieved only by posting a price smaller than $\min\{\operatorname{supp} F\}$. Moreover, the buyer may have an incentive to look for an outside option that gives him a *larger* payoff than that granted by the posted price, especially if he expects it to be high. As an extreme example, if the buyer expects the designer to post price 1, then he would have a strong incentive to look for an alternative price that does not extract his entire willingness to pay.

One possibility for the designer to achieve the social efficient allocation is to post price 0. Indeed, the buyer would effortlessly attain his maximum payoff possible, subject to the seller being individually rational, and so he would have no incentive to look for an alternative vendor.

Theorem 2. *If the Designer cares more about the buyer ($\alpha < \frac{1}{2}$), then in the unique equilibrium, the posted price is 0, the buyer chooses the distribution concentrated on 1, H_1, and trade happens with probability 1.*

Proof. First, I will show that this is an equilibrium. If the Designer posts a price of 0, then the only best reply of the buyer is to choose the cheapest distribution available: no alternative price could possibly be below 0, so the buyer should exert no effort, and choose distribution H_1. Because of Assumption 2, any other distribution that the buyer could choose costs more than $0 = C(H_1)$, and therefore H_1 is the unique best reply of the buyer. If the buyer chooses the point mass at 1, then any price p of the Designer would have trade happening with probability 1. The Designer's payoff from posting price p would be

$$\alpha p + (1-\alpha)(1-p) = (1-\alpha) + p(2\alpha - 1) < \alpha \cdot 0 + (1-\alpha)\cdot 1,$$

and so the Designer would have no incentive to deviate.

As for uniqueness: suppose by contradiction that $(\hat{F}, \hat{p})$ is another equilibrium, with $\hat{p} > 0$, so that the probability of trade is $\hat{q} = 1 - \hat{F}(\hat{p})$. I claim that $p = 0$ is a profitable deviation. Indeed, the Designer's payoff in the candidate equilibrium

$(\hat{F}, \hat{p})$ is

$$(2\alpha - 1)\hat{q}\hat{p} + (1-\alpha)\hat{q} + (1-\alpha)\int_0^{\hat{p}} b \mathrm{d}\hat{F}(b).$$

The first term in the summation is negative, so putting $p = 0$ would (strictly) improve on it. The sum of the last two terms is dominated by

$$(1-\alpha)\big(1 - \hat{F}(\hat{p})(1-\hat{p})\big) \leq (1-\alpha)\big(1 - \hat{F}(0)(1-0)\big) = 1 - \alpha,$$

so that $p = 0$ is a profitable deviation, a contradiction. □

As the second part of the proof of Theorem 2 shows, if $\alpha < \frac{1}{2}$ then in any equilibrium the designer posts price 0. The equilibrium uniqueness depends on Assumption 2, which guarantees that H_1 is the only distribution with 0 cost, and so is the only best reply to a price of 0. If Assumption 2 is violated, then it is possible that another distribution – say, G – has zero cost, in which case $(G, 0)$ would be another equilibrium. In such equilibrium too, trade would happen with certainty.[6]

1.4.2 Case II: the Designer Cares More About the Seller ($\alpha > \frac{1}{2}$)

The intuition behind Theorem 2 is simple: the Designer cares more about the privately informed party, so she gives him as much surplus as possible. Matters are not so simple when the Designer cares more about the uninformed party, however. In this case, the classical static monopoly trade-off between setting a high price and having a high probability of trade emerges. When optimally balancing this trade-off, the Designer must try to correctly guess which was the distribution over outside options that the buyer chose.

The problem is greatly simplified by the following crucial observation: any two types of the buyer who are willing to trade have the same ex post payoff. To be more

[6] I'm implicitly assuming here that a buyer with an outside option of 0 trades with the seller if offered price 0.

precise, suppose that the posted price is $p \in [0, 1)$, and consider two outside options $b > b' > p$. Then both these types accept the price they are offered, trade, and obtain a payoff of $1 - p$. But since b and b' have the same payoff, the buyer would rather concentrate all the probability on b, the largest of these two values. This would decrease the cost of the distribution chosen by the buyer, since the resulting CDF would FSD-dominate the starting one.

Lemma 1. *Under Assumptions 1 and 2, if F is a best reply of the buyer to $p \in (0, 1)$, then F assigns zero measure to the interval $[p, 1)$, or, equivalently, $F(p) = \lim_{b \uparrow 1} F(b)$.*

Proof. Let F be a best reply to p, and suppose by contradiction that F puts some probability mass between p and 1. The buyer's expected payoff is

$$\int_0^p (1 - b)\mathrm{d}F(b) + (1 - p)(1 - F(p)) - C(F).$$

Consider the distribution $\hat{F}$ defined by

$$\hat{F}(b) = \begin{cases} F(b) & b \leq p, \\ F(p) & p < b < 1, \\ 1 & b \geq 1. \end{cases}$$

In plain English, $\hat{F}$ coincides with F below p, is constant and equal to $F(p)$ on $[p, 1)$, and has an atom of size $1 - F(p)$ at 1. Clearly, $F \succsim_{FSD} \hat{F}$, so that, by Assumption 2, $C(F) > C(\hat{F})$. But then $\hat{F}$ has a larger payoff than F for the buyer, a contradiction. □

The characterization of the equilibrium is basically a corollary to Lemma 1. Indeed, if the buyer has type b and accepts a price offer of p, then $b = 1$, so that the mass of traders is concentrated on an atom at 1. Consequently, the designer, who favors the seller, has an incentive to post price 1. Thus, if an equilibrium exists, then the equilibrium price equals the buyer's valuation of the good, 1. So, when $\alpha > \frac{1}{2}$, the buyer's expected payoff from trading with the seller is 0. His best reply is the distribution F that maximizes

$$\int_0^1 F(b)\mathrm{d}b - C(F) = 1 - \mathbb{E}_F b - C(F).$$

Notice that the buyer's expected payoff in this case only depends on F through its mean, and its cost. Since the buyer is risk neutral, he chooses the cheapest distribution among those that have the same mean. So, define a function $c : [0,1] \to \mathbb{R}_+$ by

$$c(\mu) = \min\{C(F) : \mathbb{E}_F b = \mu\}.$$

The value $c(\mu)$ is the least cost needed for the buyer to generate an average outside option μ.[7] It is easy to see that the function c is convex (by convexity of C) and strictly decreasing (since C is decreasing in First Order Stochastic Dominance). This suggests that the best reply to $p = 1$ can be found by solving two problems. The first problem is to compute the best mean outside option μ^*, i.e.,

$$\mu^* \in \mathrm{argmax}_{\mu\in[0,1]}\, 1 - \mu - c(\mu).$$

The second problem is to find a distribution G^* that has mean μ^* and achieves the smallest cost among CDFs with that mean, i.e. such that $C(G^*) = c(\mu^*)$. The solution to this second problem is the distribution supported on $\{0,1\}$ with mean μ^*, $(1-\mu^*)H_0 + \mu^* H_1$. Indeed, this distribution is a mean preserving spread of any CDF F with average μ^*, and so, by Assumption 3, has a cost lower than $C(F)$. Notice that the best price that the Designer could offer against this two-point distribution is indeed 1. This discussion is summarized in the second main result of the paper.

Theorem 3. *Suppose that the cost function C satisfies Assumptions 1, 2, and 3. If the Designer cares more about the seller $\left(\alpha > \frac{1}{2}\right)$, then the equilibrium posted price is 1.*

[7] The minimization problem in the formula for $c(\mu)$ is well defined, since C is continuous and the set of all CDFs that have average μ is a compact subset of $\mathcal{D}[0,1]$.

The buyer chooses a distribution supported on $\{0,1\}$*,* $(1-\mu^*)H_0 + \mu^* H_1$*, where the mean* μ^* *solves*

$$\max_{\mu \in [0,1]} 1 - \mu - c(\mu).$$

The equilibrium of Theorem 3 is not necessarily unique, however. While in any equilibrium the price must be 1, there may be several different means $\hat{\mu}$ that solve

$$\max_{\mu \in [0,1]} 1 - \mu - c(\mu),$$

since the function c is convex, but possibly not strictly convex. For each $\hat{\mu}$, the distribution supported on $\{0,1\}$ with mean $\hat{\mu}$ is a best reply to price 1, and so $\big((1-\hat{\mu})H_0 + \hat{\mu}H_1,\ 1\big)$ is an equilibrium. Consequently, the set

$$\operatorname{argmax}_{\mu \in [0,1]} 1 - \mu - c(\mu)$$

is a compact (possibly degenerate) interval $[\mu_l^*, \mu_h^*] \subseteq [0,1]$.[8] Thus, the set of equilibria either contains a single element, or contains a continuum of elements indexed by the mean $\hat{\mu} \in [\mu_l^*, \mu_h^*]$ that the buyer chooses. Notice that the buyer obtains the same payoff,

$$1 - \hat{\mu} - c(\hat{\mu}) = 1 - \mu_h^* - c(\mu_h^*)$$

in each of these equilibria. The designer's payoff in equilibrium, instead, is

$$\alpha\hat{\mu} + (1-\alpha)(1-\hat{\mu}) = (2\alpha - 1)\mu + (1-\alpha),$$

which is strictly increasing in $\hat{\mu}$. Thus, the equilibria of the game are Pareto-ranked, and the Pareto-best one is the equilibrium in which the buyer chooses $(1-\mu_h^*)H_0 + \mu_h^* H_1$.

Before concluding the Section, consider again the cost function $\tilde{C}$ of Example

[8] Compactness follows from Berge's Maximum Theorem. It is an interval since the set of maximizers of a concave function is convex.

1. Then the optimal mean outside option that the buyer targets, $\tilde{\mu}^*$, is either 0 if $\tilde{\kappa}(0) \leq 1$ or 1 otherwise. With the cost function $\hat{C}$ of Example 2, $\hat{\mu}^*$ must satisfy the first order condition

$$0 \in \partial\Big(1 - \hat{\mu}^* - h\big((1-\hat{\mu}^*)\hat{\kappa}(0)\big)\Big),$$

where $\partial(g(x))$ is the superdifferential of g at x.

1.5 Observable Investment

The model of Section 1.3 assumes that the investment choice of the buyer is neither observable by the Designer, nor contractible. In equilibrium, unless the designer favors the buyer, a hold-up inefficiency arises: the buyer invests too much in generating outside options, and the probability of trade is below 1. In this section, I analyse two variations of the model. In the first one, I assume that the seller can observe the distribution chosen by the buyer, but not the realized outside option. I show that, even in this case, in equilibrium the buyer chooses the same distribution supported on $\{0, 1\}$ as in Theorem 3. Thus, the hold-up inefficiency arises even if the seller observes the investment decision of the buyer. Second, I study how large is the welfare loss due to hold-up. To this end, I analyse a version of the model in which investment is contractible, so the buyer and the seller sign a contract prior to the buyer's choice. In order to avoid that moral hazard issues affect the analysis, I also assume that the seller is able to observe the buyer's choice, so the contract specifies an outcome for each distribution that the buyer could choose.

1.5.1 Non-contractible Investment

(Throughout this section, assume that Assumption 1 holds.) A striking implication of Theorem 2 is that the buyer does not obtain any *information rent* in equilibrium, despite being privately informed about his choice and the value of his outside option. Indeed, either there is trade and the seller extracts the whole consumer surplus, or there is no trade and the buyer exercises his outside option. Is the unobservability of his investment choice behind this lack of information rents?

In other words, if the Designer could observe the investment choice, would the buyer's equilibrium payoff be larger? The intuition would be the same as in a Stackelberg leader-follower game: by moving first, the buyer is able to dictate the best-reply of the Designer that he prefers. Notice that the buyer cannot obtain a lower equilibrium payoff with observability than without it. After all, he could always choose the two point distribution of Theorem 3. Could he obtain a *strictly* larger payoff than that?

Condorelli and Szentes (2020) study a similar model, in which the buyer invests in a distribution over *valuations* for the good. The seller observes the buyer's distribution, and makes a take-it-or-leave-it offer p to the buyer. In equilibrium, the buyer accepts the offer if and only if it is (weakly) below his realized valuation. The authors argue that if the cost function of the buyer is strictly decreasing in mean preserving spread, the buyer's equilibrium distribution belongs to a certain family of distributions. Specifically, for $\pi \in [0,1]$, and $\rho \in [\pi, 1]$ define $F_{\pi,\rho} \in \mathcal{D}[0,1]$ by

$$F_{\pi,\rho}(v) = \begin{cases} 0 & v < 0, \\ 1 - \frac{\pi}{\rho} & v \in [0,\rho), \\ 1 - \frac{\pi}{v} & v \in [\rho, 1), \\ 1 & v \geq 1. \end{cases}$$

The seller responds to this choice by posting price ρ, as any price in the interval $[\rho, 1]$ leave her with an expected profit of π. I will now show that the Condorelli and Szentes (2020) result still holds when the buyer invests in outside options instead. To facilitate the comparison, I will assume that the Designer only cares about the seller ($\alpha = 1$) – or, equivalently, that the Designer *is* the seller. First, a preliminary lemma.

Lemma 2. *Suppose $G \in \mathcal{D}[0,1]$, let π denote the monopoly profit if the buyer chooses G, and p be the monopoly price. Then there exists a unique $\rho \in [\pi, 1]$ such that*

1. *$F_{\pi,\rho}$ is a mean preserving spread of G, and*
2. *The buyer's expected payoff is larger if he chooses $F_{\pi,\rho}$ and the seller posts price ρ than if he chooses G and the seller posts price p,*

$$1 - \rho + \int_0^{\rho} F_{\pi,\rho}(b)\mathrm{d}b - C(F_{\pi,\rho}) \geq 1 - p + \int_0^{p} G(b)\mathrm{d}b - C(G),$$

and the inequality is strict unless $G = F_{\pi_G,\rho_G}$ for some $\pi_G \in [0,1]$, $\rho_G \in [\pi_G, 1]$.

Proof. The geometric property 1 is taken *verbatim* from Condorelli and Szentes (2020), Lemma 4. As for 2, suppose G does not belong to the family, and notice that the value of the left-hand side of the inequality is

$$1 - \rho + \int_0^{\rho} 1 - \frac{\pi}{\rho}\mathrm{d}b - C(F_{\pi,\rho}) = 1 - \pi - C(F_{\pi,\rho}).$$

Since $F_{\pi,\rho}$ is a mean preserving spread of G, then, $C(F_{\pi,\rho}) \leq C(G)$. Therefore it is enough to argue that

$$\pi \leq p - \int_0^{p} G(b)\mathrm{d}b.$$

Since π is the monopoly profit, $\pi = p(1 - \lim_{b \uparrow p} G(b'))$. Since G is increasing,

$$p \lim_{b' \uparrow p} G(b') \geq \int_0^{p} \lim_{b' \uparrow p} G(b')\mathrm{d}b \geq \int_0^{p} G(b)\mathrm{d}b,$$

which concludes the proof. □

Corollary 1. *Under Assumption 3, in equilibrium the buyer chooses a distribution supported on $\{0,1\}$ with average μ^*, where*

$$\mu^* \in \operatorname{argmax}_{\mu} 1 - \mu - c(\mu).$$

Proof. Because of Lemma 2, it is without loss of generality to constrain the buyer's choice to the set

$$\mathcal{H} = \{F_{\pi,\rho} \in \mathcal{D}[0,1] : \pi \in [0,1], \rho \in [\pi,1]\}.$$

Pick any $\pi \in [0,1]$, and any $\rho \in [\pi,1]$. Then the buyer's payoff from choosing $C(F_{\pi,\rho})$ is lower than his payoff from choosing the distribution supported on $\{0,1\}$ with average π:

$$1 - \pi - C(F_{\pi,\rho}) \leq 1 - \pi - c(\pi).$$

The inequality is strict unless $\rho = 1$. In other words, the buyer's optimal distribution involves having $\rho = 1$, i.e. that he chooses a distribution supported on $\{0,1\}$, as desired. □

Thus, even if the seller could observe his investment decision, in equilibrium the buyer would still choose the maximally spread out distribution $(1-\mu^*)H_0 + \mu^* H_1$ as in Theorem 2, and the seller would respond by posting price 1. Therefore, the equilibrium outcome is the same regardless of the observability of investment, and so the hold-up inefficiency arises also in this case.

Finally, observe that the buyer does not extract any information rents from the seller also if investment is observable. The buyer's returns from investment come entirely from the value of the outside option, rather than from the ability to obtain information rents by dictating the seller's best reply.

1.5.2 Contractible Investment

Finally, I study what happens in the model if the investment choice of the buyer is both contractible, and observable. In particular, I assume that the Designer is able to offer to buyer a *complete, contingent contract* prior to his investment decision, which specifies a mechanism for each possible distribution chosen.

Consider the following timing of events. First, the Designer offers a contract to the buyer. A *contract* is a menu $\{(q_G, p_G)_{G \in \mathcal{D}[0,1]}\}$ that associates with each distribution $G \in \mathcal{F} \subseteq \mathcal{D}[0,1]$ a direct revelation mechanism $(q_G, p_G) : \operatorname{supp} G \to [0,1]^2$.

Having been offered the contract $\{(q_G, p_G)\}_{G\in\mathcal{D}[0,1]}$, the buyer decides whether to accept it or not. If he accepts, he next chooses a distribution over outside options — say, F— which the Designer observes. Finally, the buyer privately observes his outside option, $b \in \operatorname{supp} F$, and reports a type b' to the Designer, who then conducts trade according to the rules specified by the mechanism, $(q_F(b'), p_F(b'))$.

I will restrict attention to subgame perfect equilibria in truthful reporting. That is to say, I will assume that the contract only includes incentive compatible mechanisms, and the buyer reports his type truthfully, regardless of his investment choice. In equilibrium, the Designer recognizes that the buyer will respond optimally to the contract he's offered, that is, he will invest in the distribution, and hence choose the associated mechanism, which grants him the largest expected payoff. Say that contract $\{(q_G, p_G)_{G\in\mathcal{D}[0,1]}\}$ *induces* distribution F if the mechanism associated with F, (q_F, p_F), gives the buyer the largest expected payoff within the contract:

$$F \in \operatorname{argmax}_{G\in\mathcal{D}[0,1]} \int_0^1 \big(1 - b + q_G(b)(b - p_G(b)\big)\mathrm{d}G(b) - C(G).$$

Distribution F is *inducible* if there exists a contract inducing F. It is useful at this point to define the *simple contract at* F as the contract specifying a posted price mechanism at 0 if F is chosen, and a posted price mechanism at 1 otherwise. This contract rewards the buyer with the mechanism he prefers if F is chosen, and punishes him with a mechanism that extracts all his surplus whenever the buyer and the seller trade if F is not chosen.[9] As it turns out, the inducibility of F is tied to whether F is inducible by the simple contract at F.

[9]Since I will restrict the attention to ex-post individually rational mechanisms, the posted price mechanism at 1 is the harshest punishment available to the Designer to punish undesirable investment choices. In principle, the contract may specify also other punishments, for instance a large reimbursements from the buyer to the seller could be triggered if distribution F is not chosen. In that case, distribution F should still grant a payoff of at least $1 - \mu^* - c(\mu^*)$ for the buyer to enter the contract.

Lemma 3. *The following statements are equivalent:*

1. *F is inducible;*

2. *The simple contract at F induces F;*

3. *$C(F) \leq \mu^* + c(\mu^*)$, where $\mu^* \in \operatorname{argmax}_{\mu\in[0,1]} 1 - \mu - c(\mu)$.*

Proof. To show that (1) implies (2), suppose that there exists a contract inducing F, $\{(q_G, p_G)_{G\in\mathcal{D}[0,1]}\}$. Observe that the buyer's payoff from choosing F is larger under the simple contract than under $\{(q_G, p_G)_{G\in\mathcal{D}[0,1]}\}$ since

$$\int_0^1 (1 - b + q_F(b)(b - p_F(b))\, \mathrm{d}F(b) - C(F) \leq \int_0^1 (1-0)\mathrm{d}F(b) - C(F),$$

while the buyer's payoff is lower by choosing $G \neq F$ under the simple contract than under $\{(q_G, p_G)_{G\in\mathcal{D}[0,1]}\}$:

$$\int_0^1 (1 - b + q_G(b)(b - p_G(b))\, \mathrm{d}G(b) - C(G) \geq \int_0^1 (1-b)\mathrm{d}G(b) - C(G).$$

Consequently, if $\{(q_G, p_G)_{G\in\mathcal{D}[0,1]}\}$ induces F, then also the simple contract at F induces F.

To show that (2) implies (3), suppose that the simple contract at F induces F. Then the payoff from choosing F, $1 - C(F)$, is larger than the payoff from choosing any other distribution G, which is by definition no more than $1 - \mu^* - c(\mu^*)$.

To show that (3) implies (1), suppose that (1) is violated, so that F is not inducible. Therefore, for any contract $\{(\hat{q}_G, \hat{p}_G)_{G\in\mathcal{D}[0,1]}\}$ there must exist a distribution $\hat{G}$ such that

$$\int_0^1 1 - b + \hat{q}_{\hat{G}}(b)(b - \hat{p}_{\hat{G}}(b))\mathrm{d}\hat{G}(b) - C(\hat{G}) > \int_0^1 1 - b + q_F(b)(b - p_F(b))\mathrm{d}F(b) - C(F).$$

In particular, this condition is true for the simple contract at F, so that

$$\max_{G\in\mathcal{D}[0,1]} \int_0^1 1 - b + q_G(b)(b - p_G(b)\mathrm{d}G(b) - C(G) > 1 - C(F),$$

that is,

$$1-\mu^{*}-c(\mu^{*})>1-C(F),$$

so that (3) is violated. □

Having characterized the set of inducible distributions, the optimal contract can be found by answering two questions. First, which distribution, F^*, should the Designer induce? Second, which mechanism should be associated with the target distribution F^*? The optimal contract can be completed by punishing any other distribution with a posted price mechanism at 1, the harshest punishment available to the designer.[10]

In the model of Section 1.3, the Designer's objective function was a weighted average of the buyer's and seller's expected payoffs, where the buyer's payoff was net of the investment cost. The idea behind this modelling choice is that the investment cost is sunk when the buyer and the seller negotiate. However, when the parties sign a contract prior to the investment decision, the cost of it is not sunk yet, and so it should be included in the Designer's objective function.[11] The Designer's problem is

$$\max_{F,(q_F,p_F)} \quad \alpha\int_0^1 q_F(b)p_F(b)\mathrm{d}F(b)+(1-\alpha)\Bigg(\int_0^1 1-b+q_F(b)(p_F(b)-b)\mathrm{d}F(b)-C(F)\Bigg)$$

$$\text{s.t.} \int_0^1 \big(1-b+q_F(b)(b-p_F(b))\big)\mathrm{d}F(b)-C(F)\geq \max_{\mu\in[0,1]} 1-\mu-c(\mu),$$

$$\forall b,b'\in\operatorname{supp}F,\quad 1-b+q_F(b)(b-p_F(b))\geq 1-b+q(b')(b-p(b')),$$

$$\forall b\in\operatorname{supp}F,\quad q(b)>0 \implies 0\leq p(b)\leq b.$$

[10]The resulting contract is not necessarily the unique optimum for the designer, since there could be other contracts inducing F^* by associating the mechanism found above with F^*, and punishing the choice of $G\neq F^*$ with a different mechanism than the posted price mechanism at 1. For simplicity, in what follows I will refer to the contract above as *the* optimal contract.

[11]Still, the inclusion of $C(F)$ in the objective function of the Designer does not impact the results.

Theorem 4. *Suppose that investment is both observable and contractible.*

- *If the Designer favors the buyer ($\alpha < \frac{1}{2}$), then she induces distribution H_1 by associating with H_1 a price of 0.*
- *If the Designer favors the seller ($\alpha > \frac{1}{2}$), then she induces distribution H_1 by associating it with price $\mu^{**} + c(\mu^{**})$, the largest solution to the following maximization problem:*

$$\max_{\tilde{\mu}} \tilde{\mu} + c(\tilde{\mu}),$$
$$\textit{s.t. } \tilde{\mu} \in \operatorname{argmax}_{\mu \in [0,1]} 1 - \mu - c(\mu).$$

Off the equilibrium path, the optimal contract specifies a posted price mechanism at 1.

Proof. First, notice that by construction both mechanisms induce distribution H_1. Since trade is socially efficient in this model, posting any price that ensures trade with probability 1 maximizes the (unweighted) sum of the buyer and seller's payoffs. Let V_B^0 and U_S^0 be the buyer's and seller's payoff under the posted price 0. Let $\tilde{V}_B$ and $\tilde{U}_S$ be the buyer's and seller's payoff under an alternative ex-ante individually rational contract (possibly not inducing distribution H_1). Since trade is socially efficient,

$$U_S^0 + V_B^0 \geq \tilde{U}_S + \tilde{V}_B.$$

Suppose first that $\alpha < \frac{1}{2}$. Since $U_S^0 - \tilde{U}_S \geq V_B^0 - \tilde{V}_B$, it follows that $\alpha(U_S^0 - \tilde{U}_S) \geq \alpha(V_B^0 - \tilde{V}_B)$, and hence $(1-\alpha)(U_S^0 - \tilde{U}_S) > \alpha(V_B^0 - \tilde{V}_B)$. This proves that posting a price of 0 and inducing distribution H_1 is optimal if the Designer favors the seller.

Now suppose that $\alpha > \frac{1}{2}$. To simplify the exposition, suppose first that the cost function $C : \mathcal{D}[0,1] \to \mathbb{R}$ is strictly convex, so that also the function $c : [0,1] \to \mathbb{R}$ is strictly convex. Let μ^* be the only solution to the maximization problem $1 - \mu^* - c(\mu^*)$. I claim that any contract associating price $\mu^* + c(\mu^*)$ to H_1, and price 1 to any other distribution, is optimal. Let V_B^*, U_S^* be the buyer's and seller's payoff in this contract. Analogously to the previous paragraph, $U_S^* - \tilde{U}_S \geq \tilde{V}_B - V_B^*$, so that $\alpha\left(U_S^* - \tilde{U}_S\right) \geq \alpha\left(\tilde{V}_B - V_B^*\right) > (1-\alpha)\left(\tilde{V}_B - V_B^*\right)$, as desired.

To conclude the proof, suppose that the cost function C is convex, but not strictly convex, so that the set of maximizers of $1 - \mu - c(\mu)$ may contain more than one element. Notice that the argument in the previous pargraph shows that the optimal contract associates with H_1 a price that makes the buyer indifferent between choosing the targeted distribution H_1, and any different distribution. To find which price should the contract specify, define $\mu^{**} + c(\mu^{**})$ as follows:

$$\mu^{**} + c(\mu^{**}) = \max_{\tilde{\mu}} \left\{ \tilde{\mu} + c(\tilde{\mu}) : \tilde{\mu} \in \text{argmax}_{\mu \in [0,1]} \left(1 - \mu - c(\mu)\right) \right\}.$$

In other words, among the prices that make the buyer indifferent between complying with the contract and choosing the targeted distribution H_1, and not complying and be punished with a posted price at 1, $\mu^{**} + c(\mu^{**})$ is the largest one, so it maximizes the seller's payoff, and, therefore, the designer's objective function.

Let V_B^{**}, U_S^{**}, be the payoffs of the buyer and the seller under posted price $\mu^{**} + c(\mu^{**})$. Analogously to the previous paragraph, $U_S^{**} - \tilde{U}_S \geq \tilde{V}_B - V_B^{**}$, so that $\alpha \left(U_S^{**} - \tilde{U}_S \right) \geq \alpha \left(\tilde{V}_B - V_B^{**} \right) > (1 - \alpha) \left(\tilde{V}_B - V_B^{**} \right)$, as desired. □

With the optimal contract at hand, one can calculate the size of the inefficiency due to non-contractibility of the buyer's investment choice. To simplify the exposition, assume that the cost function C, and hence c, is strictly convex, so that the optimal contract associates with the targeted distribution H_1 the price $\mu^* + c(\mu^*)$, with $\mu^* = \text{argmax}_{\mu \in [0,1]} 1 - \mu - c(\mu)$. There is no inefficiency if the designer favors the buyer, as the buyer internalizes the social cost of investing in outside options even if the investment is not contractible. On the other hand, if the designer favors the seller, the social welfare under the optimal contract is

$$\alpha \left(\mu^* + c(\mu^*) \right) + (1 - \alpha) \left(1 - \mu^* - c(\mu^*) \right)$$

while the social welfare in equilibrium under non-contractibility (Section 1.3) is

$$\alpha \mu^* + (1 - \alpha)(1 - \mu^* - c(\mu^*)),$$

so that the size of the inefficiency, the difference in social welfare in the two scenarios, equals $\alpha c(\mu^*)$.

While the analysis of this Section has been conducted assuming that the buyer could generate any distribution in $\mathcal{D}[0,1]$, the result of Theorem 4 still holds if the available set is $\mathcal{F} \subseteq \mathcal{D}[0,1]$, once the least cost of generating mean μ, $c(\mu)$, is appropriately redefined as the cost of the least expensive distribution in $\mathcal{F}$ with average μ.

1.6 Discussion

1.6.1 Weakening Assumptions 1-3

How robust are the conclusions of Section 1.4 to weakening Assumptions 1-3? First, I show that nothing substantially changes if any of the Assumptions is weakened and the Designer cares more about the buyer than the seller. The proof of Theorem 2 makes it clear that neither the properties of C nor the domain of available distribution $\mathcal{F}$ matter for proving that the equilibrium price is 0. As for the choice of distribution, any distribution minimizing the cost would, together with the 0 price, form an equilibrium. If Assumption 2 holds, or, more generally, if H_1 is the only distribution with 0 cost, then there would be a unique equilibrium, $(H_1, 0)$. Either way, all equilibria are payoff equivalent: the buyer's payoff is 1, and the designer's payoff is $(1-\alpha)$.

For the rest of this Section, assume that the designer favors the buyer ($\alpha > \frac{1}{2}$). Suppose now that Assumption 1 is violated, so that the buyer has available only a (compact and convex) subset $\mathcal{F}$ of $\mathcal{D}[0,1]$, but continue assuming that C is strictly decreasing in first order stochastic dominance and mean preserving spread (Assumptions 2 and 3). The mechanics behind Lemma 1 wouldn't work without the full domain assumption, as the buyer may be unable to concentrate the probability of trade in a single type, 1. Still, if Assumption 3 holds, the best reply of the buyer to any price $p \in [0,1]$ would be maximal in the $\succsim_{MPS}$ order. Indeed, remember that the buyer's best reply to p solves the program

$$\max_{F\in\mathcal{F}} 1-p+\int_0^p F(b)\mathrm{d}b - C(F).$$

If F is a mean preserving spread of G, and $F, G \in \mathcal{F}$, then $\int_0^p \big(F(b)-G(b)\big)\mathrm{d}b \geq 0$. Therefore,

$$1-p+\int_0^p F(b)\mathrm{d}b - C(F) > 1-p+\int_0^p G(b)\mathrm{d}b - C(G),$$

so that G cannot be a best reply to p.

Lemma 4. *Suppose that the set of available distributions is $\mathcal{F}$ and that the cost function C satisfies Assumption 3. If (p^*, F^*) is an equilibrium, then the only mean preserving spread of F^* is F^* itself.*

Corollary 2. *Suppose that Assumption 3 holds. If the distribution concentrated at 0 is available to the buyer, $H_0 \in \mathcal{F}$, and the designer favors the seller, then all the equilibria are of the form $(0,(1-\mu^*)H_0+\mu^* H_1)$, where $\mu^* \in \operatorname{argmax}_{\mu\in[0,1]} 1-\mu-c(\mu)$.*

After this discussion, it should also be clear that if Assumption 2 is violated, but Assumptions 1 and 3 are maintained, the equilibria would still have the form $((1-\mu^*)H_0+\mu^* H_1, 1))$, where $\mu^* \in \operatorname{argmax}_{\mu\in[0,1]} 1-\mu-c(\mu)$. Indeed, by Assumption 3 the buyer's best reply to any price p belongs to the set of maximally spread out distributions

$$\{F \in \mathcal{D}[0,1] : \exists \mu \in [0,1], F = (1-\mu)H_0 + \mu H_1\}.$$

Hence, in any equilibrium the designer posts price 1, so that the buyer still has to choose the average μ^* which maximizes $1-\mu-c(\mu)$. However, assuming that the cost function be strictly decreasing in first order stochastic dominance is far from useless, as it implies (by Lemma 1) that in any equilibrium the price be 1. This is particularly relevant when Assumption 3 is weakened.

Indeed, suppose that Assumptions 1 and 2 are satisfied, but the cost function C is only *weakly* decreasing in mean preserving spread: if $F \succsim_{MPS} G$ then $C(F) \leq C(G)$ (possibly with equality). Because of Lemma 1, in any equilibrium the price would be 1. Consequently, the buyer's best reply can still be found with the two-step procedure described in Section 1.4.2: finding an optimal average outside option, μ^*, then choosing the distribution with least cost among those with cost $c(\mu^*)$.

Clearly, the two-point distribution $(1-\mu^*)H_0 + \mu^* H_1$ would be minimize the investment cost among the CDFs with average μ^*, so that $(1, (1-\mu^*)H_0 + \mu^* H_1)$ would still be an equilibrium. Suppose there is another distribution $(1, F)$. For this to happen, it must be that (1) $\mathbb{E}_F b = \mu^*$, (2) $C(F) = c(\mu^*) = C((1-\mu^*)H_0 + \mu^* H_1)$, and (3) posting price 1 is a best reply to F. Since F is a mean preserving contraction of $(1-\mu^*)H_0+\mu^* H_1$, the probability of trade in the equilibrium $(1, F)$ would be lower than in the equilibrium $(1, (1-\mu^*)H_0 + \mu^* H_1)$. Indeed trade happens if and only if the outside option of the buyer is 1, but since F is a mean preserving contraction of $(1-\mu^*)H_0 + \mu^* H_1$, the probability that buyer's type under F would necessarily be lower than that under the two-point distribution. This discussion is summarized in the following Theorem.

Theorem 5. *Suppose that Assumptions 1 and 2 are satisfied, and that the cost function C is weakly decreasing in mean preserving spread. Then in any equilibrium $(\hat{p}, \hat{F})$, the price is 1 ($\hat{p} = 1$) and the distribution of outside options $\hat{F}$ satisfies $\mathbb{E}_{\hat{F}} b = \mu^*$ and $C(\hat{F}) = c(\mu^*)$, where $\mu^* \in \operatorname{argmax}_\mu 1 - \mu - c(\mu)$. Moreover, $(1, (1-\mu_h^*)H_0 + \mu_h^* H_1)$ is another equilibrium, which Pareto dominates $(1, \hat{F})$.*

As an application of this result, consider the mean based cost function $\mathcal{L}$ of Example 3. Since $\mathcal{L}$ is strictly decreasing in FSD and weakly decreasing in MPS, Theorem 5 applies, and the least cost of generating average μ is $\ell(\mu)$. The set of equilibria contains all pairs $(1, F)$ such that (1) $\mathbb{E}_F b = \mu^* \in \operatorname{argmax}_{\mu \in [0,1]} 1 - \mu - \ell(\mu)$ and (2) 1 is a best reply of the designer to F.

However, it is hard to further weaken Assumption 3 without risking the existence of an equilibrium. Indeed, if Assumptions 1 and 2 are satisfied, so that the buyer's choice in any equilibrium would have to be a CDF F with average μ^* and cost $C(F) = c(\mu^*)$. If, for example, C is *strictly increasing* in mean preserving spread, then such distribution would be the one concentrated at μ^*, H^*_μ. But the designer's best reply to H^*_μ would be μ^*, not 1! This implies that no (pure) equilibrium exists in this context.

Theorem 6. *Suppose that Assumptions 1 and 2 hold, and that C is strictly increasing in Mean Preserving Spread. Then no pure equilibrium exists.*

1.6.2 Suboptimal Mechanisms

The main model of Section 1.3 assumes that the bargaining mechanism governing trade is ex-ante Pareto efficient. Thus, the equilibrium of the game can be found by studying the equilibrium of a fictitious game between the buyer and a *mechanism designer*, in which the designer's objective function is a weighted average of the buyer's and seller's expected payoff.

In this Section, I study how restrictive is ex-ante Pareto optimality in the model. To simplify the exposition, suppose that Assumptions 1-3 hold, and that the cost function of the buyer, $C : \mathcal{D}[0,1] \to \mathbb{R}$ is strictly convex. The first step is to adapt the definition of equilibrium, Definition 1. since the equilibrium mechanism needs not be optimal. However, the mechanism should still be defined only for outside options that are possible according to the buyer's chosen distribution, and clearly the distribution should still be optimal given the mechanism governing trade. In other words, say that a pair $(F^*, (q^*, p^*))$ is a *S-equilibrium* if

- $(q^*, p^*) : \operatorname{supp} F^* \to [0,1]^2$ is incentive compatible and ex-post individually rational, and

- F^* is a best reply for the buyer to (q^*, p^*).

Clearly, any equilibrium according to Definition 1 is a S-equilibrium, but the converse is false. To characterize the S-equilibria of the game, suppose that the buyer expects the mechanism governing trade to be $(q,p) : X \to [0,1]^2$. Then his best reply solves

$$\max_{G \in \mathcal{D}[0,1]} \int_0^1 \max\left\{\max_{x \in X}\{1-b+q(x)\,(b-p(x))\}, 1-b\right\} \mathrm{d}G(b) - C(G).$$

Notice that the integrand in the previous expression,

$$h(b) = \max\left\{\max_{x \in X}\{1-b+q(x)\,(b-p(x))\}, 1-b\right\},$$

is a convex function of $b \in [0,1]$, being a maximum of two convex functions, namely $1-b$ and $\max_{x \in X}\{1-b+q(x)(b-p(x))\}$. But then the integral

$$\int_0^1 \max\left\{\max_{x \in X}\{1-b+q(x)\,(b-p(x))\}, 1-b\right\} \mathrm{d}G(b)$$

reaches its maximum at a maximal element of the $\succsim_{MPS}$ order. Since the cost function C is strictly decreasing in MPS, the best reply of the buyer must be maximally spread out. This discussion is formalized in the following Lemma.

Lemma 5. *For any direct revelation mechanism $(q,p) : X \to [0,1]$, there exists $\nu^* \in [0,1]$ such that the best reply of the buyer to (q,p) is $(1-\nu^*)H_0 + \nu^* H_1$.*

In other words only distributions supported on $\{0,1\}$ can be a best reply to a mechanism, since they are maximal elements of the $\succsim_{MPS}$ order. It follows that in any S-equilibrium the distribution chosen by the buyer is supported on $\{0,1\}$.

Theorem 7. *A pair $(F^*, (q^*, p^*))$ is a S-equilibrium if and only if*

1. *The mechanism $(q^*, p^*) : \{0,1\} \to [0,1]^2$ satisfies*

$$q^*(1)p^*(1) \le q^*(1) - q^*(0);$$

2. There exists $\nu^* \in [0,1]$ *such that* $F^* = (1-\nu^*)H_0 + \nu^* H_1$ *with*

$$\nu^* = \text{argmax}_{\nu \in [0,1]}\, 1 - \nu + \nu q^*(1) p^*(1) - c(\nu).$$

Proof. By the previous discussion, we know that in any equilibrium the distribution chosen is supported on $\{0,1\}$, and hence the associated mechanism is defined on $\{0,1\}$. So let $((1-\nu^*)H_0 + \nu^* H_1, (q^*, p^*))$ be an equilibrium. Since (q^*, p^*) is ex-post individually rational, $p^*(0) = 0$. Thus type 1's incentive compatibility condition can be written as

$$q^*(0) \cdot (1-0) + (1-q^*(0))\,(1-1) \leq q^*(1)\,(1-p^*(1))) + (1-q^*(1))\,(1-1),$$

which implies condition 1. On the other hand, condition 2 just states that $(1 - \nu^*)H_0 + \nu^* H_1$ is a best reply to mechanism (q^*, p^*). Sufficiency is obvious. $\square$

Thus, also in any S-equilibrium the buyer chooses a distribution supported on only two points, 0 and 1, so he is unable to extract any information rent from the seller.

1.7 Conclusion

This paper studies the impact that investment in outside options can have on the outcomes of a negotiation. In the model, one of the two traders (the buyer) chooses a distribution generating a random outside option before negotiating with the other trader (the seller). The main result shows that, despite privately observing the realized value of the outside option, and even if the seller cannot observe the distribution chosen, in equilibrium the buyer is unable to extract information rents. Indeed, if the buyer has more bargaining power, then in equilibrium he does not invest in generating outside options, and the price equals the opportunity cost of the seller. Alternatively, if the seller has more bargaining power, then the price equals the willingness to pay of the buyer, and the buyer either sees all his surplus extracted, if the parties trade, or else he exercises his outside option. This result is robust to a substantial weakening of the assumptions used to derive it.

The results of the paper point to some intriguing questions. First, in the model only one party (the buyer, specifically) acquires through effort an outside option, while the other party can only trade with him. It is often more realistic to imagine that both parties in a relationship look around for alternative opportunities to improve on their bargaining position. Second, the benchmark given by the optimal contract of Section 1.5.2 is computed assuming that the effort of the buyer be both *contractible* and *observable* to the designer, who can then implement any punishment deemed appropriate in case of a deviation by the buyer. An alternative possibility regards a moral hazard framework, in which the effort of the buyer is indeed contractible, but not observable: the designer offers the contract prior to the investment choice, without the possibility of verifying it. In this case, the contract would only specify one mechanism, and the designer would decide which mechanism in order to induce a particular distribution chosen by the buyer. The problem of selecting a mechanism to influence the decision of an agent that generates his private information is also studied by Mensch (2020). The main difference with my framework is that in his model there is an underlying prior distribution of a *state of the world*, and the agent acquires a signal to learn about it.

Chapter 2

Incomplete Preferences or Incomplete Information? On Rationalizability in Games with Private Values

2.1 Introduction

The appeal of the completeness axiom of individual preferences has long been questioned, both on normative and positive grounds (von Neumann and Morgenstern, 1953; Aumann, 1962). In its standard interpretation, the axiom says that the agent is able to rank any pair of possible choices, and so its violation denotes an occasional indecisiveness by the decision maker. When incomplete orderings are used to model players' preferences in a game, however, a different interpretation is possible: that player i's preference relation is complete, but only part of that preference, captured by an incomplete order, is common knowledge in the game. In other words, i's true preference is a completion of a commonly known order, but exactly which completion is i's private information. Motivated by this double interpretation, I analyse a solution concept akin to rationalizability in a game with incomplete preferences, called Incomplete Preference Rationalizability (IPR). The main results of the paper establish that, in a wide class of games, IPR is well-defined, and its predictions coincide with those of apparently unrelated solution concepts for games with incomplete information: Belief-free Rationalizability (Battigalli et al., 2011; Bergemann and Morris, 2017), and Interim Correlated Rationalizability (Dekel et al., 2007).

In complete information games, rationalizability is the behavioral implication of rationality and common belief in rationality (Brandenburger and Dekel, 1987; Tan and Werlang, 1988). Rationality requires a player to play the best action given his belief. If players are expected utility maximizers, a rational player chooses her actions by maximizing her expected payoff, with respect to her belief. This means that a rational action is a maximum element in the ranking (over own actions) implied by the player's belief. When preferences are incomplete, however, the notion of "best", or maximum, can be too stringent to be imposed, due to occasional indecisiveness between actions. To address this, I define rational behavior as choosing an action which is maximal, given the player's belief. This is equivalent to requiring the agent choosing a maximum element of some completion of her ranking over actions. The solution concept I propose, called Incomplete Preference Rationalizability (IPR), is the largest set of actions with the *best-reply property*, that is, it only contains actions that are justified by conjectures supported on IPR. Theorem 8, shows that in a large class of games, IPR is non-empty and compact, and coincides with the set of actions surviving the iterated deletion of strictly dominated strategies.

The second set of results starts from the observation that incomplete orderings in a strategic environment can also model lack of common knowledge about players' tastes.[1] This is only possible in a *private values model*, in which i's private information does not directly affect j's payoff. Usually, to analyze a game with private values, a type space *à la* Harsanyi is used. There, player i's payoff type represents one of the possible completions of the part of her preference that is common knowledge. If action a_i is dominated by action b_i under all logically possible completions of i's preference, then it cannot be a maximal element of i's prefer-

[1]A similar observation is made by Carroll (2010).

ence. This implies that, if i is rational, she does not choose a_i. Following on this connection, I characterize Incomplete Preference Rationalizability in this context as the result of iterated deletion of actions that are strictly dominated under all of the *logically possible* completions of players' preferences. Moreover, if the commonly known part of each player's preference satisfies the expected multi-utility axioms as in Evren (2014), it is enough to look at all the expected utility completions. This facilitates the comparison between IPR and other kinds of Rationalizability developed for Bayesian Games. Theorem 9 shows that the set of incomplete preference rationalizable actions of player i is the same as the that of belief-free rationalizable actions for some expected utility completion of i's preference. The last result of the paper, Theorem 10, shows that an action is belief-free rationalizable if and only if it is interim correlated rationalizable in the universal type space. Consequently, the set of incomplete preference rationalizable actions coincides with that of interim correlated rationalizable actions.

The paper is organized as follows. Section 2.2 reviews the literature. Section 2.3 introduces the model and the main notation. Section 2.4 defines IPR and characterizes it in terms of iterated deletion of strictly dominated actions for a large class of games. Section 2.5 focuses on the incomplete information interpretation, and compares the solution concept with Belief-free Rationalizability and with Interim Correlated Rationalizability. Section 2.6 concludes the paper. All proofs are in the Appendix.

2.2 Literature Review

This paper contributes to two strands of the economic theory literature: games with incomplete preferences, and Rationalizability in games with incomplete information.

Incomplete preferences have been studied in game theory at least since the 1950's. For instance, Shapley and Rigby (1959) studies the equilibria of games in which the payoff function is vector-valued, instead of scalar valued. More recently, Bade (2005) studies equilibrium behavior in games with incomplete preferences. Her main result is that each Nash Equilibrium of such a game is an equilibrium of a game obtained by completing, in some way, the preferences of each player. Park (2015) studies equilibria in potential games with incomplete preferences. Chen et al. (2016) study Rationalizability under several possible models of individual behavior, including Bewley (2002)'s multiple priors model. Ziegler and Zuazo-Garin (2020) study Rationalizability using preferences *à la* Bewley (2002), showing how these ambiguity-sensitive preferences can be related to the notion of strategic cautiousness. Kokkala et al. (2019) also study the Rationalizable set in games with incomplete preferences, dispensing with the assumption that players hold probabilistic beliefs. In their model, player i's belief is a set of possible strategies of i's opponents, and a rational player chooses a non-dominated action given her belief.

The paper also contributes to the literature on Rationalizability. The solution concept was first proposed by Bernheim (1984) and Pearce (1984). Brandenburger and Dekel (1987) and Tan and Werlang (1988) prove that, in games with complete information, Rationalizability is the behavioral implication of rationality and common belief in rationality. In games with incomplete information, several related solution concepts have been proposed. For example, Δ-Rationalizability (Battigalli and Siniscalchi, 2003) studies the implications of common belief in rationality and in a restriction requiring that players' first order beliefs lie in a set Δ. In Bayesian games, where a type space *à la* Harsanyi is appended to the game form, Interim Correlated Rationalizability (Dekel et al., 2007) characterizes the implication of common belief in rationality and common knowledge of the type space. Finally, the notion of Belief-free Rationalizability was first proposed in Battigalli (2003), then further studied in Battigalli et al. (2011) and Bergemann and Morris (2017). Belief-free Rationalizability captures the implications of rationality and common belief in rationality, without imposing assumptions on the beliefs that players hold about each others' private information.

2.3 Model

2.3.1 Notation

A *preorder* $\succsim$ over a set X is a reflexive and transitive binary relation. Its irreflexive and reflexive parts are denoted $\succ$ and $\sim$, respectively. It is complete if it ranks each pair of elements: for all $x, y \in X$, either $x \succsim y$ or $y \succsim x$ holds. The upper contour set of $\succsim$ at x is $U_{\succsim}(x) = \{y \in X : y \succsim x\}$, the lower contour set is $L_{\succsim}(x) = \{y \in X : x \succsim y\}$. An *extension* of a preorder $\succsim$ is a binary relation $\succsim'$ such that $\succsim' \supseteq \succsim$ and $\succ' \supseteq \succ$. A *completion* of $\succsim$ is a complete extension. The set of all completions of $\succsim$, $\text{Compl}(\succsim)$, is non-empty by Szpilrajn (1930)'s Theorem. An element $x \in X$ is $\succsim$-*maximal* if there is no $y \in X$ for which $y \succ x$; it is a $\succsim$-*maximum* if for each $z \in X$, $x \succsim z$.

If X is a topological space, then $\mathbf{C}(X)$, $\mathcal{B}(X)$ and $\Delta(X)$ denote the set of continuous real functions on X, its Borel σ-algebra and the set of all Borel probability measures on X, respectively. In this paper, $\mathbf{C}(X)$ is always endowed with the sup-norm topology, while $\Delta(X)$ with the topology of weak convergence. If $Y \subseteq X$, then $\Delta(Y)$ denotes the subset of $\Delta(X)$ containing all the probability measures supported on (a subset of) Y, with a slight abuse of notation. A preorder $\succsim$ on X is *closed* if its graph is closed in the product topology of $X \times X$, and *open* if the graph of $\succ$ is open. By Schmeidler (1971)'s Theorem, a preorder on a connected set can be both open and closed if and only if it is complete.

If $(Y_j)_{j\in J}$ is an indexed collection of sets, write $Y = \bigtimes_{j\in J} Y_j$ and $Y_{-i} = \bigtimes_{j\neq i} Y_j$. A subset C of $Y = \bigtimes_{j\in J} Y_j$ is *Cartesian* if there exists $(C_j)_{j\in J}$ with $C_j \subseteq Y_j$ for each j such that $C = \bigtimes_{j\in J} C_j$. The set of all Cartesian subsets of Y is denoted $\mathfrak{C}(Y)$. If $f : Y \to X$ is a function and $y_i \in Y_i$, then $f_{y_i} : Y_{-i} \to X$ denotes the section of f at y_i, so that $f_{y_i}(y_{-i}) = f(y_i, y_{-i})$. If $\mathcal{F} : X \rightrightarrows Y$ is a correspondence, then its graph, denoted $\mathrm{Gr}(\mathcal{F})$ is the set

$$\mathrm{Gr}(\mathcal{F}) = \{(x, y) \in X \times Y : x \in \mathcal{F}(y)\}.$$

The graph of a function $f : X \to Y$, $\mathrm{Gr}(f)$, is defined analogously. Finally, $X^{\mathbb{N}}$ denotes the set of all sequences on X.

2.3.2 Games with Incomplete Preferences

In this paper, I only analyze games in normal form. A *game form* is a tuple $\langle I, \mathcal{C}, g, (A_i)_{i\in I}\rangle$ where I is a (finite) set of players, $\mathcal{C}$ is a set of consequences, A_i is the set of actions available to player i, and $g : \bigtimes_{i\in I} A_i \to \mathcal{C}$ is the consequence function. A game form is called *compact continuous* if $\mathcal{C}$ and all of the action spaces A_i's are compact metric spaces, and g is a continuous mapping. In the rest of the paper, I only study compact continuous games. Observe that, trivially, finite games are compact continuous.

I assume that each player i is endowed with a preorder $\succsim^i$ over $\Delta(\mathcal{C})$, that captures her attitudes toward risk. A *game* is a tuple obtained by appending players' preferences to the game form: $\Gamma = \langle I, \mathcal{C}, g, (A_i, \succsim^i)_{i\in I}\rangle$. The game itself, i.e. Γ, is assumed to be common knowledge: all players know the game they are playing, they know that everybody knows the game, and so on *ad infinitum*. When i's preorder $\succsim^i$ is not complete, the model admits two possible interpretations. First, it is possible that i's preferences are genuinely incomplete, that is, there are two lotteries $p, q \in \Delta(\mathcal{C})$ that i is unable to rank. Alternatively, the preorder $\succsim^i$ captures the part of i's preferences that is common knowledge among the players. In this second interpretation, i's true preference is complete, and it is some completion $\widehat{\succsim}^i$ of $\succsim^i$. Exactly which completion of $\succsim^i$ is i's private information.

I assume that the beliefs of player i about her coplayers' behavior can be captured by a single probability measure $\mu^i \in \Delta(A_{-i})$. I assume that no additional incompleteness stems from the strategic interaction with i's opponents.

Player i ranks actions according to the lottery they induce, together with the belief. Formally, let $\mu^i_{d_i} \in \Delta(\mathcal{C})$ be the pushforward measure of g_{d_i} under μ^i. Then i prefers action a_i to action b_i, written $a_i \succsim^i_{\mu^i} b_i$, if $\mu^i_{a_i} \succsim^i \mu^i_{b_i}$. Notice that this is a generalization of the expected utility case, where an agent with payoff function $u_i : \mathcal{C} \to \mathbb{R}$ and belief $\mu^i \in \Delta(A_{-i})$ ranks action a_i higher than action b_i if $\mu^i_{a_i}$ gives larger expected utility than $\mu^i_{b_i}$.

With incomplete preferences, rationality is captured by the choice of a $\succsim^i_{\mu^i}$-maximal element. Formally, if i is rational and has belief $\mu^i \in \Delta(A_{-i})$, she chooses a_i so that, for no $b_i \in A_i$, $b_i \succ^i_{\mu^i} a_i$.

Definition 3. An action $a_i \in A_i$ is *justifiable* if there exists $\mu^i \in \Delta(A_{-i})$ such that a_i is $\succsim^i_{\mu^i}$-maximal. In this case, a_i is *justified* by μ^i. The correspondence $r_i(\,\cdot\,; \succsim^i) : \Delta(A_{-i}) \rightrightarrows A_i$ defined by

$$r_i(\mu^i; \succsim^i) = \{a_i \in A_i : \nexists b_i \in A_i, b_i \succ^i_{\mu^i} a_i\}$$

is called (with a slight abuse of terminology) the *best-reply correspondence* of player i. (When it is clear from the context, I will drop the dependence on the preorder $\succsim^i$ from the notation.)

The first result of the paper establishes how the properties of player i's preference relation translate into properties of the best-reply correspondence. These properties mimic those that the best-reply correspondence has for an Expected Utility maximizer.

Lemma 6. *If $\succsim^i$ is either open or closed, then r_i is non-empty valued. Furthermore, if $\succsim^i$ is open, r_i is also upper-hemicontinuous and compact valued.*

The notion of justifiability of Definition 3 is motivated by the incomplete preference interpretation of the model. Under the incomplete information interpretation, rational behavior means choosing an action that maximizes the true completion of the (commonly known) preference relation. Therefore, an action maximizing any completion of $\succsim^i$ is justifiable under this interpretation of the model. Formally, define the correspondence $\tilde{r}_i : \Delta(A_{-i}) \rightrightarrows A_i$ by

$$\tilde{r}_i(\mu^i; \succsim^i) = \bigcup \big\{ r_i(\mu^i; \widehat{\succsim}^i) : \widehat{\succsim}^i \in \mathrm{Compl}(\succsim^i) \big\}.$$

If $\tilde{r}_i$ and the best-reply correspondence r_i coincide, then rational behavior in this model is the same regardless of its interpretation. More generally, given a preorder $\succeq$ on X, if x is $\succeq$-maximal, then there exist a completion $\widehat{\succeq}$ of $\succeq$ that has x at its maximum – see (Bade, 2005, Lemma 1) for a proof. Consequently, the actions that μ^i justifies for the incomplete preference $\succsim^i$ are exactly those that μ^i justifies for some completion of $\succsim^i$.

Lemma 7. *For any belief $\mu^i \in \Delta(A_{-i})$, $r_i(\mu^i; \succsim^i) = \tilde{r}_i(\mu^i; \succsim^i)$.*

2.4 Incomplete Preference Rationalizability

I now introduce the main solution concept of the paper, Incomplete Preference Rationalizability (IPR). As with similar solution concepts, IPR can be defined in two ways: as the largest set with the *best-reply property*, or through the iterative elimination of strictly dominated actions. Let's start from the former. Define the *rationalization operator* to be the mapping $\rho : \mathfrak{C}(A) \to \mathfrak{C}(A)$ such that

$$C \mapsto \rho(C) = \bigtimes_{i \in I} r_i\big(\Delta(C_{-i})); \succsim^i \big).$$

Say that a Cartesian set $C \in \mathfrak{C}(A)$ has the *best-reply property* if for each player i, C_i contains only actions that are justified by beliefs supported on C_{-i}, that is, if $C \subseteq \rho(C)$.

Definition 4. Action profile $a \in A$ is Incomplete Preference Rationalizable if and only if it belongs to a set $C \in \mathfrak{C}(A)$ with the best-reply property.

An immediate consequence of Definition 4 is that the set of all Incomplete Preference Rationalizable action profiles, IPR, is the largest set with the best-reply property.

Alternatively, IPR can be defined through the inductive elimination of strictly dominated actions. Define inductively the n-th power of ρ to be the operator $\rho^n : \mathfrak{C}(A) \to \mathfrak{C}(A)$ such that $\rho^n(C) = \rho(\rho^{n-1}(C))$, where $\rho^0(C) = C$. The sequence $(\rho^n(A))_{n\in\mathbb{N}}$ performs the (maximal) iterated deletion of strictly dominated actions.

Definition 5. The set of iteratively undominated action profiles is the intersection of all elements of the sequence $(\rho^n(A))_{n\in\mathbb{N}}$:

$$\widehat{\mathrm{IPR}} = \bigcap_{n\in\mathbb{N}} \rho^n(A).$$

The two definitions need not coincide, even for compact continuous games, unless players' preferences also satisfy some topological requirements. When they differ, Definition 4 should be used, as it captures the implications of rationality and common belief in rationality. The main result of this Section establishes sufficient conditions for the two definitions to coincide, and studies the properties of the rationalizable set.

Theorem 8. *If either*

1. Γ *is a finite game, or*

2. Γ *is a compact continuous game, and the preference relation of each player* $i \in I$, $\succsim^i$, *is an open preorder,*

then Definitions 4 and 5 coincide, i.e. $\mathrm{IPR} = \widehat{\mathrm{IPR}}$. *Moreover, the rationalizable set is non-empty, compact, and is a maximal fixed point of the rationalization operator, so that* $\mathrm{IPR} = \rho(\mathrm{IPR})$.

There are other, rather indirect, ways to prove that IPR is not empty. For instance, when the game admits a Nash Equilibrium a^*, then the set $\{a^*\}$ has the best-reply property, so that IPR is non-empty.[2] Another possibility is to show that a game obtained by completing players' preferences in some way has a non-empty rationalizable set.

2.4.1 Comparative Statics

In this subsection, I perform a comparative statics analysis of the Incomplete Preference Rationalizable set. The main question I ask here is: if the preference of $\succsim^i$ becomes "less incomplete", what happens to the rationalizable set?

For instance, suppose that one of the players in the game is a committee whose preferences are formed using the unanimity rule, so that consequence A is preferred to consequence B if and only if each member of the committee prefers A to B. The fewer members the committee has, the more comparisons it is able to perform, since agreeing on the relative ranking between two options would be easier. In a sense, the preference of the committee with fewer members is less incomplete. An alternative motivation of this question comes from the interpretation of the model as an incomplete information model. Suppose that more information is revealed to all players about i's preference: what can we say about the resulting rationalizable set?

To state the question formally, say that a game $\widehat{\Gamma} = \langle I, \mathcal{C}, g, (A_i, \widehat{\succsim}^i)_{i \in I} \rangle$ is a *preference extension* of $\Gamma = \langle I, \mathcal{C}, g, (A_i, \succsim^i)_{i \in I} \rangle$ if for each player i, $\widehat{\succsim}^i$ extends $\succsim^i$. It

[2]This is of course true also with complete preferences – see (Osborne and Rubinstein, 1994, p.56).

is a *preference completion* of Γ if, for each i, $\widehat{\succsim}^i$ is a completion of $\succsim^i$, and an *Expected Utility completion* if each $\widehat{\succsim}^i$ admits an expected utility representation.

Proposition 1. *If $\widehat{\Gamma}$ is a preference extension of Γ, then the rationalizable set of $\widehat{\Gamma}$ is contained in that of Γ.*

Intuitively, the more comparisons the player is able to make, the more actions he is able to rank for any given belief. Consequently, the best-reply correspondence shrinks, that is, any belief justifies fewer actions.

Proposition 1 holds, in particular, whenever $\widehat{\Gamma}$ is a preference completion of Γ. Therefore, the rationalizable set of Γ is a superset of the union of the rationalizable sets over all preference completion games of Γ. However, an incomplete preference rationalizable action profile of Γ is not necessarily rationalizable in some preference completion of Γ. In other words, the inclusion above can be strict: a simple counterexample is provided in the Appendix.[3]

Proposition 1 can also be used to argue that IPR is non-empty in some games. In particular, if one could argue that the rationalizable set of some preference completion game of Γ is non-empty, than that of Γ would be non-empty as well. One example of this is if the preference of each player admits an expected utility completion. The final result of this section is an immediate application of Proposition 3 in Dubra et al. (2004).

Corollary 3. *Let Γ be a compact-continuous game, and suppose that the preference of each player i is a closed preorder that satisfies Independence. Then the rationalizable set of Γ is non-empty.*

[3]Notice the difference with Bade (2005)'s main result, which says that a is a Nash equilibrium of game Γ if and only if it is a Nash equilibrium for *some* preference completion of Γ.

2.5 Incomplete preferences or incomplete information?

In this Section, I will focus on the second interpretation of the model: that each player i has a *complete* preference over consequences, but only a part of it, described by the incomplete order, is common knowledge. For consistency, the true preference relation of player i, which is her private information, must be a completion of the commonly known incomplete order. Normally, the analysis of private information in games is done through the Bayesian model of Harsanyi (1967). There, each player i is endowed with a *type* capturing both her private information and her beliefs on her opponents' information.

This paper proposes an alternative method to model private information: incomplete orderings, which can be used to capture the part of each player's preference that is common knowledge in the game. For example, in an auction setting it is common knowledge that player i prefers allocation A to allocation B when both assign the good to player i and the price is lower in A than in B, while her preference between allocation A and one in which she doesn't obtain the good is i's private information. This procedure allows the modeler to focus on which aspects of the game are common knowledge, while avoiding assumptions on the interactive beliefs of the players about each other's information. That said, it is important to notice that type spaces *à la Harsanyi* can be used also to model situations where i's private information directly affects j's payoff, while the incomplete orderings approach cannot.

In most type spaces models, players not only hold probabilistic beliefs about each other's information and actions, but are also expected utility maximizers. To facilitate the comparison between rationalizability in Bayesian games and Incomplete Preference Rationalizability, I will use a *multi-utility* representation for preferences under risk due to Evren (2014). Say that $\succsim^i$ admits an Evren expected multi-utility representation if there exists a compact and convex set of Bernoulli indices $\mathcal{U}_i \subseteq \mathbf{C}(\mathcal{C})$ such that

$$p \succ^i q \iff \forall u_i \in \mathcal{U}_i, \quad \mathbb{E}_p u_i > \mathbb{E}_q u_i, \quad \text{and}$$
$$p \sim^i q \iff \forall u_i \in \mathcal{U}_i, \quad \mathbb{E}_p u_i = \mathbb{E}_q u_i.$$

Each Bernoulli index $u_i \in \mathcal{U}_i$ represents an expected utility completion of $\succsim^i$. Conversely, if $\widehat{\succsim}^i$ is an expected utility completion of $\succsim^i$, then there exists a Bernoulli index $u_i \in \mathcal{U}_i$ that represents $\widehat{\succsim}^i$. The interpretation is that which $u_i \in \mathcal{U}_i$ represents i's true preference is i's private information, while $\mathcal{U}_i$ is commonly known.

Throughout the rest of the section, assume that $\Gamma = \langle I, \mathcal{C}, g, (A_i, \succsim^i)_{i \in I} \rangle$ is a compact continuous game with incomplete preferences, and that each player's preference, $\succsim^i$, admits a compact and convex multi-utility representation $\mathcal{U}_i$ as above. The interpretation is that the correct Bernoulli utility index $u_i \in \mathcal{U}_i$ is i's private information, whereas the set of *possible* Bernoulli utilities, $\mathcal{U}_i$, is common knowledge. In other words, assume (1) that player i's multi-utility representation is $\mathcal{U}_i$, (2) that player i has a complete preference relation $\widehat{\succsim}^i$, which is a completion of $\succsim^i$, and (3) that $\widehat{\succsim}^i$ admits an expected utility representation. This interpretation of the model is valid as long as action a_i is justified by $\mu^i \in \Delta(A_{-i})$ for the incomplete order $\succsim^i$ if and only if it is justified by μ^i for *some* expected utility completion of $\succsim^i$. The following result shows that this is indeed the case.

Proposition 2. *Suppose that $\succsim^i$ has a compact and convex multi-utility representation $\mathcal{U}_i$, and let $r_i(\,\cdot\,; u_i)$ be i's best reply correspondence if her Bernoulli index is u_i. Then, for each $\mu^i \in \Delta(A_{-i})$,*

$$r_i(\mu^i; \succsim^i) = \bigcup \{ r_i(\mu^i; u_i) : u_i \in \mathcal{U}_i \}.$$

In the rest of this Section, I will compare IPR with two solution concepts developed for games with incomplete information: Belief-free Rationalizability, or BFR

(Battigalli et al., 2011; Bergemann and Morris, 2017), and Interim Correlated Rationalizability, or ICR (Dekel et al., 2007). Since a preference $\succsim^i$ that admits a Evren representation must be open, both Definitions 4 and 5 can be used to characterize the Interim Correlated Rationalizable set. I will use the iterated dominance one, since it makes the comparison with other solution concepts easier. The main results of this Section show that Incomplete Preference Rationalizability characterizes all the outcomes that can emerge under either BFR or ICR.

2.5.1 Belief-free Rationalizability

The first solution concept I compare IPR to is Belief-free Rationalizability (Battigalli et al., 2011; Bergemann and Morris, 2017). It studies the implication of rationality and common belief in rationality alone, taking for granted the common knowledge of the game form (which includes the possible payoff types of each player). In other words, it is free of assumptions regarding interactive exogenous beliefs, that are implicitly imposed by a Harsanyi type structure.

BFR is a correspondence mapping payoff types $u_i \in \mathcal{U}_i$ into the actions that are belief-free rationalizable. It is defined inductively through the following iterated deletion of action-and-type pairs.[4] (*Basis Step*) For each i and each u_i, set $\mathrm{BFR}_i^0(u_i) = A_i$. (*Inductive Step*) If $\mathrm{BFR}_j^{n-1}(u_j)$ has been defined for all players j and all payoff-types $u_j \in \mathcal{U}_j$, then

$$\mathrm{BFR}_i^n(u_i) = \left\{ a_i \in A_i : \exists \nu^i \in \Delta(A_{-i} \times \mathcal{U}_{-i}), \begin{array}{cc} & (i)\ a_i \in r_i(\nu^i; u_i), \\ (ii) & \nu^i(\mathrm{Gr}(\mathrm{BFR}_{-i}^{n-1})) = 1. \end{array} \right\}$$

Condition (i) requires that ν^i justifies action a_i, while condition (ii) requires that ν^i only gives probability to action-payoff types pairs that have not yet been eliminated.

[4]The formalizations of Battigalli et al. (2011) and of Bergemann and Morris (2017) are slightly different. I adopt the definition of the former article.

The set of belief-free rationalizable actions for type u_i of player i, then, is

$$\mathrm{BFR}_i(u_i) = \bigcap_{n \in \mathbb{N}} \mathrm{BFR}_i^n(u_i).$$

Theorem 9 establishes that the set of incomplete preference rationalizable actions coincides with the set of *all* belief-free rationalizable actions, that is, with the image of the BFR correspondence.

Theorem 9. *For each player $i \in I$, a_i is incomplete preference rationalizable if and only if there exists a payoff type $u_i \in \mathcal{U}_i$ such that a_i is belief-free rationalizable for type u_i.*

It is not particularly surprising that every belief-free rationalizable action is also incomplete preference rationalizable: $\mathrm{BFR}(\mathcal{U}) = \bigtimes_{i \in I} \mathrm{BFR}_i(\mathcal{U}_i)$ is a set with the best-reply property, and so it must be a subset of IPR. The opposite inclusion is less obvious, so consider a two-player game for intuition. If $a_i \in A_i$ is incomplete preference rationalizable, then there exist a payoff type $u_i \in \mathcal{U}_i$ and a belief $\mu^i \in \Delta(A_{-i})$, giving zero probability to actions that are not incomplete preference rationalizable, such that a_i maximizes the expected value, according to μ^i, of u_i. Since each a_{-i} in the support of μ^i maximizes the expectation of some payoff type u_{-i}, one can construct a probability measure on $A_{-i} \times \mathcal{U}_{-i}$ supported on (action, payoff type) pairs such that the action is justified by some belief for that payoff type. The technical details of this proof can be found in the Appendix.

2.5.2 Interim Correlated Rationalizability

The second solution concept IPR is compared to is Interim Correlated Rationalizability (Dekel et al., 2007). It studies the behavioral implications of common belief in rationality *and* in the type space. Among its desirable properties, the set of ICR actions only depends on the hierarchies of beliefs that a type space induces, but not on the specific signals that induce a given hierarchy. As a consequence, the set of interim correlated rationalizable actions of the universal type space of Mertens and Zamir (1985) is as large as it can be: the ICR set obtained by appending any type space to the game form will be a subset of that of the universal type space.

To make things formal, let $(\mathcal{T}_i, \kappa_i)_{i\in I}$ be the universal type space, the space of all belief hierarchies such that (1) each player i knows her own type, and (2) this fact is common knowledge, endowed with the canonical homeomorphism $\kappa_i : \mathcal{T}_i \to \Delta(\mathcal{U}_{-i} \times \mathcal{T}_{-i})$. Therefore, $\kappa_i(t_i)$ is the belief that player i would have on both her opponents payoff type, as well as their beliefs. Let $\mathcal{T}_i[u_i]$ be the set of all belief hierarchies in $\mathcal{T}_i$ in which i knows that her type is u_i.[5] I will show that the set of all interim correlated rationalizable actions for types in $\mathcal{T}_i[u_i]$ coincides with the set of all belief-free rationalizable actions for the payoff-type u_i. Combining this result with Theorem 9, it follows that, regardless of the type space appended to the game form Γ, any interim correlated rationalizable action is also incomplete preference rationalizable.

Interim Correlated Rationalizability is defined by the following iterative deletion procedure. (*Basis step*) For each i, each $u_i \in \mathcal{U}_i$ and each $t_i \in \mathcal{T}_i[u_i]$, set $\mathrm{ICR}_i^0(t_i) = A_i$. (*Inductive step*) If $\mathrm{ICR}_i^{n-1}(t_i)$ has been defined, then

$$\mathrm{ICR}_i^n(t_i) = \left\{ a_i \in A_i : \exists \nu^i \in \Delta(A_{-i} \times \mathcal{U}_{-i} \times \mathcal{T}_{-i}),\ \begin{array}{c} (1)\ a_i \in r_i(\nu^i; u_i), \\ (2)\ \mathrm{marg}_{\mathcal{U}_{-i}\times\mathcal{T}_{-i}}\, \nu^i = \kappa_i(t_i), \\ (3)\ \nu^i(\mathrm{Gr}\, \mathrm{ICR}_{-i}^{n-1}) = 1. \end{array} \right\}$$

In words, condition (1) requires the action to be justified by ν^i when i's payoff function is u_i; (2) imposes consistency between the conjecture ν^i and type t_i's beliefs; and (3) requires that no probability is assigned to action type pairs that have been already eliminated. The set of interim correlated rationalizable actions of player i

[5]The formal definition of these objects is standard and is available in the Appendix.

with type t_i is then

$$\mathrm{ICR}_i(t_i) = \bigcap_{n \in \mathbb{N}} \mathrm{ICR}_i^n(t_i).$$

Theorem 10 shows that the set of *all* belief-free rationalizable actions for a payoff type u_i coincides with that of all interim correlated rationalizable for types $t_i \in \mathcal{T}_i[u_i]$. Consequently, the set of incomplete preference rationalizable actions coincides with the image of the ICR correspondence.

Theorem 10. *For each player $i \in I$ and each payoff function $u_i \in \mathcal{U}_i$,*

$$\mathrm{BFR}_i(u_i) = \mathrm{ICR}_i(\mathcal{T}_i[u_i]).$$

Consequently,

$$\mathrm{IPR} = \mathrm{ICR}(\mathcal{T}).$$

The intuition behind the Theorem is straightforward: the universal type space contains all the belief hierarchies, and so it does not impose assumptions on interactive beliefs. The proof of this result follows the lines of that of Theorem 9, and can be found in the Appendix.

2.6 Concluding Remarks

1. Not all possible types of incomplete information can be modeled using an incomplete preorder. First, this approach can only be used in games with private values. Second, even in games with private values, the information which is common knowledge about player i's preference, may be finer than that expressible by an incomplete preorder. For instance, it can be common knowledge that either $a \succ^i b \succ^i c$ or $c \succ^i b \succ^i a$, while the only preorder that can capture this information is the empty set.

2. In a recent paper, Ziegler (2021) reaches a similar result to that of Theorem 10, showing that, in finite games, the set of Δ-rationalizable actions (Battigalli

and Siniscalchi, 2003) coincides with the set of interim correlated rationalizable actions of type spaces consistent with the restrictions contained in Δ; the first sentence of Theorem 10 follows as a corollary.

3. While it is outside of the scope of this paper, the first sentence of Theorem 10 and its proof can be adapted to a larger class of incomplete information games, including those where players' private information affects their opponents payoffs.

Appendix A

Appendix to Chapter 1

In this Appendix, I prove Theorem 1, that states that for each $F \in \mathcal{D}[0,1]$ and each $\alpha \in [0,1]$, there is a posted price mechanism which is a best reply to F. The proof of the Theorem relies on the well-known convex analysis fact that a continuous and linear functional W on a compact and convex set $\mathcal{M}$ attains its maximum value at an extreme point of $\mathcal{M}$. However, some care must be put before applying this fact to the Designer's problem since her objective function is not linear in the choice variable, the mechanism. So the main proof consists in four steps. First, redefine the space of mechanisms that the Designer chooses from, so to make her objective function linear. Second, define an appropriate Banach space containing all the mechanisms the Designer could choose, and such that the Designer's objective is also continuous in the norm topology. Third, checking that the set of (direct revelation) mechanisms that are incentive compatible and ex-post individually rational is a compact and convex subset of the Banach space found above. And finally, show that the extreme points of the set of feasible mechanisms coincide with the posted price mechanisms. Throughout the proof, fix $F \in \mathcal{D}[0,1]$, the CDF the Designer is responding to.[1] I will refer to elements of $\operatorname{supp} F$ as *types*.

Step 1. Let (q,p) be a direct revelation mechanism on $\operatorname{supp} F$, and define the function $t : \operatorname{supp} F \to [0,1]$ by $t(b) = q(b)p(b)$. From now on, a *mechanism* is a

[1] This outlined argument is the standard one used when F admits a density– see Borgers (2015), Chapter 2. The innovation is to extend this argument to *fully general* cumulative distribution functions, not just discrete or continuous ones.

pair (q,t). For technical reasons, it is useful to extend the mechanism (q,t) to the whole interval $[0,1]$, by declaring it 0 outside of $\operatorname{supp} F$. So say that a mechanism $(q,t):[0,1]\to[0,1]^2$ is F-feasible if it is Incentive Compatible,

$$\forall b,b'\in\operatorname{supp} F,\ \ 1-b+q(b)b-t(b)\geq 1-b+q(b')b-t(b'),$$

ex-post Individually Rational,

$$\forall b\in\operatorname{supp} F,\ \ t(b)\in[0,q(b)b],$$

and identically 0 on $[0,1]\setminus\operatorname{supp} F$,

$$\forall x\notin F,\ \ q(b)=t(b)=0.$$

Let $\mathcal{M}(F)$ be the set of F-feasible mechanisms. The Designer's problem can then be written as

$$\max_{(q,t)\in\mathcal{M}(F)}\alpha\int_0^1 t(b)\mathrm{d}F(b)+(1-\alpha)\int_0^1\big(1-b+t(b)-q(b)b\big)\mathrm{d}F(b).$$

Step 2. Let $(\mathcal{B}[0,1]^2,\|\cdot\|_1)$ be the (Banach) space of all pairs of bounded real functions on $[0,1]$ endowed with the norm

$$\|(f,g)\|_1=\int_0^1|f|\mathrm{d}\lambda+\int_0^1|g|\mathrm{d}\lambda,$$

where λ is the Lebesgue measure. Therefore, the functional $W_\alpha:\mathcal{B}[0,1]^2\to\mathbb{R}$ defined by

$$(f,g)\mapsto W_\alpha(f,g)=\alpha\int_0^1 g(b)\mathrm{d}F(b)+(1-\alpha)\int_0^1\big(1-b+g(b)-f(b)b\big)\mathrm{d}F(b)$$

is linear and continuous in $\|\cdot\|_1$.

Step 3. I claim that $\mathcal{M}(F)$ is compact and convex. Convexity is easy: if $(q,t),(\tilde{q},\tilde{t}) \in \mathcal{M}(F)$, and $\xi \in [0,1]$, then for all $b, b' \in \operatorname{supp} F$,

$$\xi\big(q(b)b - t(b)\big) + (1-\xi)\big(\tilde{q}(b)b - \tilde{t}(b)\big) \geq \xi\big(q(b')b - t(b')\big) + (1-\xi)\big(\tilde{q}(b')b - \tilde{t}(b')\big),$$

that is, $\xi\big(q,t\big) + (1-\xi)\big(\tilde{q},\tilde{t}\big)$ is Incentive Compatible,

$$0 \leq \xi t(b) + (1-\xi)\tilde{t}(b) \leq \xi q(b)b + (1-\xi)\tilde{q}(b)b,$$

that is, it is ex-post Individually Rational, and is clearly 0 outside of $\operatorname{supp} F$. Thus, $\xi(q,t\big) + (1-\xi)\big(\tilde{q},\tilde{t}\big) \in \mathcal{M}(F)$, as desired.

In order to show that $\mathcal{M}(F)$ is compact, I will use Helly's selection theorem, which states that a uniformly bounded sequence of monotone functions admits a pointwise convergent subsequence. To apply Helly's theorem, I will first show that any F-feasible mechanism is non-decreasing on the support of F. Then, I will consider a modified version of $\mathcal{M}(F)$, which will be denoted by $\widehat{\mathcal{M}}(F)$, such that each memeber of this set is non-decreasing on the whole interval $[0,1]$, and coincides with some F-feasible mechanism on $\operatorname{supp} F$. By Helly selection theorem, $\widehat{\mathcal{M}}(F)$ is compact, and, by constructing a continuous mapping between $\widehat{\mathcal{M}}(F)$ and $\mathcal{M}(F)$ I show that also the latter is compact.

So pick an F-feasible mechanism (q,t). I claim that both q and t are non-decreasing on $\operatorname{supp} F$. Take $b, b' \in \operatorname{supp} F$, with $b \geq b'$. By Incentive Compatibility,

$$q(b)b - t(b) \geq q(b')b - t(b'), \quad \text{and} \quad q(b')b' - t(b') \geq q(b)b' - t(b),$$

so adding the two conditions,

$$(b-b')(q(b) - q(b')) \geq 0,$$

so that q is non-decreasing. Rearranging the Incentive Compatibility condition of b', I obtain

$$t(b) - t(b') \geq b'\big(q(b) - q(b')\big) \geq 0,$$

so that also t is non-decreasing on $\operatorname{supp} F$.

Define the set $\widehat{\mathcal{M}}(F)$ by

$$\widehat{\mathcal{M}}(F) = \left\{ (\hat{q}, \hat{t}) : \exists (q,t) \in \mathcal{M}(F), (\hat{q}, \hat{t})(b) = \begin{cases} (q,t)(b) & b \in \operatorname{supp} F, \\ \sup_{\operatorname{supp} F \cap [0,b]} (q,t)(b') & b \notin \operatorname{supp} F. \end{cases} \right\}$$

In other words, if $(\hat{q}, \hat{t}) \in \widehat{\mathcal{M}}(F)$, then $(\hat{q}, \hat{t})$ coincides with some F-feasible mechanism (q,t) on $\operatorname{supp} F$, and assumes the "last" value that (q,t) took, outside of $\operatorname{supp} F$. Notice that $\widehat{\mathcal{M}}(F)$ only contains pairs of non-decreasing and uniformly bounded functions. I claim that $\widehat{\mathcal{M}}(F)$ is compact. Take a sequence $(\hat{q}_n, \hat{t}_n)_{n \in \mathbb{N}} \in (\mathcal{B}[0,1]^2)^{\mathbb{N}}$ such that $(\hat{q}_n, \hat{t}_n) \in \widehat{\mathcal{M}}(F)$ for each n. By Helly's selection theorem there exists a pointwise converging subsequence $(\hat{q}_{n_k}, \hat{t}_{n_k})_{n_k}$, with limit $(\bar{q}, \bar{t})$. By the Dominated Convergent theorem, $(\bar{q}, \bar{t})$ is also the norm limit of $(\hat{q}_{n_k}, \hat{t}_{n_k})_{n_k}$. Standard argument then establish that $(\bar{q}, \bar{t}) \in \widehat{\mathcal{M}}(F)$, which is therefore compact.

Next, consider the mapping $\Gamma : \widehat{\mathcal{M}}(F) \to (\mathcal{B}[0,1])^2$, defined by

$$\Gamma(\hat{q}, \hat{t})(b) = \begin{cases} (\hat{q}, \hat{t})(b) & \text{if } b \in \operatorname{supp} F, \\ 0 & \text{else.} \end{cases}$$

By construction, $\Gamma(\hat{q}, \hat{t})$ is an F-feasible mechanism: $\Gamma(\hat{q}, \hat{t}) \in \mathcal{M}(F)$. The mapping Γ is also Lipschitz continuous, with Lipschitz constant equal to 1. Indeed, take $(\hat{q}, \hat{t}) \in \widehat{\mathcal{M}}(F)$. Then

$$\|\Gamma(\hat{q}, \hat{t})\|_1 = \int\limits_{\operatorname{supp} F} |\hat{q}| \mathrm{d}\lambda + \int\limits_{\operatorname{supp} F} |\hat{t}| \mathrm{d}\lambda \leq \int\limits_0^1 |\hat{q}| \mathrm{d}\lambda + \int\limits_0^1 |\hat{t}| \mathrm{d}\lambda = \|(\hat{q}, \hat{t})\|_1,$$

where the first equality follows form $\Gamma(\hat{q},\hat{t})$ being 0 outside of $\operatorname{supp} F$. Thus, Γ is (Lipschitz) continuous, so it maps compact sets into compact sets, which proves that $\mathcal{M}(F) = \Gamma(\widehat{\mathcal{M}}(\mathcal{F}))$ is compact.

Step 4. The last step consists into proving the following statement.

Claim. *An F-feasible mechanism (q,t) is an extreme point of $\mathcal{M}(F)$ if and only if there exists $x \in \operatorname{supp} F \cup \{0\}$ such that (q,t) is a posted price mechanism (PPM) at x.*

To prove the "if" direction, consider a posted price mechanism (q,t) at $x \in \operatorname{supp} F\{0\}$. Suppose that $(q,t) = \frac{1}{2}(\bar{q},\bar{t}) + \frac{1}{2}(\underline{q},\underline{t})$, for some $(\bar{q},\bar{t})$, $(\underline{q},\underline{t})$ in $\mathcal{M}(F)$. Fix $b \in [0,1]$. Since $q(b)$ is either 0 or 1, then $\bar{q} = \underline{q} = q$. Moreover, if $\bar{t}(b)$ and $\underline{t}(b)$ do not coincide with $t(b)$, then either one– say, $\bar{t}(b)$– must be larger than $t(b)$, violating ex-post Individual Rationality. Thus, (q,t) is an extreme point of $\mathcal{M}(F)$.

For the converse direction, fix $(q,t) \in \mathcal{M}(F)$ and suppose there is no $x \in \operatorname{supp} F \cup \{0\}$ such that (q,t) is a posted price mechanism at x. There are two possible cases.

Case I. The mechanism (q,t) is a posted price mechanism, but the price posted is not an element of the support of F, nor is 0. That is, there is $\tilde{x} \in \operatorname{supp} F$, and $y \notin \operatorname{supp} F$ such that

$$q(b) = \begin{cases} 0 & b \in [0,\tilde{x}) \cup (\operatorname{supp} F)^c, \\ 1 & b \in [\tilde{x},1]; \end{cases} \quad \text{and} \quad t(b) = \begin{cases} 0 & b \in [0,\tilde{x}) \cup (\operatorname{supp} F)^c, \\ y & b \in [\tilde{x},1]. \end{cases}$$

Since (q,t) is incentive compatible, it must be that $[y,\tilde{x}) \cap \operatorname{supp} F = \emptyset$, as otherwise a type in that interval would have an incetive to lie and report $\tilde{x}$ instead of his true type. Pick $\epsilon \in (0, \min\{y - \sup\{b \in \operatorname{supp} F : b < \tilde{x}\}, \tilde{x} - y\})$,[2] and consider the

[2]The interval in which ϵ lies is non-empty, since the support of F is a closed set, and therefore y, being outside of it, cannot be equal to $\sup\{b \in \operatorname{supp} F : b < \tilde{x}\}$.

following posted price mechanisms, $(\bar{q},\bar{t})$ and $(\underline{q},\underline{t})$, defined by $\bar{q}=\underline{q}=q$ and

$$\underline{t}(b)=\begin{cases}0 & b\in[0,\tilde{x})\cup(\operatorname{supp}F)^c,\\ y-\epsilon & b\in[\tilde{x},1].\end{cases}\quad\text{and}\quad \bar{t}(b)=\begin{cases}0 & b\in[0,\tilde{x})\cup(\operatorname{supp}F)^c,\\ y+\epsilon & b\in[\tilde{x},1].\end{cases}$$

By construction, both $(\underline{q},\underline{t})$ and $(\bar{q},\bar{t})$ are F-feasible, and $\frac{1}{2}(\underline{q},\underline{t})+\frac{1}{2}(\bar{q},\bar{t})=(q,t)$, so that (q,t) is not an extreme point of $\mathcal{M}(F)$.

Case II. The mechanism (q,t) is not a posted price mechanism, and, in particular, there exists $b\in\operatorname{supp}F$ such that $q(b)\in(0,1)$. Define the following types:

- $b_1'=\inf\{b\in\operatorname{supp}F:q(b)=1\}$ and $b_1''=\sup\{b\in\operatorname{supp}F:q(b)<1\}$;
- $b_{\frac{1}{2}}'=\inf\{b\in\operatorname{supp}F:q(b)\geq\frac{1}{2}\}$ and $b_{\frac{1}{2}}''=\sup\{b\in\operatorname{supp}F:q(b)<\frac{1}{2}\}$.

(Of course, if the support of F has no gaps $b_1'=b_1''$ and $b_{\frac{1}{2}}'=b_{\frac{1}{2}}''$.) As a matter of notation, let U be the expected payoff that the highest type obtains in the mechanism (q,t), that is, $U=1-\bar{b}_F+q(\bar{b}_F)\bar{b}_F-t(\bar{b}_F)$, where $\bar{b}_F=\max\{\operatorname{supp}F\}$. Notice that there may or may not be a type b which trades with probability 1, i.e. the set $[b_1',1]\cap\operatorname{supp}F$ may or may not be empty. If there is such a type, however, then he must trade at price $1-U$ by incentive compatibility.

I will now define two mechanisms, $(\underline{q},\underline{t})$ and $(\bar{q},\bar{t})$ with midpoint (q,t), that is, $(q,t)=\frac{1}{2}(\underline{b},\underline{t})+\frac{1}{2}(\bar{b},\bar{t})$. For $\kappa\geq 0$, define

$$(\underline{q},\,\underline{t})(b)=\begin{cases}(0,\,0) & b\in(\operatorname{supp}F)^c\cup[0,b_{\frac{1}{2}}''),\\ (2q(b)-1,\,U+\kappa-1+2t(b)) & b\in[b_{\frac{1}{2}}'',b_1')\cap\operatorname{supp}F,\\ (1,\,1+\kappa-U) & b\in[b_1',1]\cap\operatorname{supp}F,\end{cases}$$

and

$$(\bar{q},\, \bar{t})(b) = \begin{cases} (0,\, 0) & b \in (\operatorname{supp} F)^c, \\ (2q(b),\, 2t(b)) & b \in [0, b''_{\frac{1}{2}}) \cap \operatorname{supp} F, \\ (1,\, 1 - U - \kappa) & b \in [b''_{\frac{1}{2}}, 1] \cap \operatorname{supp} F. \end{cases}$$

By construction, $\frac{1}{2}(\bar{q}, \bar{t}) + \frac{1}{2}(\underline{q}, \underline{t}) = (q, t)$. After considerable tedium, one can show that there exists κ that makes both $(\bar{q}, \bar{t})$ and $(\underline{q}, \underline{t})$ incentive compatible and ex-post individually rational, so that (q, t) is not an extreme point of $\mathcal{M}(F)$, as desired.

Appendix B

Appendix to Chapter 2

B.1 Auxiliary results

In this Appendix, I state and prove a few auxiliary results that are needed for the main theorems. Many results depend on the fact that preorders on compact sets admit a maximal element, if they are either open or closed.

Lemma 8 (Bergstrom (1975))**.** *Let X be a compact Hausdorff space. If $\succsim \subseteq X \times X$ is an open preorder, then it has a maximal element.*

Lemma 9. *Let X be a compact Hausdorff space. If $\succsim \subseteq X \times X$ is a closed preorder, it has a maximal element.*

Proof of Lemma 9. [1] Let $Y \subseteq X$ be a $\succsim$-chain, that is, a set which is totally ordered by $\succsim$. Let $\{y_1, \ldots, y_n\} \subseteq Y$ be a finite subset of Y. Without loss of generality, assume that $y_1 \succsim \cdots \succsim y_n$, so that $y_1 \in \bigcap_{k=1}^{n} U_{\succsim}(y_k)$. It follows that the collection $\{U_{\succsim}(y)\}_{y \in Y}$ has the finite intersection property. Since $\succsim$ is closed, each of the $U_{\succsim}(y)$'s is compact. Thus, $\bigcap_{y \in Y} U_{\succsim}(y)$ is non-empty, so that every $\succsim$-chain has an upper bound. By Zorn's Lemma, $\succsim$ must have a maximal element, as desired. □

The next Lemma generalizes a well-known result of finite dimensional linear programs.

Lemma 10. *Let X be a compact metric space, let Y be a compact subset of $\Delta(X)$. Let V be a continuous, affine functional on $\Delta(X)$, and suppose p maximizes V over $\overline{\text{conv}}(Y)$, the closed convex hull of Y. Then there exists $q \in C$ such that $V(p) = V(q)$.*

[1] This proof was suggested by Simone Cerreia-Vioglio.

Proof. Since $p \in \overline{\text{conv}}(Y)$, there exists a sequence of lotteries $(p_n)_n$ living in the convex hull of Y, such that $p_n \to p$. Since, at each n, $p_n \in \text{conv}(C)$, affinity of V implies that there exists $s_n \in Y$ such that $V(s_n) \geq V(p_n)$. Since Y is compact, $(s_n)_n$ has a convergent subsequence, $(s_{n_k})_{n_k}$. Denote by $s \in Y$ its limit. Then continuity of V again implies $V(s_{n_k}) \to V(s)$, and since $V(s_{n_k}) \geq V(p_{n_k})$ for each n_k, it must be that $V(s) \geq V(p)$, as desired. □

For the rest of the appendix, fix a compact continuous game $\Gamma = \langle I, \mathcal{C}, g, (A_i, \succsim^i)_{i\in I}\rangle$. The next result is a corollary of Theorem 3.6.1. in Bogachev (2007), p.191. As a matter of notation, remember that $\mu^i_{a_i}$ is the pushforward of g_{a_i} under belief $\mu^i \in \Delta(A_{-i})$.

Lemma 11. *Consider a pair $(\mu^i, a_i) \in \Delta(A_{-i}) \times A_i$. A measurable function $f : \mathcal{C} \to \mathbb{R}$ is integrable with respect to $\mu^i_{a_i}$ if and only if $f \circ g_{a_i}$ is integrable with respect to μ^i. Moreover,*

$$\int_{\mathcal{C}} f \,\mathrm{d}\mu^i_{a_i} = \int_{A_{-i}} f \circ g_{a_i} \,\mathrm{d}\mu^i.$$

I also need to argue that the topology of the A_i's spaces, and the topology of weak convergence on lotteries of consequences interact well, in the following sense: for any fixed belief $\mu^i \in \Delta(A_{-i})$, given some converging sequence of actions, the corresponding sequence of lotteries converges as well. This is the content of the following two Lemmas.

Lemma 12. *Fix player $i \in I$, and let $(a^n_i)_{n\in\mathbb{N}} \in A^{\mathbb{N}}_i$ be a sequence of actions such that $a^n_i \to a_i$. Let $\mu^i \in \Delta(A_{-i})$. Then the sequence of lotteries $(\mu^i_{a^n_i})_{n\in\mathbb{N}} \in \Delta(A_{-i})^{\mathbb{N}}$ converges to $\mu^i_{a_i}$.*

Proof. By the Portmanteau Theorem (Theorem 15.3 in Aliprantis and Border, 2005, p.508), I only need to show that $\int_{\mathcal{C}} f \mathrm{d}\mu^i_{a^n_i} \to \int_{\mathcal{C}} f \mathrm{d}\mu^i_{a_i}$ for all real continuous functions f. So pick $f \in \mathbf{C}(\mathcal{C})$. Since both f and g are continuous, $f \circ g_{a^n_i}$ is continuous for all n. By continuity again, $f \circ g_{a^n_i} \to f \circ g_{a_i}$ everywhere, and by compactness of A, $f \circ g_{a^n_i}$ is bounded by $\sup_{a\in A} f(g(a))$, which is integrable with respect to μ^i since the latter is a probability measure. By the Dominated Convergence Theorem

(Theorem 11.21 in Aliprantis and Border, 2005, p.415),

$$\lim_n \int_{A_{-i}} f \,\mathrm{d}\mu^i_{a_i^n} = \lim_n \int_{\mathcal{C}} f \circ g_{a_i^n} \,\mathrm{d}\mu^i = \int_{\mathcal{C}} \lim_n f \circ g_{a_i^n} \,\mathrm{d}\mu^i = \int_{\mathcal{C}} f \circ g_{a_i} \,\mathrm{d}\mu^i = \int_{A_{-i}} f \,\mathrm{d}\mu^i_{a_i},$$

as desired. □

Lemma 13. *Fix player $i \in I$. Let $(a_i^n)_{n\in\mathbb{N}} \in A_i^{\mathbb{N}}$ and $(\mu^{i,n})_{n\in\mathbb{N}} \in \Delta(A_{-i})^{\mathbb{N}}$ be two sequences of actions and beliefs, respectively, such that $a_i^n \to a_i$ and $\mu^{i,n} \to \mu^i$. Then $\mu^{i,n}_{a_i^n} \to \mu^i_{a_i}$. In particular, this implies that the mapping $(a_i, \mu^i) \mapsto \mu^i_{a_i}$ has the closed graph property.*

Proof. Given a sequence $(a_i^n)_{n\in\mathbb{N}} \in A_i^{\mathbb{N}}$ with $a_i^n \to a_i$, the sequence of sections of the consequence function $g_{a_i^n}$ converges uniformly to g_{a_i}. This implies that for any continuous function $f : \mathcal{C} \to \mathbb{R}$, the sequence $f \circ g_{a_i^n}$ converges uniformly to $f \circ g_{a_i}$. By Corollary 15.17 p.511 in Aliprantis and Border (2005), the evaluation map $(h, \nu) \to \int_{A_{-i}} h \,\mathrm{d}\nu$ on $\mathbf{C}(A_{-i}) \times \Delta(A_{-i})$ is jointly continuous, where $\mathbf{C}(A_{-i})$ is the space of real continuous functions on A_{-i}, endowed with the $\sup$-norm topology. This implies that for each $f \in \mathbf{C}(\mathcal{C})$,

$$\int_{\mathcal{C}} f \,\mathrm{d}\mu^{i,n}_{a_i^n} = \int_{A_{-i}} f \circ g_{a_i^n} \,\mathrm{d}\mu^{i,n} \to \int_{A_{-i}} f \circ g_{a_i} \,\mathrm{d}\mu^i = \int_{\mathcal{C}} f \,\mathrm{d}\mu^i_{a_i},$$

so that $\mu^{i,n}_{a_i^n} \to \mu^i_{a_i}$ by the Portmanteau theorem, as desired. □

The last two auxiliary results regard the preference extension of Section 2.4.1. The first one establishes that the maximal element of an extension of a preorder must be maximal in the preorder itself. The second shows how the extension of the preferences in a compact continuous game is related to the extension of the ordering over actions given a belief, i.e. to $\succsim^i_{\mu^i}$.

Lemma 14. *Let $\succsim \subseteq X \times X$ be a preorder. Then $x \in X$ is $\succsim$-maximal if and only if there is an extension $\widehat{\succsim}$ of $\succsim$ such that x is $\widehat{\succsim}$-maximal, if and only if there is a completion $\succsim'$ of $\succsim$ such that x is a maximum of $\succsim'$.*

Proof. Suppose x is dominated, so that for some $y \in X$, $y \succ x$. Then in *any* extension, and *any* completion, of $\succsim$, x must be dominated by y. This establishes both of the "if" directions. As for the "only if" parts: the first one is an immediate consequence of the fact that $\succsim$ is an extension of itself. The second one is Lemma 1 in Bade (2005). $\square$

Lemma 15. *Let $\widehat{\Gamma} = \langle I, \mathcal{C}, g, (A_i, \widehat{\succsim}^i)_{i\in I}\rangle$ be a preference extension of the compact continuous game $\Gamma = \langle I, \mathcal{C}, g, (A_i, \succsim^i)_{i\in I}\rangle$. Then for each $i \in I$ and each $\mu^i \in \Delta(A_{-i})$, $\hat{\succsim}^i_{\mu^i}$ extends $\succsim^i_{\mu^i}$. Furthermore, $\hat{R}$ is a completion of $\succsim^i_{\mu^i}$ if and only if there exists a completion $\widehat{\succsim}^i$ of $\succsim^i$ such that $\widehat{\succsim}^i_{\mu^i} = \hat{R}$.*

Proof. Pick $\mu^i \in \Delta(A_{-i})$, and suppose $a_i \succsim^i_{\mu^i} b_i$. Then $\mu^i_{a_i} \succsim^i \mu^i_{b_i}$, so that $\mu^i_{a_i} \hat{\succsim}^i \mu^i_{b_i}$. Similarly, if $a_i \succ^i_{\mu^i} b_i$ then $\mu^i_{a_i} \succ^i \mu^i_{b_i}$ so that $\mu^i_{a_i} \hat{\succ}^i \mu^i_{b_i}$, as desired.

For the second part of the lemma, the *if* direction is obvious. As for the *only if* part, suppose $\hat{R}$ is a completion of $\succsim^i_{\mu^i}$. Define $\succsim' \subseteq \Delta(\mathcal{C}) \times \Delta(\mathcal{C})$ by

$$\succsim' = \mathrm{Trans}\left(\left\{(\mu^i \circ g^{-1}_{a_i}, \mu^i \circ g^{-1}_{b_i}) : a_i \hat{R} b_i\right\} \cup \succsim^i\right),$$

where $\mathrm{Trans}\, S$ is the transitive closure of the relation S. Standard arguments can be used to establish that $\succsim'$ extends $\succsim$. Let now $\hat{\succsim}^i$ be a completion of $\succsim'$, and, hence, of $\succsim^i$. By construction, $\hat{\succsim}^i_{\mu^i} = \hat{R}$, proving the claim. $\square$

B.2 Proof of the Main Results

In this Appendix, I prove the main results in the paper. First, a preliminary result is needed.

Lemma 16. *Let $\succsim^i \subseteq \Delta(\mathcal{C}) \times \Delta(\mathcal{C})$ be player i's preference, and let $\mu^i \in \Delta(A_{-i})$ be her belief. If $\succsim^i$ is a preorder, $\succsim^i_{\mu^i}$ is a preorder too. Furthermore, if $\succsim^i$ is open (resp. closed), then $\succsim^i_{\mu^i}$ is open (resp. closed) too.*

Proof of Lemma 16. Let $\succsim^i$ be a preorder. It is obvious that $\succsim^i_{\mu^i}$ is reflexive. Pick $a_i, b_i, c_i \in A_i$ so that $a_i \succsim^i_{\mu^i} b_i$ and $b_i \succsim^i_{\mu^i} c_i$. Then $\mu^i_{a_i} \succsim^i \mu^i_{b_i}$ and $\mu^i_{b_i} \succsim^i \mu^i_{c_i}$, so that $\succsim^i_{\mu^i}$ is transitive. If $\succsim^i$ is complete, for each $a_i, b_i \in A_i$, either $\mu^i_{a_i} \succsim^i \mu^i_{b_i}$ or $\mu^i_{b_i} \succsim^i \mu^i_{a_i}$, so that either $a_i \succsim^i_{\mu^i} b_i$ or $b_i \succsim^i_{\mu^i} a_i$, that is, $\succsim^i_{\mu^i}$ is complete.

Next, suppose that $\succsim^i$ is open. Let $(a^n_i)_n, (b^n_i)_n \in A^{\mathbb{N}}_i$ be two sequences of actions such that $a^n_i \to a_i$, $b^n_i \to b_i$ and, for each $n \in \mathbb{N}$, $\neg(b^n_i \succ^i_{\mu^i} a^n_i)$. Thus, for each n, $\neg(\mu^i_{b^n_i} \succ^i \mu^i_{a^n_i})$. Since $\succsim^i$ is open, and since, by Lemma 13, $\mu^i_{a^n_i} \to \mu^i_{a_i}$ and $\mu^i_{b^n_i} \to \mu^i_{b_i}$, it follows that $\neg(\mu^i_{b_i} \succ^i \mu^i_{a_i})$, that is, $\neg(b_i \succ^i_{\mu^i} a_i)$. Therefore, $\succsim^i_{\mu^i}$ is open, as desired. The proof for the case in which $\succsim^i$ is closed is analogous, thus omitted. □

Proof of Lemma 6. By Lemma 16, for any belief $\mu^i \in \Delta(A_{-i})$, if $\succsim^i$ is an open preorder, also $\succsim^i_{\mu^i}$ is an open preorder. Since A_{-i} is a compact metric space by Tychonoff Product Theorem, $\Delta(A_{-i})$ is a compact metric space in the topology of weak convergence (see Theorem 6.4 p.45 in Parthasarathy, 1967). By Lemma 8, $\succsim^i_{\mu^i}$ has a maximal element, so that $r_i(\mu^i; \succsim^i)$ is not empty. An analogous argument, using Lemma 9, establishes that $r_i(\mu^i; \succsim^i)$ is non-empty when $\succsim^i$ is a closed preorder.

Finally, pick a sequence of beliefs $(\mu^{i,n})_n \in \Delta(A_{-i})^{\mathbb{N}}$ with $\mu^{i,n} \to \mu^i$, and a sequence of actions $(a^n_i)_n \in A^{\mathbb{N}}_i$ with $a^n_i \to a_i$. Suppose that for each $n \in \mathbb{N}$, $a^n_i \in r_i(\mu^{i,n}; \succsim^i)$. Thus, for each $b_i \in A_i$ and each natural number n, $\neg(\mu^i_{b_i} \succ^i \mu^{i,n}_{a^n_i})$. Using Lemma 13 and the fact that $\succsim^i$ is open, in the limit, it is false that $(\mu^i_{b_i} \succ^i \mu^i_{a_i})$. Since this holds for any $b_i \in A_i$, action a_i must be a best-reply to belief μ^i: $a_i \in r_i(\mu^i; \succsim^i)$. Thus the best-reply correspondence has the closed graph property. Being A_i compact, r_i must be upper hemicontinuous, as desired. □

Proof of Lemma 7. Let $a_i \in r_i(\mu^i)$. Then a_i is $\succsim^i_{\mu^i}$-maximal. Therefore, a_i is $\hat{R}$-maximal for some $\hat{R} \in \operatorname{Compl}(\succsim^i_{\mu^i})$. By Lemma 15, there exists a completion $\widehat{\succsim}^i$ of $\succsim^i$ so that $\widehat{\succsim}^i_{\mu^i} = \hat{R}$, and therefore $a_i \in r_i(\mu^i; \widehat{\succsim}^i)$, so that $a_i \in \tilde{r}_i(\mu^i)$.

For the opposite inclusion, suppose there exists $\widehat{\succsim}^i \in \operatorname{Compl}(\succsim^i)$ such that $a_i \in r_i(\mu^i; \widehat{\succsim}^i)$. Then a_i is a maximum for $\widehat{\succsim}^i_{\mu^i}$, which is a completion of $\succsim^i_{\mu^i}$, and hence by Lemma 14, it is $\succsim^i_{\mu^i}$-maximal. Thus $a_i \in r_i(\mu^i; \succsim^i)$, as desired. □

The next Lemma establishes that the rationalization operator ρ is monotone.

Lemma 17. *If $C, D \in \mathfrak{C}(A)$ are such that $C \subseteq D$, then $\rho(C) \subseteq \rho(D)$.*

Proof of Lemma 17. If $a \in \rho(C)$, for each $i \in I$ there exists a belief μ^i such that $\operatorname{supp} \mu^i \subseteq C_{-i}$ and $a_i \in r_i(\mu^i)$. Since $C \subseteq D$, $\Delta(C) \subseteq \Delta(D)$, so that $a_i \in r_i(\Delta(D_{-i}))$, as needed. □

Corollary 4. *If $C, D \in \mathfrak{C}(A)$ have the best-reply property, then $C \cup D$ has the best-reply property.*

Proof. If C and D have the best-reply property, then $C \subseteq \rho(C)$ and $D \subseteq \rho(D)$, so that $C \cup D \subseteq \rho(C) \cup \rho(D) \subseteq \rho(C \cup D)$, by the monotonicity of ρ. □

I can now prove the main result of Section 2.4, Theorem 8.

Proof of Theorem 8. Suppose first that Γ is a finite game. I will show that the set of iteratively undominated actions, $\widehat{\mathrm{IPR}}$, is non-empty and has the best reply property, and therefore that the set of incomplete preference rationalizable actions, IPR, must also be non-empty. Second, I will show that $\widehat{\mathrm{IPR}}$ and IPR must coincide.
Consider player i, and suppose she has belief $\mu^i \in \Delta(A_{-i})$. Then $\succsim^i_{\mu^i}$ is a preorder on A_i, a finite set, and thus it has a maximal element. Thus $r_i(\mu^i; \succsim^i)$ is non-empty for any belief μ^i. It follows that, for any Cartesian set $C \in \mathfrak{C}(A)$, $\rho(C)$ is non-empty. An elementary inductive argument can thus establish that $(\rho^k(A))_{k\in\mathbb{N}}$ is a nested sequence of non-empty, finite, sets. As such, there must exist $K \in \mathbb{N}$ such that $\rho^K(A) = \rho^{K+n}(A) \neq \emptyset$ for all $n \in \mathbb{N}$, and therefore $\widehat{\mathrm{IPR}} = \rho^K(A)$ is non-empty, as desired. Since $\rho(\widehat{\mathrm{IPR}}) = \widehat{\mathrm{IPR}}$, the latter has the best reply property, and therefore IPR must be non-empty. This also establishes that $\widehat{\mathrm{IPR}} \subseteq \mathrm{IPR}$. Finally, let a be incomplete preference rationalizable. Then there exists a Cartesian set $C \in \mathfrak{C}(A)$ with the best-reply property such that $a \in C$. Since C is a subset of A and since the rationalization operator ρ is monotone, $\rho(C) \subseteq \rho(A)$, so that $a \in \rho(A)$. An easy inductive argument establishes that, for each n, $a \in \rho^n(C) \subseteq \rho^n(A)$, so that $a \in \bigcap_{n\in\mathbb{N}} \rho^n(A) = \widehat{\mathrm{IPR}}$. This proves that $\widehat{\mathrm{IPR}} \supseteq \mathrm{IPR}$, and therefore Definitions 4 and 5 coincide for finite games.

Next, suppose Γ is a compact continuous game, and that the preference of each player $i \in I$, $\succsim^i$, is open. Once more, I will show that $\widehat{\mathrm{IPR}}$ is non-empty and has the best reply property. By Lemma 6, r_i is upper hemicontinuous and non-empty valued, so that $r_i(\Delta(A_{-i}); \succsim^i)$ is non-empty and compact. By Tychonoff Product Theorem, $\rho(A)$ is thus non-empty and compact. By induction, it's easy to show that $\rho^n(A) = \rho(\rho^{n-1}(A))$ is non-empty and compact for each $n \in \mathbb{N}$. By Lemma 17, $(\rho^n(A))_{n\in\mathbb{N}}$ is a nested sequence of non-empty compact sets. Thus, by the Finite Intersection Property of compact sets,

$$\widehat{\mathrm{IPR}} = \bigcap_{n\in\mathbb{N}} \rho^n(A)$$

is non-empty and compact. I claim that $\widehat{\mathrm{IPR}}$ has the best reply property, and, in particular, that $\rho(\widehat{\mathrm{IPR}}) = \widehat{\mathrm{IPR}}$.
Since for each $n \in \mathbb{N}$, $\widehat{\mathrm{IPR}} \subseteq \rho^n(A)$, it must be that $\rho(\widehat{\mathrm{IPR}}) \subseteq \bigcap_{n\geq 0} \rho^n(A)$ (where I adopt the convention that $\rho^0(A) = A$). Thus $\rho(\widehat{\mathrm{IPR}}) \subseteq \widehat{\mathrm{IPR}}$. For the opposite

inclusion, consider an iteratively undominated action profile $a \in \widehat{\mathrm{IPR}}$, and pick $i \in I$, so that $a_i \in \widehat{\mathrm{IPR}}_i$. Thus there is a sequence of beliefs $\mu^{i,n} \in \Delta(A_{-i})^{\mathbb{N}}$ such that, for each $n \in \mathbb{N}$, $\operatorname{supp} \mu^{i,n} \subseteq \operatorname{proj}_{A_{-i}} \rho^{n-1}(A)$. Since $\Delta(A_{-i})$ is a compact metric space, the sequence has a convergent subsequence. Thus, assume without loss of generality that $\mu^{i,n} \to \mu^i$. Then, for each $n \in \mathbb{N}$, $\mu^i(\operatorname{proj}_{A_{-i}} \rho^n(A)) = 1$, and hence

$$\lim_{n\to\infty} \mu^i(\operatorname{proj}_{A_{-i}} \rho^n(A)) = \mu^i(\lim_n \operatorname{proj}_{A_{-i}} \rho^n(A)) = \mu^i(\operatorname{proj}_{A_{-i}} \widehat{\mathrm{IPR}}) = 1.$$

Since r_i has the closed graph property, $a_i \in r_i(\mu^i; \succsim^i)$ and therefore $a_i \in r_i(\widehat{\mathrm{IPR}}_{-i})$. Therefore, $\widehat{\mathrm{IPR}} = \rho(\widehat{\mathrm{IPR}})$, as claimed. Thus $\widehat{\mathrm{IPR}}$ has the best-reply property, and since it is non-empty, IPR must also be non-empty. Finally, I need to show that $\mathrm{IPR} = \widehat{\mathrm{IPR}}$. Consider an incomplete preference rationalizable action profile $a \in \mathrm{IPR}$. Then there exists a set C with the best-reply correspondence that contains a. The same inductive argument used for part 1 of the Theorem establishes that $a \in \rho^n(C) \subseteq \rho^n(A)$ for each $n \in \mathbb{N}$, so that $a \in \widehat{\mathrm{IPR}}$, as desired. □

Proof of Proposition 1. By Lemma 14 and Lemma 15, for each player $i \in I$ and each belief $\mu^i \in \Delta(A_{-i})$, it holds that $r_i(\mu^i; \hat{\succsim}^i) \subseteq r_i(\mu^i; \succsim^i)$. It follows immediately that if $C \in \mathfrak{C}(A)$ has the best-reply property when $(\widehat{\succsim}^j)_{j\in I}$ is the vector of players' preferences, then it must also have the best-reply property under the original preferences. □

Proof of Corollary 3. For each player i, $\succsim^i$ admits an expected multi-utility representation $\mathcal{U}_i$ such that all the $u_i \in \mathcal{U}_i$ are Aumann utilities (Gorno, 2017), i.e. they represent (all of) the expected utility completions of $\succsim^i$. Let $\big(\widehat{\succsim}^i\big)_{i\in I}$ be a profile of expected utility completions of players' preferences. Then $\widehat{\Gamma} = \langle I, \mathcal{C}, g, \big(A_i, \widehat{\succsim}^i\big)_{i\in I}\rangle$ is a preference completion of Γ. Since a complete preorder is closed if and only if it is open, by Theorem 8, the rationalizable set of $\widehat{\Gamma}$ is non-empty and compact. By Theorem 1, then, IPR must also be non-empty. □

Proof of Proposition 2. Fix $\mu^i \in \Delta(A_{-i})$. I need to establish that

$$r_i(\mu^i; \succsim^i) = \bigcup \big\{ r_i(\mu^i; u_i) : \ u_i \in \mathcal{U}_i \big\}.$$

Indeed, by results in Evren (2014), maximal elements of $\widehat{\succsim}^i$, over compact convex sets are maximizers of *some* of the Aumann utilities in the Evren representation. and elements of $r_i(\mu^i; \succsim^i)$ correspond to maximal lotteries over the set $M = \{\mu^i_{a_i} :$

$a_i \in A_i\}$. Since the consequence function is continuous, M is compact, so one can consider its closed convex hull, $\overline{\text{conv}}\, M$. Take a $\succsim^i$-maximal element over $\overline{\text{conv}}\, M$, call it p. Then there is a Aumann utility u_i of i such that $p \in \arg\max_{q \in \overline{\text{conv}}\, M} \int u_i \mathrm{d}q$. The operator $V(r) = \int u \mathrm{d}r$ is affine and continuous in the measure, so by Lemma 10, it is also maximized by some lottery in M, that is, by an element $\mu^i_{a_i}$ for some $a_i \in A_i$, as desired. □

Before I can prove the main result of Section 2.5, I need an auxiliary lemma of independent interest.

Lemma 18. BFR^n_i *has a closed graph for each* $i \in I$ *and each* $n \in \mathbb{N}$.

Proof. The proof is by induction. The basis step is immediate since $\mathrm{BFR}^i_0(u_i) = A_i$ for each player i and each payoff-type u_i, and A_i is compact. So suppose BFR^m_j has a closed graph for all players $j \in I$ and all $m \leq n$, for some $n \in \mathbb{N}$.

Take a sequence $(a^k_i, u^k_i)_{k \in \mathbb{N}} \in (A_i \times \mathcal{U}_i)^{\mathbb{N}}$ such that $a^k_i \to a_i$, $u^k_i \to u_i$ and for each $k \in \mathbb{N}$, $a^k_i \in \mathrm{BFR}^{n+1}_i(u^k_i)$. I need to show that $a_i \in \mathrm{BFR}^{n+1}_i(u_i)$. For each k there exists a belief $\nu^k_i \in \Delta(A_{-i} \times \mathcal{U}_{-i})$ such that

1. For all $b_i \in A_i$, $\int_{A_{-i}} u^k_i(g(a^k_i, \cdot)) \mathrm{d}\, \mathrm{marg}_{A_{-i}} \nu^k_i \geq \int_{A_{-i}} u^k_i(g(b_i, \cdot)) \mathrm{d}\, \mathrm{marg}_{A_{-i}} \nu^k_i$;
2. $\mathrm{supp}\, \nu^k_i \subseteq \{(a_{-i}, u_{-i}) : a_{-i} \in \mathrm{BFR}^n_{-i}(u_{-i})\}$.

Define $h^k : A_{-i} \to \mathcal{C}$ by $h^k(a_{-i}) = u^k_i(g(a^k_i, a_{-i}))$. Then $h^k \to h$ uniformly, where $h(a_{-i}) = u_i(g(a_i, a_{-i}))$. Since $\Delta(A_{-i} \times \mathcal{U}_{-i})$ is compact, assume without loss of generality that $\nu^k_i \to \nu_i \in \Delta(A_{-i} \times \mathcal{U}_{-i})$. Then $\nu_i(\{(a_{-i}, u_{-i}) : a_{-i} \in \mathrm{BFR}^n_{-i}(u_{-i})\}) = 1$, and moreover $\mathrm{marg}_{A_{-i}} \nu^k_i \to \mathrm{marg}_{A_{-i}} \nu_i$. Since the evaluation map on $\mathbf{C}(A_{-i}) \times \Delta(A_{-i})$, $(H, \mu) \mapsto \int_{A_{-i}} H \mathrm{d}\mu$ is jointly continuous,

$$\int\limits_{A_{-i}} u^k_i(g(a^k_i, \cdot)) \mathrm{d}\, \mathrm{marg}_{A_{-i}} \nu^k_i \to \int\limits_{A_{-i}} u_i(g(a_i, \cdot)) \mathrm{d}\, \mathrm{marg}_{A_{-i}} \nu_i,$$

so that $a_i \in \mathrm{BFR}^{n+1}_i(u_i)$ as desired. □

Proof of Theorem 9. I will show by induction that, for each $m \in \mathbb{N}$, and each $i \in I$

$$\mathrm{BFR}^m_i(\mathcal{U}_i) = \rho^m_i(A) = r_i(\Delta(\rho^{m-1}_{-i}(A)); \succsim^i).$$

For $m = 0$, the statement is true by construction, so suppose it holds for any $m \leq n$, for some natural number n.

For the "$\subseteq$" inclusion, consider player $i \in I$ and take $a_i \in \mathrm{BFR}_i^{n+1}(\mathcal{U}_i)$. Then there exists a utility $u_i \in \mathcal{U}_i$ and a belief $\nu^i \in \Delta(A_{-i} \times \mathcal{U}_{-i})$ such that

1. $a_i \in \arg\max_{b_i} \int_{A_{-i}} u_i(g(b_i, \cdot)) \mathrm{d} \operatorname{marg}_{A_{-i}} \nu^i$,
2. $\operatorname{supp} \nu^i \subseteq \big\{(a_{-i}, u_{-i}) : a_{-i} \in \mathrm{BFR}_{-i}^n(u_{-i})\big\}$.

From condition 1 and Proposition 2, it follows again that $a_i \in r_i(\operatorname{marg}_{A_{-i}} \nu^i)$. Hence, it suffices to show that $\operatorname{supp} \operatorname{marg}_{A_{-i}} \nu^i \subseteq \rho_{-i}^n(A)$. To this end, pick $a_{-i} \in \operatorname{supp} \operatorname{marg}_{A_{-i}} \nu^i$. For each $j \neq i$ there exists $u_j \in \mathcal{U}_j$ such that $a_j \in \mathrm{BFR}_j^n(u_j)$. By the inductive hypothesis, $a_{-i} \in \rho_{-i}^n(A)$, so that $\operatorname{marg}_{A_{-i}} \nu^i(\rho_{-i}^n(A)) = 1$. Consequently, $a_i \in \rho_i^{n+1}(A)$, as desired.

For the "$\supseteq$" inclusion, consider $a_i \in \rho_i^{n+1}(A)$. There exists $\mu^i \in \Delta(\rho_{-i}^n(A))$ such that $a_i \in r_i(\mu^i; \succsim^i)$. By Proposition 2, there exists $u_i \in \mathcal{U}_i$ such that $a_i \in r_i(\mu^i; u_i)$. I claim that $a_i \in \mathrm{BFR}_i^{n+1}(u_i)$. To this end, I need to find a belief $\nu^i \in \Delta(A_{-i} \times \mathcal{U}_{-i})$ such that conditions 1 and 2 above hold. By Lemma 18, $\mathrm{BFR}_{-i}^n : \mathcal{U}_{-i} \rightrightarrows A_{-i}$ has a closed graph. Thus the correspondence $\mathcal{J} : \mathrm{BFR}_{-i}^n(\mathcal{U}_i) \rightrightarrows \mathcal{U}_{-i}$ defined by

$$\mathcal{J}(a_{-i}) = \{u_{-i} \in \mathcal{U}_{-i} : a_{-i} \in \mathrm{BFR}_{-i}^n(u_{-i})\}$$

also has a closed graph. Therefore, $\mathcal{J}$ is measurable (Aliprantis and Border, 2005, Theorem 18.20 p.606), and by a version of the Kuratowski-Ryll-Nardzewski theorem (Aliprantis and Border, 2005, Theorem 18.13, p.600) it admits a measurable selection. Let $f : \mathrm{BFR}_i^n(\mathcal{U}_{-i}) \to \mathcal{U}_{-i}$ be a measurable selection, so that for each $a_{-i} \in \mathrm{BFR}_{-i}^n(\mathcal{U}_{-i})$, $a_{-i} \in \mathrm{BFR}_{-i}^n(f(a_{-i}))$. By a result by Buckley (1974), the graph of f is a measurable subset of $A_{-i} \times \mathcal{U}_{-i}$. Therefore, for each measurable set $D \in \mathcal{B}(A_{-i} \times \mathcal{U}_{-i})$, $D \cap \mathrm{Gr}(f)$ is measurable. Define the map $\tilde{\mu}^i : \mathcal{B}(A_{-i} \times \mathcal{U}_{-i}) \to [0,1]$ by

$$D \mapsto \tilde{\mu}^i(D) = \mu^i(\operatorname{proj}_{A_{-i}}(D \cap \mathrm{Gr}(f))).$$

It can be easily checked that $\tilde{\mu}^i$ is a well-defined probability measure on $A_{-i} \times \mathcal{U}_{-i}$. Moreover, by construction and the inductive hypothesis, it holds that

$$\operatorname{supp} \tilde{\mu}^i \subseteq \{(a_{-i}, u_{-i}) : a_{-i} \in \mathrm{BFR}_{-i}^n(u_{-i})\}.$$

Finally, notice that by construction $a_i \in r_i(\operatorname{proj}_{A_{-i}} \tilde{\mu}^i; u_i)$. Thus conditions 1 and 2 above are satisfied, and $a_i \in \mathrm{BFR}_i^{n+1}(u_i)$, as desired. □

Before proving Theorem 10, I need to formally define the type space. I will here first review the construction of Brandenburger and Dekel (1993). Define inductively a sequence of sets $(X_n)_{n\in\mathbb{N}}$ by $X_0 = \bigtimes_{i\in I}\mathcal{U}_i$ and $X_n = X_{n-1}\times\Delta(X_{n-1})$. The set X_n is the set of n-th order beliefs of each player. Define $\hat{T} = \bigtimes_{n\geq 0}\Delta(X_n)$, and notice that $\hat{T}$ is compact by Tychonoff product theorem. Create $|I|$ copies of $\widehat{T}$, called $\widehat{T}_1, \ldots, \widehat{T}_{|I|}$, with the interpretation that $\hat{T}_i$ is the set of all types of player i. Say that a *type* $t_i = (t_i^n)_{n\in\mathbb{N}} \in \hat{T}_i$ is *coherent* if for each $n \geq 2$, $\mathrm{marg}_{X_{n-2}}\, t_i^n = t_i^{n-1}$. Let $\hat{T}_{1,i}$ be the set of all coherent types of player i. It is easy to prove that $\hat{T}_{1,i}$ is closed for each player i, hence compact. Proposition 1 in Brandenburger and Dekel (1993) says that there is a (canonical) homeomorphism $\hat{\kappa}_i : \hat{T}_{1,i} \to \Delta(\mathcal{U}\times\hat{T}_{-i})$. For each natural number $n \geq 2$ and each player i, inductively define $\hat{T}_{n,i}$ by

$$\hat{T}_{n,i} = \{t_i \in \hat{T}_{n-1,i} : \hat{\kappa}_i(t_i)(\mathcal{U}\times\widehat{T}_{n-1,-i}) = 1\}.$$

Define $\tilde{T}_i = \bigcap_{n\in\mathbb{N}}\hat{T}_{n,i}$. This is the space of all types of player i that are coherent and have common belief in coherency of each opponent. By Proposition 2 in Brandenburger and Dekel (1993), there is a homeomorphism $\tilde{\kappa}_i : \tilde{T}_i \to \Delta(\mathcal{U}\times\tilde{T}_{-i})$. Then $\tilde{T}_i$ contains as a subset what I have been calling the universal type space, the space of all belief hierarchies of player i such that (1) each player knows her type, and (2) condition (1) is common belief. The inductive construction is standard. Define

$$\mathcal{T}_i^1 = \{t_i \in \tilde{T}_i : \exists p \in \Delta(\mathcal{U}_{-i}), \exists u_i \in \mathcal{U}_i, t_i^1 = \delta_{u_i}\otimes p\},$$

where t_i^1 is the first-order belief of type t_i, and δ_x is the Dirac measure at x. In other words, $\mathcal{T}_i^1$ contains all types that know their own payoff function. Notice that $\mathcal{T}_i^1$ is a closed, hence compact, subset of $\tilde{T}_i$. Next, use the canonical homeomorphism $\tilde{\kappa}_i : \tilde{T}_i \to \Delta(\mathcal{U}_i\times\tilde{T}_{-i}))$ to inductively define $\mathcal{T}_i^n$, as follows. If $\mathcal{T}_i^{n-1}$ has been defined for each $i \in I$ and some $n \in \mathbb{N}$, then

$$\mathcal{T}_i^n = \{t_i \in \mathcal{T}_i^{n-1} : \tilde{\kappa}_i(t_i)(\mathcal{U}_{-i} \times \mathcal{T}_{-i}^{n-1}) = 1\}.$$

Again, it's an easy inductive argument to show that $\mathcal{T}_i^n$ is non-empty and compact. Thus,

$$\mathcal{T}_i = \bigcap_{n \in \mathbb{N}} \mathcal{T}_i^n,$$

is also non-empty and compact, and is the (belief-closed) subspace of $\tilde{T}_i$ satisfying conditions (1) and (2) above. In the following, I will refer to $(\mathcal{T}_i, \kappa_i)_{i \in I}$ as the type space, where κ_i is the composition of $\tilde{\kappa}_{i|\mathcal{T}_i}$ with the canonical homeomorphism between $\{t_i \in \hat{\mathcal{T}}_i : \tilde{\kappa}_i(t_i)(\mathcal{U}_i \times \tilde{T}_{-i}) = 1\}$ and $\Delta(\mathcal{U}_{-i} \times \tilde{\mathcal{T}}_{-i})$. Finally, define $\mathcal{T}_i[u_i]$ as

$$\mathcal{T}_i[u_i] = \{t_i \in \mathcal{T}_i : \exists p \in \Delta(\mathcal{U}_{-i}), t_i^1 = \delta_{u_i} \otimes p\}.$$

This is the set of types of player i whose payoff function is $u_i \in \mathcal{U}_i$.

Proof of Theorem 10. I present the proof for a game Γ with only two players. This is not restrictive, but makes the notation much simpler. It suffices to show that for each $n \in \mathbb{N}$, each $(u_i)_{i \in I} \in \bigtimes_i \mathcal{U}_i$,

$$\mathrm{BFR}_i^n(u_i) = \mathrm{ICR}_i^n(\mathcal{T}_i[u_i]).$$

The proof is by induction. Suppose the above equality holds for each player $j \in I$ at each natural number strictly smaller than $n \in \mathbb{N}$. Consider $a_i \in \mathrm{ICR}_i^n(\mathcal{T}_i[u_i])$. Then there exists a type $t_i \in \mathcal{T}_i[u_i]$ with payoff function $u_i \in \mathcal{U}_i$, and a conjecture $\nu^i \in \Delta(A_{-i} \times \mathcal{U}_{-i} \times \mathcal{T}_{-i})$ such that conditions (1)-(3) in the definition of ICR are met:

1. $a_i \in r_i(\nu^i; u_i)$,

2. $\mathrm{marg}_{\mathcal{U}_{-i} \times \mathcal{T}_{-i}} \nu^i = \kappa_i(t_i)$

3. $\nu^i(\mathrm{Gr}\,\mathrm{ICR}_{-i}^{n-1}) = 1$.

Let $\xi^i = \operatorname{marg}_{A_{-i}\times\mathcal{U}_{-i}} \nu^i$. By the inductive hypothesis, $\xi^i(\operatorname{Gr}\operatorname{BFR}^{n-1}) = 1$, and clearly $a_i \in r_i(\xi^i; u_i)$. Thus conditions (i) and (ii) of the definition of BFR_i^n are met: $a_i \in \operatorname{BFR}_i^n(u_i)$.

For the converse inclusion, let $a_i \in \operatorname{BFR}_i^n(u_i)$. Then there exists a belief $\xi^i \in \Delta(A_{-i} \times \mathcal{U}_{-i})$ such that

i $a_i \in r_i(\xi^i; u_i)$;

ii $\xi^i(\operatorname{Gr}(\operatorname{BFR}_{-i}^{n-1})) = 1$.

I will use an argument similar to the proof of Theorem 9 to construct a belief $\nu^i \in \Delta(A_{-i} \times \mathcal{U}_{-i} \times \mathcal{T}_{-i})$ such that conditions (1)-(3) above are satisfied for a_i. By the inductive hypothesis, for each $(a_{-i}, u_{-i}) \in \operatorname{supp}\xi^i$, there exists $t_{-i} \in \mathcal{T}_{-i}[u_{-i}]$ such that $a_{-i} \in \operatorname{ICR}_{-i}^{n-1}(t_{-i})$. Define a correspondence $\mathcal{L} : \operatorname{supp}\xi^i \rightrightarrows \mathcal{T}_{-i}$ by

$$\mathcal{L}(a_{-i}, u_{-i}) = \{t_{-i} \in \mathcal{T}_{-i}[u_{-i}] : a_{-i} \in \operatorname{ICR}_{-i}^{n-1}(t_{-i})\}.$$

Then $\mathcal{L}$ has a closed graph. Indeed, take a sequence $(a_{-i}^\ell, u_{-i}^\ell, t_{-i}^\ell)_{\ell\in\mathbb{N}}$ such that, for each $\ell \in \mathbb{N}$, $t_{-i} \in \mathcal{L}(a_{-i}, u_{-i})$, and $(a_{-i}^\ell, u_{-i}^\ell, t_{-i}^\ell) \to (a_{-i}, u_{-i}, t_{-i})$. Then for each ℓ there exists $\nu^{-i,\ell} \in \Delta(A_i \times \mathcal{U}_i \times \mathcal{T}_i)$ that satisfies conditions (1)-(3) above. Since $A_i \times \mathcal{U}_i \times \mathcal{T}_i$ is compact, the sequence $\nu^{-i,\ell}$ has a convergent subsequence, so assume without loss of generality that $\nu^{-i,\ell} \to \nu^{-i}$. Since the canonical homeomorphism κ_{-i} is continuous, $\operatorname{marg}_{\mathcal{U}_{-i}\times\mathcal{T}_{-i}} \nu^i = \kappa_{-i}(t_{-i})$. Moreover, by the upper hemicontinuity of the best-reply correspondence, $a_{-i} \in r_{-i}(\nu^{-i}; u_{-i})$. Finally, the support restriction in (3) is also satisfied by standard arguments. This establishes that $\mathcal{L}$ has a closed graph.

Since $\mathcal{L}$ has a closed graph, it is measurable, and therefore by a version of the Kuratowski-Ryll-Nardzewski theorem, it admits a measurable selection. Let $f : \operatorname{supp}\xi^i \to \mathcal{T}_{-i}$ be one such measurable selection, so that

$$(a_{-i}, u_{-i}) \mapsto f(a_{-i}, u_{-i}) \implies a_{-i} \in \operatorname{ICR}_{-i}^{n-1}(f(a_{-i}, u_{-i})) \quad \text{and} \quad f(a_{-i}, u_{-i}) \in \mathcal{T}_{-i}[u_{-i}].$$

Define a belief $\nu^i \in \Delta(A_{-i} \times \mathcal{U}_{-i} \times \mathcal{T}_{-i})$ as follows: for each $D \in \mathcal{B}(A_{-i} \times \mathcal{U}_{-i} \times \mathcal{T}_{-i})$,

$$\nu^i(D) = \xi^i(\operatorname{proj}_{A_{-i}\times\mathcal{U}_{-i}}(D \cap \operatorname{Gr} f)).$$

Finally, define a type $t_i \in \mathcal{T}_i[u_i]$ as follows: the first order belief of t_i is

$$t_i^1 = \delta_{u_i} \otimes \mathrm{marg}_{\mathcal{U}_{-i}} \nu^i,$$

and her k-th order belief by

$$t_i^k = \delta_{u_i} \otimes \mathrm{marg}_{\mathcal{U}_{-i} \times \Delta(X_{k-1})} \nu^i,$$

where $\Delta(X_{k-1})$ is the set of $k-1$-th order beliefs of player $-i$. In other words, t_i is the unique type that has payoff function u_i and beliefs ν^i on her opponent. By construction, $a_i \in \mathrm{ICR}_i(t_i)$, which concludes the proof. □

B.3 A counterexample

In this section, I show by counterexample that the rationalizable set of Γ is typically *not* equal to the union of the rationalizable sets over the preference completion games of Γ. I will assume that players' preferences satisfy Evren's multi-utility theorem, and focus only on expected utility preference completion games of Γ. This is without loss of generality by Proposition 2 and Lemma 7.

Counterexample. Consider a finite two-player game with incomplete preferences, Γ. Player 1, Ann, chooses the row, and her preferences are complete, whereas player 2, Bob, chooses the column, and his preferences are incomplete. Ann is an Expected Utility maximizer, whereas Bob's preferences admit an Evren multi-utility representation $\mathcal{U}_B = \mathrm{conv}\{u_B, \tilde{u}_B\}$. The payoff matrices for the two extreme points of $\mathcal{U}_B$ are shown below.

Ann \ Bob	L	C	R
U	$2, 2.5$	$1, 1$	$0, 0$
M	$1, 2$	$2, 1$	$1, 0$
D	$0, 1$	$1, 2$	$2, 0$

u_B

Ann \ Bob	L	C	R
U	$2, 0$	$1, 2$	$0, 1$
M	$1, 0$	$2, 3$	$1, 4$
D	$0, 0$	$1, 1$	$2, 3$

$\tilde{u}_B$

A generic element of $\mathcal{U}_B$ takes the form $\lambda u_B + (1-\lambda)\tilde{u}_B$, for some $\lambda \in [0,1]$. Thus, the generic preference completion of Γ, can be indexed by λ, and is denoted

Γ_λ. The payoff matrix of Γ_λ is shown below.

		Bob L	Bob C	Bob R
Ann	U	$2, 2.5\lambda$	$1, 2-\lambda$	$0, 1-\lambda$
Ann	M	$1, 2\lambda$	$2, 3-2\lambda$	$1, 4-4\lambda$
Ann	D	$0, \lambda$	$1, 1+\lambda$	$2, 3-3\lambda$

Γ_λ

First, I will argue that the Incomplete Preferences rationalizable set, IPR is the entire set of action profiles, A. Start with Ann's actions. U is the best reply to L, M is the best reply to C and D is the best reply to R, so that all of Ann's actions are justifiable. Next, consider Bob. His best reply to U contains both L and C (but not R), whereas his best reply to D contains both C and R (but not L). Consequently, all of Bob's actions are justifiable, and therefore $\text{IPR} = \bigcap_{n\in\mathbb{N}} \rho^n(A) = \rho^1(A) = A$.

Claim. *Let $\rho^\infty(A;\Gamma_\lambda)$ be the rationalizable set of the preference completion game Γ_λ. Then for each $\lambda \in [0,1]$, $(D,L) \notin \rho^\infty(A;\Gamma_\lambda)$.*

The proof works by showing that, for any λ, either D or L is iteratively dominated, so that (D,L) cannot be part of the rationalizable set of the preference completion game Γ_λ. First, notice that L is dominated by C if

$$2.5\lambda < 2 - \lambda, \tag{B.1}$$

$$2\lambda < 3 - 2\lambda, \tag{B.2}$$

$$\lambda < 1 + \lambda. \tag{B.3}$$

Conditions (1)-(3) are jointly satisfied if and only if $\lambda < \frac{4}{7}$. Thus, in such case, $(D,L) \notin \rho^\infty(A;\Gamma_\lambda)$, as L is not justifiable.

Next, notice that R is dominated by C if

$$1 - \lambda < 2 - \lambda, \tag{B.4}$$

$$4 - 4\lambda < 3 - 2\lambda, \tag{B.5}$$

$$3 - 3\lambda < 1 + \lambda. \tag{B.6}$$

Conditions (4)-(6) are jointly satisfied if $\lambda > \frac{1}{2}$. But now notice that, if R is eliminated, then D is iteratively dominated (by M, for instance). Thus also when $\lambda > \frac{1}{2}$,

$(D, L) \notin \rho^{\infty}(A; \Gamma_{\lambda})$, concluding the proof of the Claim, and establishing the Counterexample.

Bibliography

Abreu, Dilip and Gul, Faruk. Bargaining and Reputation. *Econometrica*, 68:85–117, 2000.

Aliprantis, Charalambos D. and Border, Kim C. *Infinite Dymensional Analysis. A Hitchhiker's Guide*. Springer-Verlag, Berlin, third edition, 2005.

Aumann, Robert J. Utility theory without the completeness axiom. *Econometrica*, 30(3):445–462, 1962.

Bade, Sophie. Nash equilibrium in games with incomplete preferences. *Economic Theory*, 26:309–332, 2005.

Battigalli, Pierpaolo. Rationalizability in infinite, dynamic games with incomplete information. *Research in Economics*, 57:1–38, 2003.

Battigalli, Pierpaolo and Siniscalchi, Marciano. Rationalization and incomplete information. *Advances in Theoretical Economics*, 3, 2003.

Battigalli, Pierpaolo, Tillio, Alfredo Di, Grillo, Edordo, and Penta, Antonio. Interactive epistemology and solution concepts for games with asymmetric information. *The B.E. Journal of Theoretical Economics*, 11, 2011.

Bergemann, Dirk and Morris, Stephen. Belief-free rationalizability and informational robustness. *Games and Economic Behavior*, 104:744–759, 2017.

Bergemann, Dirk and Välimäki, Juuso. Information in Mechanism Design. In Blundell, Richard, Newey, Whitney K., and Persson, Torsten, editors, *Advances in Economics and Econometrics: Theory and Applications, Ninth World Congress*, volume 1 of *Econometric Society Monographs*, page 186–221. Cambridge University Press, 2006. doi: 10.1017/CBO9781139052269.007.

Bergstrom, Theodore C. Maximal elements of acyclic relations on compact sets. *Journal of Economic Theory*, 10:403–404, 1975.

Bernheim, B. Douglas. Rationalizable strategic behavior. *Econometrica*, 52(4): 1007–1028, 1984.

Bewley, Truman F. Knightian decision theory. part i. *Decisions in economics and finance*, 25(2):79–110, 2002.

Board, Simon and Pycia, Marek. Outside Options and the Failure of the Coase Conjecture. *American Economic Review*, 104:656–671, 2014.

Bogachev, Vladimir I. *Measure Theory. Volume I*. Springer-Verlag, Berlin, 2007.

Borgers, Tilman. *An Introduction to the Theory of Mechanism Design*. Oxford University Press, New York, NY, 2015.

Brandenburger, Adam and Dekel, Eddie. Rationalizability and correlated equilibria. *Econometrica*, 55(6):1391–1402, 1987.

Brandenburger, Adam and Dekel, Eddie. Hierarchies of beliefs and common knowledge. *Journal of Economic Theory*, 59, 1993.

Buckley, James J. Graphs of measurable functions. *Proceedings of the American Mathematical Society*, 44:78–80, 1974.

Carroll, Gabriel. An efficiency theorem for incompletely known preferences. *Journal of Economic Theory*, 145:2463–2470, 2010.

Che, Yeon-Koo and Sákovics, József. A Dynamic Theory of Hold-up. *Econometrica*, 72:1063–1103, 2004.

Chen, Yi-Chun, Luo, Xiao, and Qu, Chen. Rationalizability in general situations. *Economic Theory*, 61:147–167, 2016.

Chung, Kai Lai. *A Course in Probability Theory*. Academic Press, London, UK, 3rd edition, 2001.

Coase, Ronald Henry. Durability and Monopoly. *Journal of Law and Economics*, 15: 143–149, 1972.

Compte, Olivier and Jehiel, Philippe. On the role of Outside Options in Bargaining with Obstinate Parties. *Econometrica*, 70:1477–1517, 2002.

Condorelli, Daniele and Szentes, Balázs. Information Design in the Holdup Problem. *Journal of Political Economy*, 128(1):681–709, 2020.

De Meza, David and Lockwood, Ben. Does Asset Ownership Always Motivate Managers? Outside Options and the Property Rights Theory of the Firm. *Quarterly Journal of Economics*, 113:361–386, 1998.

Dekel, Eddie, Fudenberg, Drew, and Morris, Stephen. Interim correlated rationalizability. *Theoretical Economics*, 2:15–40, 2007.

Dubra, Juan, Maccheroni, Fabio A., and Ok, Efe A. Expected utility theory without the completeness axiom. *Journal of Economic Theory*, 115:118–133, 2004.

Evren, Özgür. Scalarization methods and expected multi-utility representations. *Journal of Economic Theory*, 151:30–63, 2014.

Fudenberg, Drew, Levine, David K., and Tirole, Jean. Incomplete Information Bargaining with Outside Opportunities. *Quarterly Journal of Economics*, pages 37–50, 1987.

Gershkov, Alex, Moldovanu, Benny, Strack, Philipp, and Zhang, Mengxi. A Theory of Auctions with Endogenous Valuations. *Journal of Political Economy*, 129(4): 1011–1051, 2021.

Gorno, Leandro. A strict expected multi-utility theorem. *Journal of Mathematical Economics*, 71:92–95, 2017.

Grossman, Sanford J. and Hart, Oliver D. The Costs and Benefits of Ownership: A Theory of Vertical and Lateral Integration. *Journal of Political Economy*, 94: 691–719, 1986.

Gul, Faruk. Unobservable Investment and the Hold-up Problem. *Econometrica*, 69: 343–376, 2001.

Gul, Faruk, Sonnenschein, Hugo, and Wilson, Robert. Foundations of Dynamic Monopoly and the Coase Conjecture. *Journal of Economic Theory*, 39:155–190, 1986.

Harsanyi, John C. Games with incomplete information played by bayesian players. parts i-iii. *Management Science*, 14:159–182,320–344,486–502, 1967.

Hart, Oliver D. and Moore, John. Property Rights and the Nature of the Firm. *Journal of Political Economy*, 98:1119–1158, 1990.

Hermalin, Benjamin E. and Katz, Michael L. Information and the Hold-up Problem. *RAND Journal of Economics*, 40:405–423, 2009.

Hwang, Ilwoo and Li, Fei. Transparency of Outside Options in Bargaining. *Journal of Economic Theory*, 167:116–147, 2017.

Klein, Benjamin, Crawford, Robert G., and Alchian, Armen A. Vertical Integration, Appropriable Rents and the Competitive Contracting Process. *The Journal of Law and Economics*, 21:297–326, 1978.

Kokkala, Juho, Poropudas, Jirka, and Virtanen, Kai. Rationalizable Strategies in Games With Incomplete Preferences. *Theory and Decision*, 86:185–204, 2019.

Machina, Mark. "Expected Utility" Analysis without the Independence Axiom. *Econometrica*, 50:277–323, 1982.

McClellan, Andrew. Outside Options and Optimal Bargaining Dynamics. *Working paper*, 2021.

Mensch, Jeffrey. Screening Inattentive Agents. *Working Paper*, 2020.

Mertens, Jean-Francois and Zamir, Shmuel. Formulation of Bayesian Analysis for Games with Incomplete Information. *International Journal of Game Theory*, 14 (1):1–29, 1985.

Morita, Hodaka and Servátka, Maroš. Investment in Outside Options as Opportunistic Behavior: An Experimental Investigation. *Southern Economic Journal*, 85:457–484, 2018.

Myerson, Roger B. Optimal Auction Design. *Mathematics of Operations Research*, 6 (1):58–73, 1981.

Myerson, Roger B. and Satterthwaite, Mark A. Efficient Mechanisms for Bilateral Trading. *Journal of Economic Theory*, 29:265–281, 1983.

Osborne, Martin J. and Rubinstein, Ariel. *A Course in Game Theory*. The MIT Press, Cambridge, MA, 1994.

Park, Jaeok. Potential games with incomplete preferences. *Journal of Mathematical Economics*, 61:58–66, 2015.

Parthasarathy, Kalyanapuram R. *Probability Measures on Metric Spaces*. Academic Press, New York, NY, 1967.

Pearce, David G. Rationalizable strategic behavior and the problem of perfection. *Econometrica*, 52(4):1029–1050, 1984.

Persico, Nicola. Committee Design with Endogenous Information. *The Review of Economic Studies*, 71:165–194, 2004.

Quirk, James P. and Saposnik, Rubin. Admissibility and Measurable Utility Functions. *The Review of Economic Studies*, 92:140–156, 1962.

Ravid, Doron, Roesler, Anne-Katrin, and Szentes, Balazs. Learning Before Trading: On the Inefficiency of Ignoring Free Information. *Journal of Political Economy*, forthcoming.

Roesler, Anne-Katrin and Szentes, Balazs. Buyer-Optimal Learning and Monopoly Pricing. *American Economic Review*, 107:2072–2080, 2017.

Rogerson, William P. Contractual Solutions to the Hold-up Problem. *The Review of Economic Studies*, 59:777–793, 1992.

Rothschild, Michael and Stiglitz, Joseph E. Increasing Risk: I. A Definition. *Journal of Economic Theory*, 2:225–243, 1970.

Rubinstein, Ariel. Perfect Equilibrium in a Bargaining Model. *Econometrica*, 50: 97–109, 1982.

Schmeidler, David. A condition for the completeness of partial preference relations. *Econometrica*, 39(2):403–404, 1971.

Shaked, Avner and Sutton, John. Involuntary Unemployment as a Perfect Equilibrium in a Bargaining Model. *Econometrica*, 52:1351–1364, 1984.

Shapley, Lloyd S. and Rigby, Fred D. Equilibrium points in games with vector payoffs. *Naval Research Logistics Quarterly*, 6:57–61, 1959.

Shi, Xianwen. Optimal Auctions with Information Acquisition. *Games and Economic Behavior*, 74:666–686, 2012.

Szpilrajn, Edward. Sur l'extension de l'ordre partiel. *Fundamenta Mathematicæ*, 16: 386–389, 1930.

Tan, Tommy and Werlang, Sérgio Ribeiro da Costa. The bayesian foundations of solution concepts of games. *Journal of Economic Theory*, 45(2):370–391, 1988.

von Neumann, John and Morgenstern, Oskar. *Theory of Games and Economic Behavior*. Princeton University Press, Princeton, NJ, third edition, 1953.

Wolinsky, Asher. Matching, Search and Bargaining. *Journal of Economic Theory*, 42: 311–333, 1987.

Ziegler, Gabriel. Informational robustness of common belief in rationality. *Mimeo. University of Edinburgh*, 2021.

Ziegler, Gabriel and Zuazo-Garin, Peio. Strategic cautiosness as an expression of robustness to ambiguity. *Games and Economic Behavior*, 119:197–215, 2020.

CURRICULUM VITAE

[redacted]

[redacted]