

The Pennsylvania State University

The Graduate School

College of the Liberal Arts

ESSAYS IN ECONOMIC THEORY

A Dissertation in

Economics

by

Bruno Salcedo Gonzalez

© 2016 Bruno Salcedo Gonzalez

Submitted in Partial Fulfillment
of the Requirements
for the Degree of

Doctor of Philosophy

August 2016

The dissertation of Bruno Salcedo was reviewed and approved* by the following:

Edward Green
Professor of Economics and Mathematics
Dissertation Co-Advisor
Co-chair of Committee

Vijay Krishna
Distinguished Professor of Economics
Dissertation Co-Advisor

Ruilin Zhou
Associate Professor of Economics

Lisa Posey
Associate Professor of Risk Management

Barry Ickes
Professor of Economics
Head of the Department of Economics

*Signatures are on file in the Graduate School.

Abstract

The present dissertation consists of three independent essays on Game Theory and its applications to Economics. The first essay compares two standard notions of rational choice under uncertainty: dominance by *pure* actions, and dominance by *pure or mixed* actions. I show that these two notions are equivalent for agents that exhibit sufficient risk aversion. Moreover, risk aversion is a cardinal property. Thus, the different forms of dominance considered cannot be distinguished based solely on the ordinal data that can be directly inferred from observed choices.

The second essay explores the consequences of relaxing a standard assumption in the analysis of strategic situations: the assumption that choices are made *simultaneously* and *independently*. Choice interdependence can have drastic consequences. For example, it can enable cooperation in some single-shot prisoners' dilemmas without the use of binding contracts or side payments. I consider a class *partially* specified environments in which the actions available to each agent and the agents' preferences are fixed, but the sequential and information structures of choices are design variables. I propose a simple and tractable solution concept—interdependent choice equilibrium—that fully characterizes all the outcomes that can be implemented in such environments without contracts or side-payments.

The third essay proposes a model of dynamic price competition between firms that use automated *pricing algorithms* to set their prices. In my model, algorithms are fixed in the short run but can be revised at exogenous dates, and each firm is able to decode its rival's algorithm. I show that, when the revision opportunities are infrequent relative to the arrival of consumers into the market, *every* sub-game perfect equilibrium of the game leads in the long run to profits that are arbitrarily close to the Pareto frontier. That is, the use of pricing algorithms might not only facilitate collusion, but might inevitably lead to it. In contrast with the plethora of equilibria that typically arises in models of repeated competition, my model is able to generate sharp predictions about the behavior of the firms in the long run.

Contents

List of Figures	vi
Acknowledgements	vii
Chapter 1 Ordinal rationalizability	1
1.1 Single-agent choice problems	4
1.2 Pure and mixed dominance	5
1.3 Timidity	7
1.4 Dominance equivalence and risk aversion	8
1.5 Summary and discussion	12
Chapter 2 Interdependent choices	14
2.1 Cooperation in a prisoners' dilemma	18
2.2 A framework for robust predictions	20
2.3 Interdependent choice equilibrium and robust predictions	22
2.4 When is cooperation possible in a prisoners' dilemma	26
2.5 Implementation via mediated mechanisms	29
2.6 Quasi-sequential implementation	32
Chapter 3 Pricing Algorithms	37
3.1 Two-period two-price prototype model	42

3.2	Dynamic price competition with pricing algorithms	44
3.3	Inevitability of collusion	48
Appendix A Omitted proofs for Chapter 1		53
Appendix B Omitted details from Chapter 2		56
B.1	Alternative definition of extensive form mechanisms	56
B.2	Omitted proofs	58
Appendix C Omitted proofs for Chapter 3		70
C.1	Existence of equilibria	70
C.2	Inevitability of collusion	75
Bibliography		83

List of Figures

1.1	Payoff matrix for introductory example, where $\gamma \in (0, 2)$	2
1.2	vNM utility functions for Example 1.1 with $X = \{1, 2\}$	5
2.1	Payoff matrix for the prisoners' dilemma.	18
2.2	Equilibrium payoffs for a teamwork game.....	24
2.3	ICE for the prisoners' dilemma with $l = 0.5$ and different values of g	27
2.4	Environment induced by a feasible policy μ	28
2.5	A mediated mechanism for the prisoners' dilemma.....	31
3.1	Profit matrix for and “ <i>tit for tat</i> ” algorithm for the prototype model.	43
3.2	Examples of pricing algorithms.	46
B.1	Implementation of b_1 when it is the only dominated action.	61
B.2	Implementation of b_1 when b_2 is also dominated.	61
B.3	Incentives for $a_i^0 \in A_i^{\text{FICR}} \setminus A_i^*$	68
C.1	Algorithm with renegotiation offers.....	80

Acknowledgements

This dissertation is the culmination of a long path that I could not have gone through without the help of many wonderful people.

First and foremost, I wish to acknowledge my greatest fortune of all, which is my wife. I will never understand the craziness in her mind that made her leave everything behind and move to a different continent to start a life with someone she barely knew, but I will always be grateful for it. I would not be where I am otherwise.

My accomplishments do not belong to me any more than they belong to my parents who supported me, trained me, and loved me unconditionally, and did the impossible to make sure that I could reach my goals. I was also lucky to have two wonderful siblings, who were sometimes my auxiliary parents and sometimes my friends, always provided support, advice, and affection, and led the way by example.

I am extremely grateful to my two fantastic advisers Ed Green and Vijay Krishna, both of whom I respect and admire greatly. They got me started me on the path of research and their guidance, support, kindness, lucidity, and eloquence have been invaluable assets since.

I am also indebted to Ruilin Zhou, Joris Pinkse, and Sung Jae Jun, for adopting a micro-theorist as a foster student, and allowing me to work with and learn from them. I want to thank Kalyan Chatterjee for his help and advice, and for consistently providing a space for graduate students to share and discuss their work. Nageeb Ali's aid was crucial in preparing my job market paper, despite the brief duration of our shared time at Penn State. I am also grateful to the rest of the faculty that has helped me over the years including Neil Wallace, Lisa Posney, Herman Bierens, Marc Henry, Kala Krishna, Paul Grieco, Stephen Yeaple and James Jordan, as well as conference and seminar presenters and attendants who have provided valuable comments and suggestions, specially Kenneth Judd, David Rahman, Alexey Kushnir and Tymofiy Mylovanov.

One of my greatest privileges as a graduate student has been having my splendid coauthors and dear friends around: Nail Kashaev, Bruno Sultanum, and Bulat Ga-

farov. And, even if I cannot call them my coauthors *yet*, I am also grateful to other classmates who have helped me and motivated me, including Pathrikrit Basu, Suraj Shekhar, Gustavo Gudiño, Evelyn, Nunes, Lidia Kosenkova, Fatemeh Borhani, Sergii Golovko, David Jenkins, Yu Awaya, and my favorite TA Kaustav Das.

The list would be incomplete without acknowledging those who created the basis of my knowledge and made me love game theory and economics, specially: Levent Ülkü, Tridib Sharma, Andrei Gomberg, Jose Luis Morales, Juan Bautista, and my dear friend and teacher Lizette de la Garza.

I gratefully acknowledge the research support from the Michael Rosenberg and Carol Franzblau-Rosenberg Family, and from the [Human Capital Foundation](#) and particularly Andrey P. Vavilov through the [Center for the Study of Auctions, Procurements, and Competition Policy](#) at the Pennsylvania State University.

All remaining errors are my own.

To my beautiful Day

Chapter 1

Ordinal rationalizability

Abstract This essay is based on joint work with Bulat Gafarov (Gafarov and Salcedo, 2015). We find that, for sufficiently risk-averse agents, strict dominance by *pure or mixed* actions coincides with dominance by *pure* actions in the sense of Börgers (1993), which, in turn, coincides with the classical notion of strict dominance by pure actions when preferences are asymmetric. Since risk-aversion is a *cardinal* feature, all finite single-agent choice problems with *ordinal* preferences admit compatible utility functions which are sufficiently risk-averse as to achieve equivalence between pure and mixed dominance. These two results extend to some infinite environments.

Suppose that a rational agent must choose between three actions: betting that an event E occurs, betting that E does not occur, or not betting at all. The agent's preferences are represented by the von Neumann-Morgenstern (vNM) utility function summarized in Figure 1.1. Notice that the ordinal ranking of action-state pairs remains unchanged as long as $0 < \gamma < 2$. Also, not betting is not strictly dominated by any *pure* action, and, if $\gamma \geq 1$, it is also not strictly dominated by any *mixed* action. However, if $\gamma < 1$, then it becomes strictly dominated by the mixed action which mixes uniformly between betting on E and betting on not E .

Here, γ can be thought of as measuring the degree of concavity or risk-aversion of the agent's vNM utility function. Hence, we see that dominance by pure strategies coincides with dominance by mixed strategies if the agent is sufficiently risk-averse, and there exists a sufficiently risk-averse utility function which is compatible with the given ordinal preferences. In the rest of the paper, we show that these two

	E	not E
bet on E	2	0
bet on not E	0	2
do not bet	γ	γ

Figure 1.1 – Payoff matrix for introductory example, where $\gamma \in (0, 2)$.

observations continue to hold for a large class of decision problems under uncertainty with *ordinal* preferences.

We compare strict dominance by *pure or mixed* actions (M_u -dominance) with the notion of dominance by *pure* actions (P -dominance) introduced by Börgers (1993).¹ An action is P -dominated if and only if it is weakly dominated by a pure action, conditional on any given set of states.² P -dominated actions are always M_u -dominated, but the converse need not be true.

A mixed action could dominate an action that is not P -dominated, because mixing enables the agent to average good and bad outcomes corresponding to different action-state pairs. However, mixing also exacerbates the agent’s uncertainty about the outcome of the environment, by adding uncertainty about the result of using her own randomization device. Hence, the more risk-averse the agent is, the less appealing mixing will be.³ We find that M_u -dominance reduces to P -dominance for sufficiently risk-averse agents, according to a specific measure which we call *timidity* (propositions 1.3 and 1.5).⁴ In particular, the set of sufficiently timid utility functions includes all

¹We index M_u -dominance by u to highlight the fact that it depends on the *cardinal* information embedded in vNM utility functions. In contrast, P -dominance only depends on the agent’s *ordinal* state-contingent preferences over actions.

²In general, when indifference is allowed, for an action to be *strictly* dominated by a pure action implies that it is P -dominated, which implies in turn that it is *weakly* dominated by a pure action. In the generic case in which all state-contingent preferences are strict, these three notions of pure dominance coincide.

³Our research is in a Bayesian framework, so we use “risk” and “uncertainty” as synonyms. Nevertheless, our intuition is closely related to the work of Klibanoff (2001). He asks under which conditions would an uncertainty averse agent be willing to choose mixed actions. As it turns out, the trade-off between averaging outcomes (uncertainty) and increasing variance (risk) plays a prominent role.

⁴These results are similar in spirit to Lemma 1 in Chen and Luo (2012), which implies that, in “concave-like” games, an action is M_u -dominated if and only if it is strictly dominated by a pure action. However, their lemma is interesting only for uncountable environments (including mixed extensions of finite environments). In finite or countable environments—like the ones we consider—if an agent has concave-like preferences, then there exists a pure action which P -dominates every other action.

CARA functions that are sufficiently risk-averse in the familiar sense.

A vNM utility function is said to be *strongly compatible* with the environment if it represents the ordinal preferences of the agent *over action-state pairs*. Any strictly concave and strictly monotone transformation of utility preserves strong compatibility while increasing timidity. In this manner, we find that if either the action space or the state space is finite, then there exists a strongly compatible vNM utility function which generates equivalence between P -dominance and M_u -dominance (Corollary 1.4). However, the degree of timidity required grows linearly with the size of the environment, and there are countable environments in which strong compatibility precludes dominance equivalence.

By relaxing the definition of compatibility, it is still possible to obtain dominance equivalence in a large class of infinite environments. If preferences are interpreted as revealed choices, then it is meaningless to compare rankings across states. We say that a vNM utility function is *compatible* with the environment if it represents the given *state-contingent* ordinal preferences over actions. If only compatibility is required, dominance equivalence is possible in all countable environments satisfying a discreteness assumption (Corollary 1.6).

Our work is closely related to Börgers (1993). Using our language, Börgers' main result can be expressed as follows. For finite environments, if an action is not P -dominated, then there exists a strongly compatible vNM utility function—which may depend on the action—according to which the action is also not M_u -dominated. Also, while Ledyard (1986) works in a very different context, some of his results have important implications for our environment. In particular, his Corollary 5.1 implies that every finite choice environment without P -dominated actions admits a compatible vNM utility function—which may not be strongly compatible—according to which there are no M_u -dominated actions.

We extend Börgers' result by showing that a single vNM utility function can be used for all actions. While his result has the logical form: “for every action, there is a utility function such that...”; our result has the logical form: “there is a utility function such that, for every action...”. We extend Ledyard's result by showing that this is possible even if strong compatibility is imposed. Also, we establish equivalence of the entire dominance relations and not just the undominated sets, we provide tight, intuitive, and sufficient conditions on utility, and we show that dominance equivalence is attainable in some infinite environments.

After writing this paper, we encountered the recent work of Weinstein (2014) whose results complement our own. First, while we focus on extreme risk-aversion,

his results imply that, for CARA agents with extreme *risk-seeking* attitudes, every action which is not a best reply to a degenerate belief is M_u -dominated. Additionally, he finds that, in the context of a game, when agents are either extremely risk-averse or extremely risk-loving, all mixed equilibria become almost pure, in the sense that each player plays some pure action with probability arbitrarily close to 1.

Dominance relations are important for rationalizability as a solution concept for games (Bernheim, 1984, Pearce, 1984). Under standard assumptions, rationalizability is equivalent to iterated M_u -dominance. Börgers' result thus implies that, when only ordinal preferences are common knowledge, then rationalizability is equivalent to iterated P -dominance (Epstein, 1997, Bonanno, 2008).⁵ Our analysis implies that the equivalence extends to situations in which utility functions are common knowledge among the players, but only ordinal preferences are known to an outside observer. Furthermore, it also allows to relate observations arising from different situations, as in generalized revealed preference theory (Chambers et al., 2010).

1.1. Single-agent choice problems

We consider a single-agent environment characterized by (A, X, \succsim) . $X = \{x, y, \dots\}$ is a nonempty set of states of Nature, $A = \{a, b, \dots\}$ is a set of (pure) actions, and \succsim is a transitive and complete preference relation on $A \times X$. \succsim_x denotes state-contingent preferences over actions conditional on state x , i.e., $a \succsim_x b$ if and only if $(a, x) \succsim (b, x)$.

Let $[a]_x = \{b \in A \mid a \sim_x b\}$ denote the set of actions that are indifferent to a conditional on x . Throughout the paper we impose the following assumption, which essentially requires the quotient set A/\sim_x to be isomorphic to a subset of \mathbb{Z} , for every state x . While the assumption does limit the applicability of the results, it is satisfied by all finite environments, and it leaves sufficient space to accommodate many interesting infinite environments.

Assumption 1.1 The collection of equivalence classes $\{[c]_x \mid a \succ_x c \succ_x b\}$ is finite for every pair of actions a and b and every state x .

A vNM utility function $u \in \mathbb{R}^{A \times X}$ is *compatible* with the environment if it preserves state-contingent preferences, i.e., if $u(a, x) \geq u(b, x)$ if and only if $a \succsim_x b$. It is

⁵Lo (2000) extends this result to all models of preferences satisfying Savage's P3 axiom.

strongly compatible if it also preserves preferences across states, i.e., if $u(a, x) \geq u(b, y)$ if and only if $(a, x) \succcurlyeq (b, y)$. We extend the domain of utility functions to mixed actions $\alpha \in \Delta(A)$ and beliefs $\mu \in \Delta(X)$ in the usual way, and we denote payoff vectors associated to pure or mixed actions by $\vec{u}(\alpha) = (u(\alpha, x))_{x \in X}$.

Example 1.1 Going back to the motivating example from the introduction, let $X = \{1, 2\}$ and $E = \{1\}$, and let a_1 correspond to betting on E , a_2 to betting on $X \setminus E$, and a_0 to not betting. A vNM utility function u is strongly compatible if and only if it can be written as in Figure 1.1 after a positive affine transformation. Notice that this implies that $u(a_0, 1) = u(a_0, 2)$. In contrast, u is compatible as long as $u(a_x, x) < u(a_0, x) < u(a_y, x)$ for all $x, y \in X$ with $x \neq y$. Notice that this does not impose any restrictions on the differences $u(a, x) - u(b, y)$ when $x \neq y$. Figure 1.2 shows a strongly compatible vNM utility function (left panel), and a vNM utility function which is compatible but not strongly compatible (right panel).

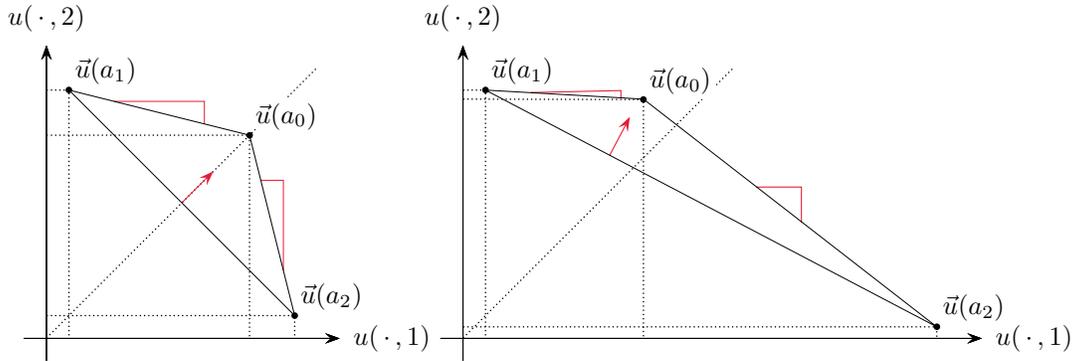


Figure 1.2 – vNM utility functions for Example 1.1 with $X = \{1, 2\}$.

1.2. Pure and mixed dominance

Loosely speaking, an action is dominated if there exist different actions yielding preferred outcomes regardless of the state. By dominance *by pure actions*, we mean the notion introduced by Börgers (1993), according to which an action is dominated if and only if it is weakly dominated conditional on each subset of states.

Definition 1.1 An action a is *P-dominated* in $B \subseteq A$, if for every nonempty set

$Y \subseteq X$ there exists some $b \in B$ such that $b \succ_y a$ for all $y \in Y$, with strict preference for at least one state $y \in Y$. $P(B)$ denotes the set of P -dominated actions in B .

P -dominance extends the classical notion of strict dominance by pure actions, and both notions coincide when preferences are asymmetric. An agent who maximizes expected utility would never choose P -dominated actions, even if they were not strictly dominated by pure actions. This is because, for an action to be a best response to some belief, it cannot be weakly dominated over the support of such beliefs. However, not being P -dominated is also not sufficient for being potentially optimal. It is well known that an action is potentially optimal if and only if it is not strictly dominated by a *pure or mixed action* according to the following definition.⁶

Definition 1.2 An action a is M_u -dominated in $B \subseteq A$ given a compatible vNM utility function u , if there exists a mixed action α such that $\text{supp}(\alpha) \subseteq B$ and $u(\alpha, x) > u(a, x)$, for every state $x \in X$. $M_u(B)$ denotes the set of M_u -dominated actions in B .

All P -dominated actions are also M_u -dominated relative to any compatible vNM utility function, but actions that are not P -dominated could still be M_u -dominated. We are interested in utility functions which guarantee that, if an action is M_u -dominated by a mixture α , then it is also P -dominated in the support of α . This requirement is equivalent to the following definition.

Definition 1.3 (Dominance equivalence) A compatible vNM utility function u generates *dominance equivalence* if $P(B) = M_u(B)$ for all $B \subseteq A$.

What conditions over vNM utility functions imply dominance equivalence? When does there exist a compatible or strongly compatible vNM utility function satisfying such conditions? The answer to these questions is closely related to risk aversion, as measured by the *timidity* coefficient introduced in the next section. Before proceeding, it is instructive to revisit our motivating example to illustrate the role of risk aversion.

Example 1.2 Consider once again the example from the introduction. Clearly, there are no P -dominated actions, and every action other than a_0 is optimal conditional on some state. Hence, dominance equivalence holds if and only if $a_0 \notin M_u(A)$, which holds if and only if $\vec{u}(a_0)$ is above the line containing $\vec{u}(a_1)$ and $\vec{u}(a_2)$, see Figure 1.2.

⁶This result can be traced back to Wald (1947).

This simply means that the upper boundary of the set of feasible payoffs is concave, or, equivalently, that $u(\cdot, x)$ exhibits *decreasing differences* (on average for compatibility, and always for strong compatibility). In environments with more states, dominance equivalence requires that finite differences should decrease sufficiently fast.

1.3. Timidity

Given a compatible vNM utility function u and a state x , we use the following notation. The set of possible payoffs given x is denoted by $U(x) = \{u(a, x) \mid a \in A\}$, and its supremum and infimum are denoted by $\bar{u}(x) = \sup U(x)$, and $\underline{u}(x) = \inf U(x)$. Also, let $u^-(a, x) = \sup\{u_0 \in U(x) \mid u_0 < u(a, x)\}$ denote the best possible payoff conditional on x which is worse than $u(a, x)$. Similarly, let $u^+(a, x) = \inf\{u_0 \in U(x) \mid u_0 > u(a, x)\}$. Assumption 1.1 implies that $u^-(a, x) < u(a, x)$ whenever $u(a, x) > \underline{u}(x)$, and $u^+(a, x) > u(a, x)$ whenever $u(a, x) < \bar{u}(x)$.

Definition 1.4 Given a compatible vNM utility function u and an action state pair (a, x) , the *timidity* coefficient of u at (a, x) is the number $\tau_u(a, x)$ given by $\tau_u(a, x) = +\infty$ if $u(a, x) \in \{\underline{u}(x), \bar{u}(x)\}$, and otherwise given by:

$$\tau_u(a, x) = \frac{u(a, x) - u^-(a, x)}{\bar{u}(x) - u(a, x)}. \quad (1.1)$$

In order to understand what timidity entails, it is useful to compare it with familiar measures of risk-aversion. Using finite differences instead of derivatives, the analogue of the Arrow-Pratt coefficient of absolute risk aversion for our discrete setting could be expressed as follows (see for instance [Bohner and Gelles \(2012\)](#)):

$$\rho_u(a, x) = 1 - \frac{u^+(a, x) - u(a, x)}{u(a, x) - u^-(a, x)}. \quad (1.2)$$

This coefficient is large when the *local* gain ($u^+(a, x) - u(a, x)$) is small compared with the local loss ($u(a, x) - u^-(a, x)$). In contrast, timidity compares the *global* gain ($\bar{u}(x) - u(a, x)$) with the local loss ($u(a, x) - u^-(a, x)$). Timidity requires that the potential loss of getting a slightly worse outcome should be more important than the potential gain of switching to the best possible outcome. A timid agent would refuse to spend a single dollar on a lottery ticket that promises to pay (with sufficiently low

probability) more money that she could spend during a hundred lifetimes.

Example 1.3 Suppose each action-state results in a monetary prize given by a function $z \in \mathbb{Z}_{++}^{A \times X}$ such that $\{z(a, x) \mid a \in A\} = \mathbb{N}$ for all $x \in X$. Further suppose that the agent's preferences only depend on her preferences over money, represented by $v \in \mathbb{R}^{\mathbb{R}_{++}}$. In this case, $u = v \circ z$ is a strongly compatible vNM utility function. For a CARA agent with $v(m) = -\exp(-rm)$, $r > 0$, the agent also exhibits constant timidity:

$$\tau_u(a, x) = \frac{-\exp(-r z(a, x)) + \exp(-r (z(a, x) - 1))}{\exp(-r z(a, x))} = \exp(r) - 1. \quad (1.3)$$

Before proceeding to the main results, we conclude our analysis of timidity by noting that it satisfies one of Pratt's classic criteria. The following proposition implies that, if an agent becomes uniformly more risk-averse as measured by ρ_u , then she also becomes uniformly more timid.

Proposition 1.1 *Fix an action a , a state x and two compatible vNM utility functions u and v . If the set of mixed actions that are preferred to a given u and x is contained in the set of mixed actions that are preferred to a given v and x , then u is more timid than v at (a, x) .*

1.4. Dominance equivalence and risk aversion

Let $W_x(a) = \{b \in A \mid a \succ_x b\}$ denote the set of actions that are worse than a conditional on x , and consider any compatible vNM utility function u . The following lemma states that if u is sufficiently timid at (a, x) , then a is not M_u -dominated by any mixed action α that assigns sufficient probability to $W_x(a)$. The rest of our results looks for conditions on u that guarantee that this can be done whenever a is not P -dominated. The conditions essentially require u to be sufficiently timid relative to the size of the environment.

Lemma 1.2 *Given a compatible vNM utility function u , a pure action a , and a mixed action α , if there exists a state x such that $(\tau_u(a, x) + 1) \cdot \alpha(W_x(a)) \geq 1$, then a is not dominated by α given u .*

1.4.1. Finite environments

Let $K = \min\{\|A\|, \|X\|\}$. When K is finite, Caratheodory's theorem (Rockafellar, 1996, Theorem 17.1) implies that an action is M_u -dominated if and only if it is dominated by a mixed action which mixes at most K distinct actions. For every such α , there exists some action a such that $\alpha(a) \geq 1/K$. Therefore, the condition of Lemma 1.2 holds for all P -undominated actions, whenever the timidity coefficient is weakly greater than $K - 1$.

Proposition 1.3 *Given a compatible utility function u , if $\tau_u(a, x) \geq K - 1$ for all x and a , then u generates dominance equivalence.*

Suppose that K is finite and $A \times X$ is countable. Then there exist strongly compatible vNM utility functions $n^* \in \mathbb{Z}^{A \times X}$ which only take integer values. For example, if $A \times X$ were finite, n^* could be the rank function defined by:

$$\text{rank}(a, x) = \left\| \left\{ (b, y) \mid (a, x) \succcurlyeq (b, y) \right\} \right\|. \quad (1.4)$$

In words, the rank of an action-state pair is the number of action-state pairs that are weakly worse than it. Let u^* be the utility function defined by:

$$u^*(a, x) = -\exp\left(-\log(K) n^*(a, x)\right). \quad (1.5)$$

If we thought of $n^*(a, x)$ as a monetary prize, then u^* would represent the preferences of a CARA agent with coefficient of risk aversion equal to $\log(K)$. Since u^* is a strictly monotone transformation of n^* , it is strongly compatible. Furthermore, we have that $u^*(a, x) \leq 0$ and $u^*(a, x) \geq K u^*(a, x)$ for all a and x , which implies that $\tau_{u^*} \geq K - 1$. Therefore, Proposition 1.3 implies the following corollary:

Corollary 1.4 *If either X or A is finite and $A \times X$ is countable, then u^* is a strongly compatible vNM function, and yields dominance equivalence.*

When both X and A are infinite, u^* is not well defined. The following example provides a countable environment which does not admit any strongly compatible vNM utility function generating dominance equivalence.

Example 1.4 Let $X = \mathbb{N}$, and suppose the agent must choose a lottery from $A = \{a_0\} \cup \{a_x \mid x \in X\}$. Lottery a_0 represents an outside option corresponding to keeping

her initial wealth. Lottery a_x represents a fair bet of one dollar *against* state x , i.e., it pays 1 if the true state is different from x and -1 otherwise. Further suppose that the agent has state-independent strictly monotone preferences over monetary holdings. After a positive affine normalization, any strongly compatible vNM utility function u can be written as:

$$u(a, x) = \begin{cases} \gamma & \text{if } a = a_0 \\ 0 & \text{if } a = a_x \\ 1 & \text{otherwise} \end{cases} . \quad (1.6)$$

for some $\gamma \in (0, 1)$. Take any such u , and consider any belief $\mu \in \Delta(\mathbb{N})$. For all $x \in \mathbb{N}$, we have that $u(a_0, \mu) = \gamma$ and $u(a_x, \mu) = (1 - \mu(x))$. If a_0 were a best response to μ , then it would be the case that $u(a_0, \mu) \geq u(a_x, \mu)$ and, consequently, $\mu(x) \geq \gamma > 0$ for all x . This would contradict the fact that μ is a probability measure. Hence, it follows that a_0 is M_u -dominated, despite the fact that it is not P -dominated.

1.4.2. Countable environments

When both X and A are infinite, guaranteeing dominance equivalence requires unbounded degrees of timidity. This is possible for countable environments if we do not require strong compatibility, because we may choose utility functions whose degree of timidity is always finite, but diverges to infinity along a sequence of states.

Proposition 1.5 *Given a compatible vNM utility function u , if X is countable and:*

$$\sum_{x \in X} \frac{1}{1 + \tau_u(a, x)} \leq 1, \quad (1.7)$$

for every action a , then u generates dominance equivalence.

The following example shows that the proposition is tight, in that, given any finite or countable X and a sufficiently large action space A , there always exist preferences such that: a compatible vNM utility function generates dominance equivalence *if and only if* it satisfies (1.7) for every action.

Example 1.5 Let X be any finite set with at least two elements, and let A and \succsim be

as in Example 1.4. Let u be any compatible vNM utility function such that:

$$\frac{1}{T} \equiv \sum_{x \in X} \frac{1}{1 + \tau_u(a_0, x)} > 1. \quad (1.8)$$

Simple algebra shows that a_0 is strictly dominated by the mixed action α given by $\alpha(a_0) = 0$ and $\alpha(a_x) = T/(1 + \tau_u(a_0, x))$ for $x \in X$.

If X is countable, then there exists an injective function $h \in \mathbb{N}^X$. Also, by Assumption 1.1, there exists a compatible vNM utility function $n^{**} \in \mathbb{Z}^{A \times X}$ which only takes integer values. Fix any such functions, and let u^{**} be the vNM utility function given by:

$$u^{**}(a, x) = -\exp\left(-h(x)n^{**}(a, x)\right). \quad (1.9)$$

Clearly, u^{**} is also compatible with the environment. Furthermore, we have that $u^{-**}(a, x) \geq e^{h(x)}u^{**}(a, x)$ and $u^{**} < 0$ for all a and x . Therefore:

$$\sum_{x \in X} \frac{1}{1 + \tau_{u^{**}}(a, x)} = \sum_{x \in X} \frac{\bar{u}^{**}(x) - u^{**}(a, x)}{\bar{u}^{**}(x) - u^{-**}(a, x)} < \sum_{k \in \mathbb{N}} \frac{1}{e^k} < 1. \quad (1.10)$$

Proposition 1.5 thus implies that:

Corollary 1.6 *If X is countable, then u^{**} is a compatible vNM function and yields dominance equivalence.*

1.4.3. Uncountable state space

Dominance equivalence is still possible in some environments with countable action spaces, and uncountable state spaces. Suppose that $(X, \mathcal{X}, \lambda)$ is a measure space, and let $Z(a, b) = \{x \in X \mid a \succ_x b\} \in \mathcal{X}$. We only require the two following assumptions.

Assumption 1.2 $(\exists \delta > 0)(\forall a, b \in A)(Z(a, b) \neq \emptyset \Rightarrow Z(a, b) \in \mathcal{X} \wedge \lambda(Z(a, b)) \geq \delta)$.

Assumption 1.3 There exists a measurable function $f : X \rightarrow (0, 1)$ such that $\int_X f d\lambda \leq 1$.

Assumption 1.2 requires that if an action a is preferred to an action b in at least one state, then it has to be preferred to it in a sufficiently large set. It plays a similar role as

our requirement that A should be discrete, but it is significantly weaker. Assumption 1.3 makes it possible to have timidity grow “sufficiently fast” as to guarantee that the condition of lemma 1.2 is satisfied at some point. It is satisfied, for instance, when $X \subseteq \mathbb{R}$ and λ is the Lebesgue measure.

Since A is assumed to be countable, there exists some compatible vNM utility function $n \in \mathbb{R}^{A \times x}$. Let $h^*(x) = -\log(\delta f(x))$, and define the vNM utility function u^* by:

$$u^*(a, x) = -\exp(-h^*(x) n(a, x)) \tag{1.11}$$

Proposition 1.7 *Under assumptions 1.2 and 1.3, u^* is a compatible vNM utility function which generates dominance equivalence.*

1.5. Summary and discussion

A vNM utility function guarantees that P -dominance and M_u -dominance coincide if it is sufficiently timid. For countable environments with discrete action spaces, it is always possible to find a sufficiently timid vNM utility function that is compatible with ordinal preferences over actions conditional on states. For finite environments with ordinal preferences over action-state pairs, it is always possible to find a sufficiently timid vNM utility function that is *strongly* compatible. In what follows, we discuss the application of the results to multi-agent environments, as well as some lines for further inquiry.

Rationalizability.— A strategic form game can be thought of as a collection of simultaneous single-agent decision problems. Rationalizability is then equivalent to the iterated removal of strategies that are not M_u -dominated. Our results then imply that, given any finite collection of finite games with ordinal payoffs, there exists a profile of compatible vNM utility functions such that, in each game, the set of rationalizable strategies corresponds to the set of strategies surviving the iterated removal of P -dominated strategies. In this sense, rationalizability and iterated P -dominance are equivalent in the absence of cardinal information.

Worst case vs. average bounds.— The degree of timidity assumed in our main results guarantees that dominance equivalence holds even in pathological scenarios with intri-

cate preferences. In particular, it guarantees that a P -undominated action that yields the second worst outcome conditional on every state is potentially optimal, even if there are other actions yielding very good outcomes in all but one state.

An interesting problem not solved in this paper is to look for *expected* (rather worst-case scenario) bounds on timidity. For instance, one could ask for the probability that a uniformly generated utility function u will generate equivalence. One step further, having fixed only the size of the environment, one could ask for the expectation of this probability given uniformly generated preferences.

Uncountable environments.— Our proofs depend crucially on Assumption 1.1. For practical purposes, this limits the scope of our results to environments with countable action spaces. This is because the definition of timidity heavily relies on the fact that there exists some $\delta > 0$ such that $|u(a, x) - u(a', x)| \geq \delta$ whenever $u(a, x) \neq u(a', x)$.

Chapter 2 |

Interdependent choices

Abstract I propose Interdependent-choice Rationalizability (ICR) and Interdependent-choice Equilibrium (ICE) as simple and tractable solution concepts that allow for interdependent choices. They can be interpreted in two ways: (1) as robust predictions that do not depend on specific assumptions about the details of the sequential and informational structure, or (2) as a characterization of self-enforceable contracts in environments in which the timing of actions is flexible, and actions are verifiable. I show that ICE enables cooperation in some but not all prisoners' dilemmas, and I use this fact to find the optimal deal that a district attorney should offer the prisoners.

When agents make choices independently and simultaneously, standard notions of rationality postulate that each agent chooses a myopic best response to a fixed belief about his opponents' behavior. The story is quite different when the decisions of some agents may depend on the actual behavior of others. In such settings, it is not sufficient for agents to consider the material consequences of their acts taking the behavior of their opponents as given. Agents must also consider the way that their own choices might affect those of others. A particularly important form of reaction or counterfactual reasoning is reciprocity: each agent may believe that others will be good to him *if and only if* he is good to others.

Choice interdependence is particularly relevant in moral hazard environments with no Pareto efficient Nash equilibria. With complete information, the problem of moral hazard can be solved if agents can enforce complete contracts (Coase theorem), or if they interact repeatedly and are patient enough (folk theorems). This is possible

because written contracts or publicly observed histories serve as coordination devices allowing interdependence. This chapter abstracts the notion of choice interdependence from such settings to investigate the extent to which its power remains in environments without *binding commitment*, *repetition*, or *monetary transfers*.

I consider situations in which each agent has to choose and perform an *irreversible action* at a *time of his choosing*, and such actions are *verifiable* but *not contractible*. Examples of such environment can be, citizens casting a vote in an election, prisoner choosing whether to accept or reject a sentence reduction in exchange for a confession, or competing firms choosing which features to include in their product. Suppose that a researcher knows the actions available to each agent and the agents' preferences over act profiles, but does *not* know any additional details about the game being played. What kind of predictions could she make about the outcome of the environment? Could anything other than a correlated equilibria arise in equilibrium? Will a folk-theorem-like result apply?

I propose two tractable and simple solution concepts to help address these questions. *Interdependent-choice Rationalizability* (ICR) is a form of rationalizability that allows for agents to believe that the choices of others depends on their own. *Interdependent-choice equilibrium* (ICE) with respect to a set of credible threats, is a simple notion of equilibrium defined by a finite set of affine inequalities. Define an *extensive form mechanism* (EFM) to be an extensive form game that is compatible with the information the researcher has about the environment. Proposition 2.3 states that, if an outcome is implementable as a Nash equilibrium of an EFM, then it must be an ICE. Proposition 2.4 states that, if an outcome is an ICE with respect to the set of ICR actions, then it is implementable as a PB equilibrium of an EFM. Proposition 2.5 states that, in generic 2×2 environments, an outcome can arise as a sequential equilibrium of an EFM if and only if it is an ICE with respect to the set of ICR actions.

Even without contracts, ICE can go well beyond the set of correlated equilibria. However, it yields sharper predictions than standard models with full commitment. The set of interdependent belief systems that are consistent with an EFM are restrictive enough to rule out a folk-theorem-like result. In particular, in sections 2.1 and 2.4, I show that cooperation possible in some but not all prisoners' dilemmas. In particular, Proposition 2.6 characterize exactly how and when cooperation is possible in a prisoners' dilemma. I use this characterization to find the optimal deal that a district attorney should offer to the prisoners (Proposition 2.8).

There are many different EFMs that implement the same outcome as an equilibrium. In section 2.5, I introduce a class of mechanisms called *mediated mechanisms* that are canonical in the sense of Forges (1986). That is, every outcome that can be implemented in some EFM can *also* be implemented in a mediated mechanism. In a mediated game, a non-strategic mediator manages the play through private recommendations. The salient features are that the mediator can choose the timing of the recommendations, and make her recommendations contingent on the behavior of the agents.¹ For these mechanisms to be feasible, it is important that the timing of the actions is flexible, and that actions can be directly observed or verified by an impartial third party. The set of mediated mechanisms is not the only canonical class of EFM. However, it is a natural choice in accordance with the principle that implementation is easier when the mediator can observe everything, while the agents have as little information as possible (Myerson, 1986).

Being an ICE with respect to ICR actions is a sufficient but not a necessary condition to be PB implementable. This is because, after observing an unexpected event in an extensive form game, a player might believe that the past choices of his opponent were not rational. Hence, despite the fact that ICR is equivalent to common knowledge of rationality with interdependent beliefs (cf. Halpern and Pass (2012)), it is possible for rational players to choose off-the-equilibrium-path actions which are not ICR. To deal with this issue, section 2.6 introduces the notions of *forward-looking interdependent-choice rationalizability* (FICR), and *quasi-sequential equilibrium* (QSE). QSE is a notion of equilibrium for extensive form games that lies in-between PB equilibrium and sequential equilibrium. Proposition 2.10 asserts that an outcome is QS implementable if and only if it is an ICE with respect to FICR actions.

Comparison with the literature

Choice interdependence is a common theme across a variety literatures. Besides the well established literatures on repeated games and games with contracts, different literatures allow for different forms of implicit repetition, commitment or transfers. The literature on counterfactual variations can be thought of as a reduced representa-

¹These mediators are remarkably powerful in comparison to the ones in classic studies that only allow mediators to engage simultaneous pre-play recommendations such as Aumann (1987), or simultaneous communication at fixed stages of the game such as Myerson (1986). At the same time, they are less powerful than mediators who can take actions on behalf of the players as in Moulin and Vial (1978), Monderer and Tennenholtz (2009) and Forgó (2010).

tion of repeated games (Kalai and Stanford, 1985). Commitment can be traced back to Moulin and Vial (1978) and Kalai (1981), which allow players to delegate choices to a mediator, or make *binding* preplay announcements. With unrestricted commitment, one obtains folk theorems (Kalai et al., 2010). Recent relevant works in this area include papers on commitment (Bade et al., 2009, Renou, 2009) and delegation (Forgó, 2010). Also related is the recent literature on revision games (Kamada and Kandori, 2009, 2011), which allow for a specific form of pre-play binding communication, see also Calcagno et al. (2014) and Iijima and Kasahara (2015).

Other literatures allow for counterfactual reasoning without always being explicit about the source of interdependence. Seminal examples include Rapoport (1965) and Howard (1971). More recently, Halpern and Rong (2010) and Halpern and Pass (2012), analyze equilibrium and rationalizability with counterfactual beliefs. Once again, folk theorems hold if no further restrictions are imposed. There is a question as to when and which forms of counterfactual reasoning are consistent with the idea that players have independent free wills (Gibbard and Harper, 1980, Lewis, 1979). The present work thus considers only those counterfactual beliefs which can arise from different sequential and informational structures of an EFM.

In this regard, the current work is closely related to (Nishihara, 1997, 1999). Nishihara proposes a mechanism that allows for cooperation in some prisoners' dilemma games. An important difference is that my mechanism is timeable in the sense of Jakobsen et al. (2016). See Section 2.5.1 for more details. Also, I go beyond cooperation in the prisoners' dilemma and characterizes every outcome which can be implemented in any finite game.

Interdependent choices have been studied in other specific contexts. Eisert et al. (1999) shows that cooperation is possible in a prisoners dilemma, when players can condition their choices on certain *quantum* randomization devices with *entangled* states. Tennenholtz (2004)'s program equilibrium generates choice interdependence for games *between computer programs* by allowing them to read each other's code before taking an action. Similarly, Levine and Pesendorfer (2007)'s self-referential equilibrium allows player's to receive a signal about their opponent's intentions before making their own choice. See also Block and Levine (2015), Block (2013). I rule out this kind of signals as they represent a form of commitment: from the moment of *deciding* which action to play, to the moment of actually *performing* it. ICE can arise in settings where choices are instantaneous, and players can hide their intentions.

2.1. Cooperation in a prisoners' dilemma

Consider the following variation of the prisoners' dilemma. Two prisoners awaiting trial are offered a sentence reduction in exchange for a confession. As usual, this is a one-shot interaction and there are no contracts nor monetary transfers. However, I do not assume that the prisoners must make their choices independently nor simultaneously. Instead, they can choose the timing of their actions and hire a non-strategic lawyer (she) to help them coordinate. Formally, suppose the following:

- (i) Each prisoner $i = 1, 2$, must submit an official signed statement at a time of his choosing $t_i \in [0, 1]$.
- (ii) The statement specifies whether the prisoner chooses to cooperate with his accomplice (C), or defect by accepting the deal (D). Once submitted, the statement cannot be modified or withdrawn.
- (iii) The lawyer can make private non-binding recommendations to each prisoner at any moment in time.
- (iv) Each prisoner can show the lawyer a certified copy of his submitted statement as *hard proof* that he cooperated or defected.
- (v) The prisoners' preferences are summarized in the following payoff matrix, where $g, l > 0$ are fixed parameters.

	C	D
C	1, 1	$-l, 1 + g$
D	$1 + g, -l$	0, 0

Figure 2.1 – Payoff matrix for the prisoners' dilemma.

If $g < 1$, it is possible for the prisoners to cooperate in equilibrium. They could instruct the lawyer to proceed as follows. She will randomly and privately choose two times $r_i \in [0, 1]$. At date r_i , she will recommend prisoner i to immediately submit his statement with some recommended action. In equilibrium, both prisoners will follow such recommendations and immediately report back showing the copy of their statements as proof of compliance. The lawyer will always recommend C along the equilibrium path, and D after any detectable deviation.

The recommendation dates should be drawn as $r_i = 1 - 1/n_i$ where $n_1, n_2 \in \mathbb{N}$ are chosen as follows. First, the lawyer will draw $n \in \mathbb{N}$ from a geometric distribution with parameter $\rho \in (0, 1 - g)$. This is possible only because of the assumption $g < 1$.² Then, with probability $1/2$, she sets $n_1 = n$ and $n_2 = n + 1$, and with probability $1/2$ she sets $n_2 = n$ and $n_1 = n + 1$.

I will now show that following the lawyer's recommendations constitutes a subgame-perfect equilibrium of the induced game. There are different types of histories to consider. First, suppose that prisoner i is recommended to cooperate at time $t_i = 0$. In this case, i knows for sure that he is the first player to receive a recommendation. Hence, if he cooperates and shows evidence of this to the lawyer, then $-i$ will also cooperate and his payoff will be $u_i(C, C) = 1$. Otherwise, if he deviates by defecting, waiting, or refusing to show evidence of his cooperation to the lawyer, then his accomplice will defect and his payoff will be no better than $u_i(D, D) = 0$. Hence, in this case, it is optimal for i to be obedient and cooperate.

Now suppose that i is recommended to cooperate at some time $t_i = 1 - 1/n_i > 0$. As before, if he cooperates and shows evidence of this to the lawyer, then $-i$ will also cooperate and his payoff will be $u_i(C, C) = 1$. If he deviates by defecting, his payoff will equal $(1 + g)$ times the probability that player $-i$ has already cooperated and won't have an opportunity to react to i 's deviation. This probability is given by

$$\Pr(t_i > t_{-i} | t_i) = \frac{1/2 \cdot \Pr(n = n_i - 1)}{1/2 \cdot \Pr(n = n_i) + 1/2 \cdot \Pr(n = n_i - 1)} = \frac{1}{2 - \rho}.$$

Hence, i 's expected utility from defecting is bounded above by

$$\frac{1}{2 - \rho}(1 + g) < \frac{1}{2 - (1 - g)}(1 - g) = 1,$$

which implies that cooperating is optimal.

Now consider a history when prisoner i is recommended to defect. Because this recommendation only arises after a detectable deviation, i can believe that his accomplice has already confessed. And, therefore, it is optimal for i to do the same. Finally, it is easy to see that there is no reason why a prisoner would like to submit his statement *before* being instructed to do so by the lawyer. Hence, always following the lawyer's recommendations constitutes a sequential equilibrium that results in full cooperation.

²The condition $g \leq 1$ is both sufficient and necessary to be able to implement cooperation as an ICE. See Corollary 2.7 in section 2.4.

2.2. A framework for robust predictions

Each player $i \in I = \{1, 2\}$ is to choose and perform one and only one irreversible action a_i from a finite set A_i .³ The restriction to two players is for exposition purposes. The model can be extended to n -player environments, but the required notation is cumbersome. i 's preferences over action profiles are represented by $u_i : A \rightarrow \mathbb{R}$. No form of contract or binding agreement with regard to these actions can be enforced. Each agent must freely choose which action to take.

2.2.1. Extensive form mechanisms

The tuple $E = (I, A, u)$ is only a *partial* characterization of the environment. It says nothing about the order in which choices will be made, nor about the information that each player will have when deciding which action to take.⁴ In particular, choices need not be independent nor simultaneous. Instead, players could be play some extensive form game that generates choice interdependence. For instance, players could condition their choices on correlated random signals. Alternatively, it could be the case that players take their actions sequentially in a fixed order, so that the decisions the later movers depends on the choices of those who moved first. I assume that the agents could be playing *any* extensive form game that is consistent with our partial description of the environment, and with the no-commitment assumption.

Extensive form games are defined as in Osborne and Rubinstein (1994), with some differences in notation. An extensive form game is a tuple $G = (M, X, j, \mathcal{H}, s_0, v)$. M is the *finite* set of moves and $X \subseteq_{t \in \mathbb{N}} M^t$ is the *countable* set of nodes. $j(x) \in I \cup \{0\}$ is the agent moving at x , where 0 represents Nature (or a mediator). Z and Y_i are the sets of terminal nodes and i 's decision nodes, respectively. \mathcal{H}_i partitions Y_i into information sets and satisfies *perfect recall*. Let $M(x)$ and $M(H)$ denote the sets of moves available at a node $x \in X$ or an information set $H \in \mathcal{H}$, respectively. $s_0 : Y_0 \rightarrow \Delta(M)$ specifies the players' common prior beliefs about Nature's behavior. Finally, $v_i : Z \rightarrow \mathbb{R}$ represents i 's preferences over terminal nodes.

³I employ the notation $-i$ for i 's opponent, $a = (a_i, a_{-i}) \in A = \times_i A_i$ for action profiles, $\alpha \in \Delta(A)$ for joint distributions, $\alpha_i \in \Delta(A_i)$ for marginal distributions, and $\alpha(\cdot | a_i) \in \Delta(A_{-i})$ for conditional distributions. Also, let $\mathcal{A} = \{\times_{i \in I} A'_i \mid A'_i \subseteq A_i\}$ denote the set of action subspaces.

⁴ Here, the term "information" refers to information about the actions taken by other players. The environment E is assumed to be common knowledge. In contrast, Bergemann and Morris (2013) derive robust predictions to assumptions on the information that the players have about the payoff structure.

The following definition captures the set of extensive form games that are consistent with our partial description of the environment, while ruling out side-payments, enforceable contracts, delegation, and any other form of *binding* commitment regarding their choices about actions a_i .

Definition 2.1 An *extensive form mechanism* (EFM) for E is an extensive form game G such that for every terminal node z there exists an action profile a^z and a tuple of decision nodes $(x_i^z)_{i \in I} \in \times_{i \in I} Y_i$ such that

- (i) $v(z) = u(a^z)$.
- (ii) For each player i , $A_i \subseteq M_i(x_i^z)$ and both x_i^z and (x_i^z, a_i^z) are predecessors of z .
- (iii) For each player i , decision node $x'_i \in Y_i$, and action $a'_i \in A_i$, if both x'_i and (x'_i, a'_i) are predecessors of z , then $x'_i = x_i^z$ and $a'_i = a_i^z$.

The first requirement is that each terminal node z can be associated with an action profile a^z , and utility over terminal nodes is given by the utility from the corresponding action profiles. The second requirement is that each player i actually chooses his own action a_i^z . Also, it rules out partial commitment by requiring that, at the moment of choosing a_i^z , player i could have chosen *any* other action in A_i . The third requirement guarantees that choices are irreversible. Also, it implies that for every every action profile $a \in A$, there exists some terminal node z such that $a^z = a$.

For exposition purposes, Definition 2.1 is more restrictive than necessary. In particular, it requires that there is at most a single way to perform each action in each information set. This precludes features that might seem relevant, such as allowing players to decide on the spot whether to take an action publicly or privately. Appendix B.1 proposes a (slightly) more general definition of extensive form mechanisms. As it turns out, all the results of the paper remain true with either definition.

2.2.2. Robust predictions

Given an extensive form mechanism, a (behavior) strategy for player i is a function that assigns a distribution $s_i(\cdot | H) \in \Delta(M(H))$ to each information set $H \in \mathcal{H}_i$. Let S_i denote the set of i 's strategies, and $S = \times_i S_i$. Let $\zeta(x|s, s_0)$ denote the probability that the game will reach node x given that Nature chooses according to s_0 and players follow s . Each strategy profile s induces a distribution $\alpha^s \in \Delta(A)$ over acts of the environment given by $\alpha^s(a) = \sum_{z \in Z(a)} \zeta(z|s, s_0)$, where $Z(a) = \{z \in Z \mid a^z = a\}$.

Definition 2.2 A distribution over acts $\alpha \in \Delta(A)$ is (Nash, sequentially, PB, ...) *implementable*, if there exists an extensive and a corresponding (Nash, sequential, PB, ...) equilibrium s^* that induces it.

The set of implementable outcomes is exactly the set of robust predictions that a researcher could make if he knew the partial description of the environment E , and he believed that the agents could be playing any equilibrium of any EFM. Alternatively, it can be interpreted as the set of outcomes that the players could implement as self-enforceable agreements without the use of binding commitment devices. The main objective of the paper is to characterize the sets of implementable outcomes under different notions of equilibrium and rationality.

2.3. Interdependent choice equilibrium and robust predictions

2.3.1. Interdependent choice equilibrium

This section characterizes the set of outcomes that are implementable in the sense of Definition 2.2. This is done via the notion of interdependent choice equilibrium defined as follows.

Definition 2.3 A distribution over action profiles $\alpha \in \Delta(A)$ is an *interdependent-choice equilibrium* (ICE) with respect to $B \in \mathcal{A}$, if there exists some $\theta : A \rightarrow \Delta(I)$ such that for all $i \in I$ and $a' \in A_i$

$$\sum_{a_{-i} \in A_{-i}} \alpha(a) \left(u_i(a) - (1 - \theta(i|a)) u_i(a'_i, a_{-i}) - \theta(i|a) \underline{w}_i(a'_i | \alpha, B) \right) \geq 0, \quad (2.1)$$

where $\underline{w}_i(a'_i | B_{-i}^*) := \min \{ u_i(a'_i, a_{-i}) \mid a_{-i} \in B_{-i}^* \}$ and $B_{-i}^* = \text{supp}(\alpha_{-i}) \cup B_{-i}$. For the case $B = A$, I omit the reference to B and simply say that α is an ICE.

Some intuition about the way ICE is defined can be acquired by analyzing two extreme cases. Condition (2.1) can be thought of as an incentive constraint for player i , requiring that a_i should be more profitable than deviating to a'_i . If $\theta(i|a) = 0$, this condition coincides with the incentive constraints for correlated equilibrium. In this case, player i computes the expected utilities from a_i and a'_i using the same

distribution $\alpha_{-i}(\cdot | a_i)$. In other words, i believes that the behavior of $-i$ does not depend on his own choice. In the opposite extreme, if $\theta(i|a) = 1$, player i believes that if he deviates from a_i to a'_i , his opponent will react by choosing the harshest punishment in B_{-i}^* .⁵ In this case, if $B^* = A$, condition (2.1) coincides with the definition of individual rationality.

The condition that $\theta(\cdot | a)$ should be a probability measure captures the fact that choice interdependence in my model should be consistent with an equilibrium of some EFM.⁶ For that to be the case, $-i$'s action can *only* depend on i 's action if $-i$ makes his choice *after* observing a signal about $-i$'s action. Of course, it cannot be the case that i moves before $-i$ and, at the same time, $-i$ moves before i . One can think of $\theta(i|a)$ to be the probability that i is the first player to move, conditional on being a path of play in which the chosen action profile is a .

Proposition 2.1 (ICE properties) *For any $B \in \mathcal{A}$, the set of ICE with respect to B is non-empty and contains the set of correlated equilibria and is contained in the set of individually rational outcomes. The set of ICE with respect to A is a non-empty closed and convex polytope.*

Example 2.1 Two partners decide whether to work (W) or shirk (S) in a joint-venture, their payoffs are depicted in Figure 2.2. The figure also shows the sets of payoffs corresponding to individual rationality, Nash equilibrium with public randomization, correlated equilibrium, and ICE. In this example, all the Pareto efficient outcomes correspond to ICE.

(W, W) is not a Nash equilibrium because, whenever an agent is working, his opponent prefers to shirk. It is an ICE because, a player who considers shirking knows that with some probability, his opponent will learn of this defection and react by also shirking. The payoff vector (1, 1) cannot be attained as an ICE, because it requires players to shirk with high probability. Since each player always prefers that his opponent works, this leaves too little room to punish deviations.

⁵The definition of ICE does not include any incentive constraints for such punishments, *as if* agents could commit to punish any deviations. This is *not* an assumption, but rather a technical tool to keep the model tractable. As it turns out, ICE can be used to characterize implementation under different notions of perfection, by choosing B adequately.

⁶This is the point where my model differs from models with full commitment, such as Kalai et al. (2010), Halpern and Pass (2012), Tennenholtz (2004) or Block and Levine (2015). This restriction is the reason why the set ICE does not result in a folk-theorem-like result. In particular, setting $\theta(i|a) = \delta \in (0, 1)$ results in a condition similar to the recursive characterization of SPNE of repeated games due to Abreu et al. (1990). With that analogy in mind, requiring $\theta(\cdot | a)$ to be a probability measure is tantamount to assuming that the average discount factor cannot exceed 1/2.

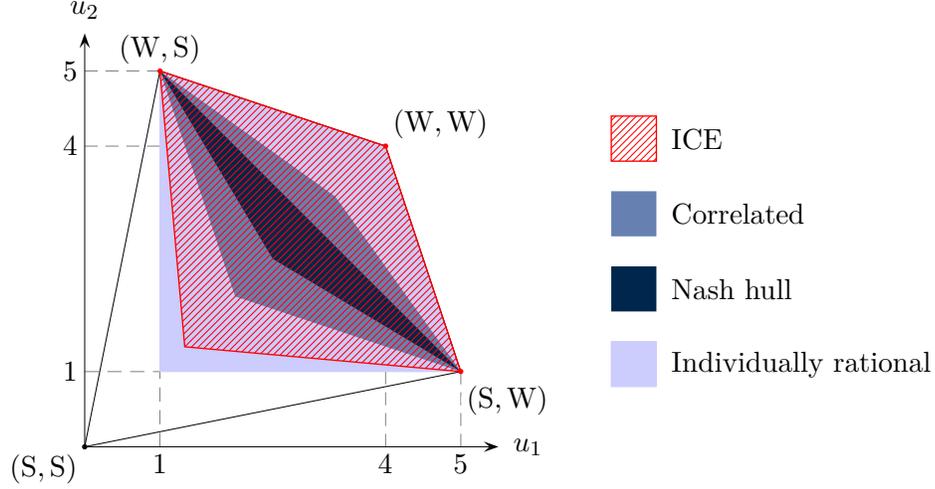


Figure 2.2 – Equilibrium payoffs for a teamwork game.

2.3.2. Interdependent choice rationalizability

The salient feature of environments with choice interdependence, is that rational agents do not make their choices given a fixed belief. Instead, they compute expected utility with respect to *counterfactual beliefs* $\lambda_i : A_i \rightarrow \Delta(A_{-i})$, where $[\lambda_i(a_i)](a_{-i})$ represents i 's assessed likelihood that his opponent will choose a_{-i} if i plays a_i . This section proposes a notion of rationalizability that incorporate this idea.

Before doing so, I will introduce some notation. Expected utility with respect to counterfactual beliefs is denote by $U_i(a_i, \lambda_i) = \sum_{a_{-i} \in A_{-i}} u_i(a_i, a_{-i}) \cdot [\lambda_i(a_i)](a_{-i})$. An action a_i is said to be a best response to λ_i if $U_i(a_i, \lambda_i) \geq U_i(a'_i, \lambda_i)$ for all $a'_i \in A_i$. For $A' \in \mathcal{A}$, $\Lambda_i(A')$ is the set of counterfactual beliefs such that $[\lambda_i(a_i)](A'_{-i}) = 1$ for every $a_i \in A_i$.

Definition 2.4 (Interdependent choice rationalizability (ICR))

- $a_i^* \in A_i$ is *IC-rationalizable* (ICR) with respect to $A' \in \mathcal{A}$, if and only if it is a best response to some $\lambda_i \in \Lambda_i(A')$.
- $A' \in \mathcal{A}$ is *self-IC-rationalizable* if and only if every action profile in A' consists of actions that are ICR with respect to A' .
- The set of *ICR action profiles*, A^{ICR} , is the largest self-IC-rationalizable set.

ICR is analogous to standard definitions of (correlated) rationalizability (cf. Pearce (1984), Bernheim (1984)), simply replacing correlated beliefs with counterfactual beliefs. It was developed independently and simultaneously by Halpern and Pass (2012),

who additionally show that it is equivalent to rationality and common certainty of rationality in epistemic models with counterfactual reasoning.

Let $\text{ICR}_i(A')$ denote the set of i 's actions that are ICR with respect to A' . A^{ICR} is guaranteed to exist because $\text{ICR}(\cdot)$ is \subseteq -monotone. Consequently, the union of all self self-IC-rationalizable sets is also self-IC-rationalizable. It is nonempty because it always contains the set of rationalizable action profiles. A^{ICR} can be found in a tractable way using the notion of absolute dominance defined ahead.

Definition 2.5 Given two actions $a_i, a'_i \in A_i$, a_i *absolutely dominates* a'_i with respect to $A' \in \mathcal{A}$ if and only if $\min_{a_{-i} \in A'_{-i}} u_i(a_i, a_{-i}) > \max_{a_{-i} \in A'_{-i}} u_i(a'_i, a_{-i})$.

In other words, a_i absolutely dominates a'_i , if and only if the best possible payoff from playing a'_i is strictly worse than the worst possible payoff from playing a_i . Absolute dominance is much simpler than strict dominance in computational terms because a player can conjecture different reactions for each alternative action, and thus mixed actions need not be considered. The following proposition ensures that $\text{ICR}(A')$ results from eliminating absolutely dominated actions, and A^{ICR} results from repeating this process iteratively.

Proposition 2.2 *An action is ICR with respect to A' if and only if it is not absolutely dominated in A' , and the iterated removal of absolutely dominated actions is order independent and converges in finite time to A^{ICR} .*

2.3.3. Robust predictions under choice interdependence

So far, two types of solution concepts have been introduced. First, the set of implementable outcomes is the set of robust predictions that a researcher can make that do not depend on the details of the environment. Second, ICE and ICR are tractable solution concepts that are easy to compute. this section shows that ICE and ICR can help to characterize the sets of implementable outcomes under different standard notions of equilibrium for extensive form games. The first result is that, every outcome that can arise as an equilibrium of an EFM is an ICE.

Proposition 2.3 *If a distribution over action profiles is Nash implementable then if it is an ICE.*

Proposition 2.3 implies that an outcome which is *not* an ICE requires some form of commitment, side payments, or repetition to be implemented in equilibrium. It suggests that, in situations in which (i) action spaces and preferences are known, (ii) agents cannot use binding commitment devices regarding their choices from A_i , and (iii) agent's choices are expected to be in equilibrium, it is safe to say that the outcome will be an ICE even if the details of the timing and information structures are unknown.

The converse of Proposition 2.3 is also true. Every ICE is Nash implementable. However, because EFM are extensive form games, plausible equilibrium predictions should involve sequential rationality constraints. In particular, the definition of ICE does not involve incentive constraints for additional off-the-equilibrium-path punishments in $B^* \setminus \{\text{supp}(\alpha)\}$. The following proposition provides sufficient conditions for implementation as perfect Bayesian (PB) equilibrium.

Proposition 2.4 *Every ICE with respect to ICR is PB implementable.*

The condition in Proposition 2.4 is sufficient but *not* necessary for PB implementation in general games. Section 2.6.2 explains it is not necessary, and provides condition that is both sufficient and necessary for general games. As it turns out, ICE with respect to ICR is necessary and “almost” sufficient for sequential implementation in generic 2×2 environments.

Proposition 2.5 *In 2×2 environments without repeated payoffs, the set of ICE with respect to ICR is dense in the set of sequentially implementable outcomes.*

2.4. When is cooperation possible in a prisoners' dilemma

Some versions of the prisoners' dilemma story fit my model closely. A signed statement is an irreversible action with flexible timing. Contracts between the prisoners are not enforceable by law. The right to have an attorney present provides a mechanism for actions to be verifiable. So, as long as repetition or transfers can be ruled out, the set of sequentially implementable outcomes provides a reasonable set of robust predictions for behavior in a prisoners' dilemma with rational prisoners. Applying proposition 2.5 results in the following complete characterization.

Proposition 2.6 *In the prisoners' dilemma from Figure 2.1, a joint distribution $\alpha \in \Delta(A)$ is sequentially implementable if and only if:*

$$\left(\frac{1-g}{l}\right) \alpha(C, C) \geq \alpha(D, C) + \alpha(C, D). \quad (2.2)$$

Let us further analyze condition (2.2). First, notice that there are no constraints on $\alpha(D, D)$, this is because defecting is already a dominant strategy. Second, notice that the right hand side is always non-negative, and can always be made equal to 0 by setting $\alpha(D, C) = \alpha(C, D) = 0$. This implies that the cooperative strategy profile (C, C) can be played with positive probability only if $g \leq 1$, and, if $g \leq 1$, then it can be played with full probability.

Corollary 2.7 *If $g \leq 1$, then (C, C) is an ICE and, if $g > 1$, the only ICE is (D, D) .*

Finally, notice that the probabilities of the asymmetric outcomes are bounded above by a linear function of the probability of the cooperative outcome. This is a natural condition, because a player can only benefit from cooperating if his opponent is also cooperating. In particular, when $g = 1$, the left hand side of condition (2.2) equals 0, which means that players cannot assign any probability to asymmetric outcomes. For $g \in (0, 1)$, this bound is relaxed so that the set of ICE payoffs fans out, see figure 2.3. However, the expected utility of players who do not confess has to be strictly greater than 1, because otherwise the threat of opponent's defection would not have any bite. This implies that there always exist individually rational payoffs which cannot be achieved by any Nash implementable distribution.

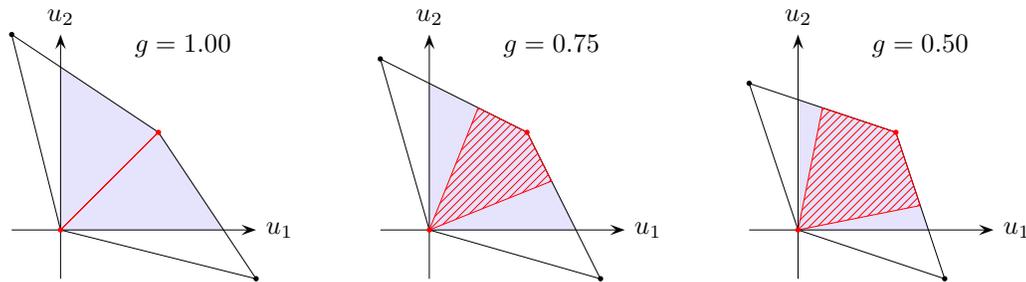


Figure 2.3 – ICE for the prisoners' dilemma with $l = 0.5$ and different values of g .

2.4.1. Optimal sentence reduction

The analysis can be taken one step forward. Consider the problem of a DA that must choose which deal to offer the prisoners as to maximize the total amount of time served. Proposition 2.6 shows that cooperation may or may not be possible depending on the specific payoffs. Hence, the DA's problem is non trivial as she must figure out the worst offer that the prisoners will accept in equilibrium. For simplicity, I assume that the prisoners have linear preferences over the amount of time served.

Suppose that, if n prisoners were to confess, the DA could secure a maximum sentence of at most $\bar{\mu}_n \in \mathbb{N}$ days in prison for each prisoner, where $\bar{\mu}_2 > \bar{\mu}_1 > \bar{\mu}_0 > 1$ are exogenous parameters.⁷ The DA can commit to offering an anonymous sentencing policy $\mu = (\mu_0, \mu_1^-, \mu_1^+, \mu_2) \in \mathbb{Z}_+^4$. If nobody confesses, each prisoner will be sentenced to μ_0 days. If both prisoners confess, each will be sentenced to μ_2 days. If only one prisoner confesses, he will be sentenced to μ_1^- days and his accomplice to μ_1^+ days. A policy μ is feasible if $\mu_0 \leq \bar{\mu}_0$, $\mu_1^+, \mu_1^- \leq \bar{\mu}_1$ and $\mu_2 \leq \bar{\mu}_2$.

The timing is as follows. First, the DA chooses a feasible policy μ . This induces the environment E in figure 2.4. Then, the prisoners play that symmetric equilibrium which minimizes their total time served. Here, I consider two scenarios. In the first scenario, the DA can force agents to make choices independently, so that the set of equilibria is just the set of NE of the environment. In the second scenario, the DA cannot prevent the prisoners from coordinating their choices, so that the relevant set of equilibria is the set of ICE.

	C	D
C	$-\mu_0, -\mu_0$	$-\mu_1^-, -\mu_1^+$
D	$-\mu_1^+, -\mu_1^-$	$-\mu_2, -\mu_2$

Figure 2.4 – Environment induced by a feasible policy μ .

Proposition 2.8 *The maximum total sentence time that the DA can guarantee under independent choices is $2\bar{\mu}_1 - 2$. Under interdependent choices, the time that the DA can guarantee is $2 \min\{2\bar{\mu}_0, \bar{\mu}_1\} - 2$. Hence, the cost to the DA of choice interdependence is $2 \max\{0, \bar{\mu}_1 - 2\bar{\mu}_0\}$.*

⁷I restrict attention to the natural numbers in order to guarantee existence of an optimal policy.

Under independent choices, the optimal mechanism is to set the maximum possible sentence for those people who do not confess, and the minimal sentence reduction in exchange for a confession. This results in a *per capita* sentence close to $\bar{\mu}_1$. The resulting game is a prisoners' dilemma. Using the same mechanism, the condition $g > 1$ from Corollary 2.7 is satisfied if $\bar{\mu}_1 \leq 2\bar{\mu}_0$. Which means that mutual defection is also the only ICE and the prisoners would accept the deal, even under choice interdependence. Hence, in that case, there are no interdependence rents for the agents.

In contrast, when $\bar{\mu}_1 > 2\bar{\mu}_0$, using the same mechanism would not be optimal because mutual cooperation would be an ICE that would result in a total sentencing time per capita of at most $\bar{\mu}_0$. In appendix B.2.4 I show that, in that case, the optimal mechanism is given by $\mu_0^* = \bar{\mu}$, $\mu_1^+ = 0$, $\mu_1^- = \bar{\mu}_1$, and $\mu_2^* = \min\{2\bar{\mu}_0, \bar{\mu}_1\} - 1$. This is the harshest prisoners' dilemma in which mutual cooperation is *not* an ICE. Interestingly, the maximum possible per capita sentence in that case is only twice the amount of time that the DA could convict the prisoners for without a confession.

2.5. Implementation via mediated mechanisms

There are many different EFMs that implement the same outcome as an equilibrium. One possible way to implement the set of ICE is via a simple class of extensive form games in which a non-strategic mediator manages the players through private recommendations. A *mediated mechanism* is characterized by a tuple (α, θ, B) $\alpha \in \Delta(A)$ is a distribution over action profiles to be implemented. $\theta : A \rightarrow \Delta(I)$ specifies a distribution over the order in which players will move, conditional on the action profile to be implemented. $B = \times_i B_i$ specifies actions that can be recommended as *additional* credible threats. The *effective* set of credible threats, $B_i^* = B_i \cup \text{supp } \alpha_i$, also includes the actions played along the equilibrium path. Consider an ICE α with respect to a set of additional threats B ,

The game begins with the mediator privately choosing the action profile a^* that she wants to implement (according to α), and the player i^* to move first (according to $\theta(\cdot | a^*)$). She then “visits” each of the players one by one, visiting i^* first and $-i^*$ second. When visiting each player i , the mediator recommends an action a_i^r , and observes the action actually taken a_i^p . At the moment of making their choices, the

players do not possess any information other than the recommendation they receive. The mediator always recommends the intended action to the first player, i.e. $a_{i^*}^r = a_{i^*}^*$. She recommends the intended action to the second player if the first player complied, and one of the worst available punishments in $B_{-i^*}^*$ otherwise, i.e.:

$$a_{-i^*}^r = a_{-i^*}^* \quad \text{if } a_{i^*}^p = a_{i^*}^*,$$

$$a_{-i^*}^r \in \arg \min_{a_{-i^*} \in B_{-i^*}^*} u_{i^*}(a_{i^*}^p, a_{-i^*}) \quad \text{if } a_{i^*}^p \neq a_{i^*}^*.$$

Note that following the mediator's recommendations constitutes a Nash equilibrium if and only if condition (2.1) is satisfied. Hence every ICE is Nash implementable via a mediated mechanism. This means that mediated mechanisms are a canonical class for Nash implementation in the sense of Forges (1986).

2.5.1. Mediated mechanism for the prisoners' dilemma and timeability

To see an example of a mediated mechanism, consider once again the the prisoners' dilemma from section 2.1 with a few differences. Now suppose that the lawyer can directly control the timing of meetings, and observe the actions of the players, and the prisoners have no notion way of measuring the pass of time. Then, the prisoners could instruct the lawyer as follows:

“You must uniformly randomize the order of our meetings. If the DA offers us a (prisoners' dilemma) deal you must always recommend that we do not confess, unless one of us has already confessed, in which case you must recommend that we do confess. Other than those recommendations, you must not provide us with any additional information.”

The resulting situation is a mediated mechanism corresponding to the extensive form game in Figure 2.5. The red arrows represent the actions recommended by the lawyer at each information set. Indeed, following such recommendations constitutes a sequential equilibrium if and only if $g \leq 1$. This mechanism is equivalent to the one analyzed in Nishihara (1997, 1999).

An crucial feature of this mechanism is that, along the equilibrium path, prisoners are completely uniformed about the order of play. Each of the prisoners cannot distinguish between the node where he is the first one to move, and the node where he is the second one and his accomplice cooperated before him. And he assign the same

There might be real world situation in which agents do not have access to such a powerful non-strategic third-party willing to act as a mediator. It is thus important to keep in mind that mediated mechanisms are one possible way to implement an ICE, but there are be many others. For example, cooperation in a prisoners' dilemma could be implemented via the mediated mechanism from Figure 2.5. And it could also be implemented via the mechanism from Section 2.1, in which prisoners are free to choose the timing of their actions and whether they want to keep their actions public or show them to the lawyer. While this mechanism also involves a lawyer, it is possible to consider environments in which similar mechanisms arise naturally without the intervention of a mediator. The concept of mediated mechanism is itself useful, as it captures the essence of choice interdependence in a clear and stark manner.

Moreover, a lot of different outcomes can still be implemented via mechanisms with weaker mediators. For example, suppose that any deviation from the equilibrium path is publicly observed by everyone and not just the mediator. The outcomes that could be implemented by mediated mechanism under such conditions can be characterized by adjusting the worst punishment functions \underline{w} in the definition of ICE. In this case, the relevant punishment function for player i would be $\underline{w}'_i(a'_i) = \min_{a_{-i} \in \text{BR}_{-i}(a'_i)} u_i(a'_i, a_{-i})$, where BR_{-i} is $-i$'s best response correspondence.

Alternatively, instead of assuming that the mediator controls the order of choices, suppose that she can control the order of her recommendations but players can choose to act before or after they encounter her. In such cases, the mediator could not recommend action-specific punishments. A player who intended to deviate would make his choice after the mediator has left, and thus the mediator could no longer observe the specific deviation. The set of implementable outcomes under these conditions could be characterized by replacing \underline{w}_i with the constant minimax punishment $\underline{w}'_i(a'_i) = \min_{\alpha_{-i} \in \Delta(B^*_{-i})} \max_{a_i \in A_i} U_i(a_i, \alpha_{-i})$.

2.6. Quasi-sequential implementation

2.6.1. Quasi-sequential equilibrium

Sequential equilibrium is defined in terms of *sequential rationality* and *belief consistency*. Sequential rationality requires choices to be optimal at the interim stage

for every information set in the game. Off the equilibrium path, belief consistency requires players to update their beliefs in accordance with some prior assessment of the relative likelihoods of different trembles or mistakes. Furthermore, it requires that these prior assessments should be common to all players. *Quasi-sequential equilibrium* (QSE) imposes sequential rationality and requires beliefs to be consistent with trembles, but allows players to disagree about which deviations are more likely.

For two player environments, it is useful to allow Nature to assign zero probability to some of its available moves. This is because, when faced with a null event, a player can believe that it was Nature who made a mistake instead of necessarily believing that an opponent deviated from the equilibrium.⁸ In order to define consistent beliefs, it is necessary to introduce new notation to denote players' beliefs about Nature's choices, other than s_0 . Let S_0 and S_0^+ denote the sets of mixed strategies and strictly mixed strategies for Nature.

A conditional belief system for i in an extensive form game G , is a function ψ_i mapping i 's information sets to distributions over nodes. $\psi_i(y|H)$ is the probability that i assigns in information set H to being in node y . Let Ψ_i denote the set of i 's conditional belief systems. An assessment is a tuple $(\psi, s) \in \Psi \times \Sigma$ that specifies both players interim and prior beliefs (or strategies). An extended assessment is a tuple $(\psi, s, s'_0) \in \Psi \times \Sigma \times \Sigma_0$ that also specifies prior beliefs on Nature's choices. Given an assessment (ψ, s) , an information set H , and an available move m , $V_i(m|H)$ denotes i 's expected payoff from choosing m at H . The expectation is taken given his interim beliefs $\psi_i(H)$ regarding the current state of the game, and assuming that future choices will be made according to s .

Definition 2.6 (Quasi-sequential equilibrium) An assessment $(\psi^*, s^*) \in \Psi \times \Sigma$ is:

- *Weakly consistent* if and only if for every player there exists a sequence of strictly mixed extended assessments (ψ^n, s^n, s'^n_0) satisfying Bayes' rule, and such that (ψ^n_i, s^n, s'^n_0) converges to (ψ^*_i, s^*, s_0) .
- *Sequentially rational* if and only if $V_i(m|H) \geq V_i(m'|H)$ for every player i , information set $H \in \mathcal{H}_i$ and moves $m, m' \in M(H)$ such that $[s^*_i(H)](m_i) > 0$.
- A *quasi-sequential equilibrium* (QSE) if it is both weakly consistent and sequentially rational.

⁸ It is often assumed that Nature assigns positive probability to all of its available moves, but I am unaware of any good arguments to maintain this assumption. Consider for instance the following quote from (Kreps and Wilson, 1982, pp. 868): "To keep matters simple, we henceforth assume that the players initial assessments [on Nature's choices] are strictly positive".

Sequential rationality requires that the choices that occur off the equilibrium path should be optimal. This implies that players must always believe that the *future* choices of their opponents will be rational, and this fact is common knowledge. However, off the equilibrium path, QSE imposes no restrictions on beliefs about *past* choices, nor agreement of beliefs across different players. In that sense, the difference between QSE and Nash equilibrium can be thought of as a form of *future-looking* rationalizability off the equilibrium path.⁹

The only difference between QSE and sequential equilibrium, is that the former imposes a stronger notion of consistency. Namely, *the same* sequence of strictly mixed assessments should work for all players. Loosely speaking, sequential equilibrium requires choices and beliefs to be in equilibrium, not only along the equilibrium path, but also in every ‘subgame’. In contrast, QSE requires equilibrium along the equilibrium path, but only imposes a form of rationalizability in null ‘subgames’.

The focus on QSE is partially motivated by the fact that it is the finer refinement for which I can provide a complete characterization. However, there may be situations for which it is more appealing than sequential equilibrium. In general, equilibrium is not a straightforward consequence of rational behavior. In order to guarantee equilibrium one must assume mutual or common knowledge of conjectures (Aumann and Brandenburger, 1995), which may be hard to justify off the equilibrium path.

In this respect, focal point arguments may be questioned because of the complexity of determining whether an equilibrium is sequential. Communication can be questioned along similar lines, because planning for all possible contingencies or agreeing on their likelihood may be too complex. Finally, precedence may provide a justification for equilibrium, but repetition provides no experience about events which only happen off the equilibrium path (Fudenberg and Levine, 1993). Hence there might be situations in which (i) it makes sense to assume agreement *exclusively* along the equilibrium path; and yet (ii) rationality and common certainty of rationality may also be defended in every subgame.

2.6.2. Forward-looking interdependent-choice rationalizability and QS implementation

The preceding discussion helps to clarify why it is that being an ICE with respect to A^{ICR} is *not* a necessary condition for PB implementation, despite the fact that

⁹This idea closely resembles the notion of common belief in future rationality from Perea (2013).

ICR is equivalent to common knowledge of rationality with interdependent beliefs (cf. Halpern and Pass (2012)). After observing an unexpected event in an extensive form game, a player might believe that the past choices of his opponent were not rational. Hence, even in a sequential equilibrium of an EFM, it is possible for player choose actions that are rational but not ICR off the equilibrium path.

In fact, two kind of actions can *always* be enforced as credible punishments for QS implementation. ICR punishments are admissible because QSE implementation does not require agreement off the equilibrium path. Hence, the player performing the punishment may very well have counterfactual beliefs which rationalize it. Moreover, since QSE only imposes belief of rationality for future choices, beliefs about *past* can be chosen freely. Best responses to arbitrary *degenerate* conjectures are thus also admissible. These two ideas are embodied in the notion of *future-looking interdependent-choice rationalizability* (FICR).

Definition 2.7 (Future-looking interdependent-choice rationalizability)

- $a_i^* \in A_i$ is *FICR* with respect to $A' \in \mathcal{A}$ if and only if there exists a belief $\lambda_i^0 \in \Delta(A_{-i})$, a counterfactual belief $\lambda_i^1 \in \Lambda(A')$, and some $\mu \in [0, 1]$ such that a_i^* maximizes expected utility with respect to the counterfactual belief $\lambda_i = \mu\lambda_i^0 + (1 - \mu)\lambda_i^1 \in \Lambda_i(A)$. Let $\text{FCR}_i(A')$ denote the set of profiles consisting of FICR actions with respect to A' .
- $A' \in \mathcal{A}$ is *self-FICR* if and only if $A' \subseteq \text{FCR}(A')$.
- The set of *FICR* action profiles $A^{\text{FICR}} \in \mathcal{A}$ is the largest self-FICR set.

As before, A^{FICR} is guaranteed to exist because $\text{FCR}(\cdot)$ is \subseteq -monotone, and thus the union of all self-FC-rationalizable sets is self-FC-rationalizable. Also, it is non-empty because it always contains the set of ICR action profiles.

Intuitively, one can think of λ_i^0 as the arbitrary beliefs (degenerate conjectures) over past deviations, and think of λ_i^1 as the conjectures about future FC-rationalizable choices. With this interpretation, an action a_i is FICR with respect to an action space A' if it is a best response to some conjecture λ_i that assigns full probability to actions in A'_{-i} , *only for choices that occur in the future*. λ_i can assign positive probability to any action, provided that this probability is independent from i 's choice. The set of FICR actions is exactly the set of credible threats that characterizes QS implementation.

Proposition 2.10 *A distribution over action profiles is QS implementable if and only if it is an ICE with respect to A^{FICR} .*

There are two interesting corollaries of this result. First, since sequential implementability implies QS implementability, Proposition 2.10 provides as a necessary condition for sequential implementation in arbitrary environments. Second, since ICR actions are FICR, in games with no absolute dominance a distribution is QS implementable if and only if it is an ICE. This means that requiring QSE instead of Nash equilibrium has a small impact, because most games of interest have no absolutely dominated actions.

Corollary 2.11 *All sequentially implementable outcomes are ICE with respect to A^{FICR} .*

Corollary 2.12 *When there are no absolutely dominated actions, a distribution is quasi-sequentially implementable if and only if it is an interdependent-choice equilibrium.*

This section concludes with a characterization of the FCR operator. Loosely speaking, the following proposition shows that it is equivalent to the elimination of strictly dominated actions in an auxiliary game. Hence, computing A^{FICR} is no more complicated than finding the set of rationalizable actions of a finite game.

Proposition 2.13 *An action $a_i \in A_i$ is FC-rationalizable with respect to an action subspace $A' \in \mathcal{A}$ if and only if there is no $\alpha_i \in \Delta(A_i)$ such that:*

- (i) $\max \{u_i(a_i, a_{-i}) \mid a_{-i} \in A'_{-i}\} < \min \{U_i(\alpha_i, a_{-i}) \mid a_{-i} \in A'_{-i}\}$
- (ii) $u_i(a_i, a_{-i}) < U_i(\alpha_i, a_{-i})$ for all $a_{-i} \in A_{-i} \setminus A'_{-i}$

Chapter 3 |

Pricing Algorithms

Abstract There is an increasing tendency for firms to use pricing algorithms that speedily react to market conditions, such as the ones used by major airlines and online retailers like Amazon. I consider a dynamic model in which firms commit to pricing algorithms in the short run. Over time, their algorithms can be revealed to their competitors and firms can revise them, if they so wish. I show how pricing algorithms not only facilitate collusion but inevitably lead to it. To be precise, within certain parameter ranges, in *any* equilibrium of the dynamic game with algorithmic pricing, the joint profits of the firms are close to those of a monopolist.

In the Spring of 2011, two online retailers offered copies of Peter Lawrence’s textbook *The Making of a Fly* on Amazon for \$18,651,718.08 and \$23,698,655.93 respectively. This was the result of both sellers using automated *pricing algorithms*. Everyday, the algorithm used by seller 1 set the price of the book to be 0.9983 times the price charged by seller 2. Later in the day, seller 2’s algorithm would adjust its price to be 1.27059 times that of seller 1. Prices increased exponentially and remained over one million dollars for at least ten days (!), until one of the sellers took notice and adjusted its price to \$106.23.¹

Automated pricing algorithms—presumably better than the ones outlined in the previous story—are now ubiquitous in many different industries including airlines (Borenstein, 2004), online retail (Ezrachi and Stucke, 2015) and high-frequency trading (Boehmer et al., 2015). Also, while not necessarily automated, algorithms also

¹See *Amazon’s \$23,698,655.93 book about flies*, <http://michaeleisen.org/blog/?p=358>.

feature in hierarchical firms in which top managers design protocols for lower level employees to implement. Optimal pricing algorithms can be highly profitable, as they would be sophisticated enough to recognize and take advantage of profitable collusion opportunities.

In this paper, I formalize these ideas via a model of dynamic competition in continuous time, in which two firms use algorithms to set prices. These algorithms react both to demand conditions and to rivals' prices. At exogenous stochastic times, the current algorithm of a firm becomes apparent to the other firm—because it has either been inferred or “decoded”. When a firm has decoded its rival’s algorithm, it can revise its own algorithm in response.² It is important in my model that this decoding takes (stochastically) longer than the arrival of demand shocks. For instance, demand shocks could arrive weekly on average, while it may take a firm, again on average, up to six months to decode its rival’s algorithm and to implement a new algorithm itself. My main result is the following³

Theorem: When demand shocks occur much more frequently than algorithm revisions, the long-run profits from *any* equilibrium are close to those of a monopolist.

My main result differs sharply from our traditional understanding of dynamic competition. In most models of dynamic competition, there is a plethora of equilibria—some collusive and some not—specially when firms are patient. In the context of repeated games, such results are dubbed “folk theorems” (see [Mailath and Samuelson \(2006\)](#)). This lack of predictability is often viewed as a weakness of the theory.⁴ An early model in which cooperation the unique prediction is due to [Aumann and Sorin \(1989\)](#). Players in their model have limited memory—this is also a feature of all algorithms—and a small amount of incomplete information. Their uniqueness result, however, applies only to games of common interest. My model predicts cooperation in a more general class of games. The main message of this paper is that, when firms

²This revision structure is similar to that of revision games ([Kamada and Kandori, 2009](#)) or Calvo’s pricing model ([Calvo, 1983](#)). An important difference is that, in these papers, agents’ choices consists of a single act. In contrast, agents in my model choose algorithms that can react to market conditions in the short run even before new revision opportunities arise.

³A formal statement is contained in Section 3.3 as Theorem 3.2.

⁴For example, [Green et al. \(2014\)](#) state that “folk theorems deal with the implementation of collusion, and have nothing to say about its initiation. The folk theorem itself does not address whether firms would choose to play the strategies that generate the monopoly outcome nor how firms might coordinate on those strategies.” Then, they proceed with the following quote from [Ivaldi et al. \(2003\)](#) “While economic theory provides many insights on the nature of tacitly collusive conducts, it says little on how a particular industry will or will not coordinate on a collusive equilibrium, and on which one.”

compete via algorithms that are fixed in the short run but can be revised over time, collusion is not only possible but rather, it is *inevitable*.

Key strategic intuitions

The key mechanism underlying my main result can be understood by means of a simple example. Suppose both firms always adopt the simple algorithm which always prices competitively, say, at the level dictated by the one-shot Bertrand equilibrium. Let me argue that this cannot be an equilibrium of the dynamic game. At the first opportunity, firm 1 can deviate to an algorithm that also prices competitively—thus best responding in the short run—but is programmed to match any price increases by its rival. Such a deviation may be thought of as a “proposal” to collude, which will be understood by firm 2 once it decodes firm 1’s algorithm. When this happens, firm 2 will understand that, by raising its own price, it can transition to a more profitable high-price regime. Moreover, suppose that it will then take some time for firm 1 to decode firm 2’s algorithm and revise its own in turn. Firm 2 would then expect the high-price regime to last long enough for the transition to be profitable. Hence, there cannot be equilibria in which firms set low prices forever. While this argument shows that competitive pricing is not an equilibrium, it does not say that equilibrium prices must be close to monopoly prices. This requires a more detailed analysis which is carried out in Section 3.3.

The model has four key features leading to the result that collusion is inevitable. First, and foremost, is that *commitment* is feasible. Because it takes time to revise an algorithm, firms are committed in the short run to a pricing policy.⁵ If algorithms could be revised very frequently, firm 2 might not accept firm 1’s “proposal”—there is the danger that firm 1 would change its algorithm right away. Interestingly, the second key feature is that commitment, while feasible, is *imperfect*. The result would not hold if firms could never revise algorithms. With full commitment, the standard folk theorem arguments would apply. In particular, it would be an equilibrium for both firms to choose a simple algorithm that always implemented the Bertrand price. Hence, the result relies crucially on the fact that firms can revise their algorithms and hence commitment is not perfect.

Third, it is important that pricing algorithms are *responsive* to market outcomes.

⁵Other than the time it takes a firm to decode its rival’s algorithm, similar forms of commitment could arise because of the cost or time it takes to design and implement new algorithms, or the limited attention of the people in charge of doing so.

Without this feature, firm 1 would not be able to offer a price increase while at the same time best-responding in the short run. Finally, it is also important that pricing algorithms are directly *observable* or decodable by rivals. Otherwise, firm 2 would not understand firm 1’s “proposal” to switch regimes.

The proof of the main result mimics the argument of the previous example. . .

Antitrust implications

My findings suggest that pricing algorithms are an effective tool for *tacit* collusion. Hence, as more firms employ pricing algorithms, there may be a need to adjust how anti-trust regulations are enforced. According to [Harrington \(2015\)](#),

“In the U.S., unlawful collusion has come to mean that firms have an *agreement* to coordinate their behavior. . . [E]videntiary standards for determining liability are based on *communications* that could produce mutual understanding and *market behavior* that is the possible consequence of mutual understanding.”

If collusion is reached through the use of pricing algorithms without explicit agreement or direct communication, it might fall outside the scope of the current regulatory framework. Indeed, according to [Ezrachi and Stucke \(2015\)](#),

“Absent the presence of an agreement to change market dynamics, most competition agencies may lack enforcement tools, outside merger control, that could effectively deal with the change of market dynamics to facilitate tacit collusion through algorithms.”

See [Ezrachi and Stucke \(2015\)](#) and [Mehra \(2015\)](#) for a detailed discussion of the legal aspects and challenges of antitrust enforcement regarding industries that use pricing algorithms.

Related literature

Communication and collusion.— The role of dynamic incentives to facilitate collusion has been thoroughly studied since the seminal work of [Friedman \(1971\)](#). However, most of the literature has focused on the possibility, rather than the inevitability of collusion. [Green et al. \(2014\)](#) argue that, in many cases cases, there is reason to believe

that communication might be necessary to *initiate* collusion.⁶ Harrington (2015) finds conditions on firm’s beliefs that are *sufficient* to guarantee collusive outcomes without further communication. Specifically, he assumes that it is common knowledge among firms that “any price increase will be at least matched by the other firms, and failure to do so results in the competitive outcome”. My model suggests that such common understanding could arise naturally when firms set prices through algorithms that can be decoded over time and cannot be continuously revised.

Renegotiation.— The mechanics of my model has some resemblance with the literature on renegotiation. Early work on *implicit* renegotiation imposed axiomatic restrictions on the set of equilibria based on the idea that different equilibria cannot be Pareto ranked, because players would always renegotiate and opt for the Pareto dominant alternative (Pearce, 1987, Bernheim and Ray, 1989, Farrell and Maskin, 1989). More recent work has focused on explicit renegotiation protocols.

Miller and Watson (2013) consider a model of explicit bargaining with transfers, and obtain a completely Pareto unranked set of equilibria assuming that play under disagreement does not vary with the manner in which bargaining broke down. Safronov and Strulovici (2014) study a different protocol that does not satisfy such condition. Under this protocol, Pareto ranked equilibria can coexist because players can be punished by the simple fact of making an offer, and, thus, profitable proposals might be deterred in equilibrium. Such punishments cannot occur in my model because, in the limit when revision opportunities are infrequent, the set of equilibrium continuation values after each specific proposal collapses to a singleton. Hence, the renegotiation protocol induced by the use of algorithms does guarantee that the set of equilibrium continuation values converges to a Pareto unranked set.

Asynchronous moves.— In my model, firms never revise their algorithms exactly at the same time. There are other papers in which asynchronous timing helps to reduce the set of equilibria. In particular, Lagunoff and Matsui (1997) obtain a unique equilibrium for *perfect coordination* asynchronous repeated games.⁷ More recently, Calcagno et al. (2014) consider a model with asynchronous revision opportunities,

⁶Communication among the firms might also play a role in the implementation phase after collusion has already been initiated. This might be the case when firms cannot monitor each other’s prices perfectly (Rahman, 2014, Awaya and Krishna, 2015), or when they have private information about market conditions (Athey and Bagwell, 2001). However, none of these features are present in my model.

⁷A perfect coordination game is one in which players incentives are perfectly aligned, i.e., they share exactly the same preferences. Yoon (2001) shows that a folk theorem applies for any game that is not a perfect coordination game, even if choices are asynchronous.

and establish uniqueness of equilibria for *single-shot* common interest games, and 2×2 opposing interest games. Regarding the study of oligopolies, Maskin and Tirole (1988a,b) study a class of asynchronous pricing games, and obtain unique equilibria restricting attention Markov strategies. Eaton and Engers (1990) study a different pricing game, but they focus on the possibility rather than the necessity of collusion. A salient aspect of my model is that firms do not choose simple acts, but rather algorithms that can react to market outcomes. This allows me to obtain sharp predictions for dynamic duopoly games without restricting attention to Markov strategies.

Ambrus and Ishii (2015) analyze a two player repeated coordination game in which players can only revise their actions at exogenous random times. They show that, in some parameter regions, every Markov-perfect equilibrium leads to efficient coordination in the long run. A crucial difference with the current work is that players in their model choose *acts* rather than contingent plans. As a result, their results do not apply to pricing competition games like the one I consider.

Choices via algorithms.— Rubinstein (1986) and Abreu and Rubinstein (1988) analyze repeated games in which players play via finite automata chosen at the beginning of time, and strictly prefer automata with fewer states. While, in some cases, these considerations narrow the set of equilibria, non-collusive equilibria remain. An important difference is that firms in my model are indifferent regarding the complexity of the algorithms they use. Tennenholtz (2004) studies a model in which players implement strategies via computer programs that can read each other codes. The set of Nash-equilibrium payoffs of this model coincides with the set of feasible and individually rational payoffs. Hence, in the context of oligopolies, such a program could *enable* collusion, but need not guarantee it. Also, all the three mentioned papers focus on the case in which players choose programs or automata at the beginning of the game and never have a chance to revise them.

3.1. Two-period two-price prototype model

Before proceeding to the full model, the following example illustrates the mechanism through which algorithms help to eliminate inefficient equilibria. Consider a two-firm, two-price, two-period environment modelled as a twice repeated prisoner's dilemma with perfect monitoring in which firms set prices via automated algorithms.

The stage game profits are given in the left panel of figure 3.1.

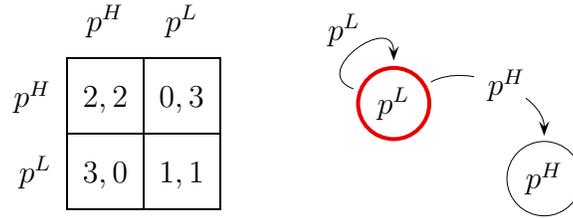


Figure 3.1 – Profit matrix for and “*tit for tat*” algorithm for the prototype model.

An algorithm for firm j is a machine that specifies a price for the first period, and a contingent price for the second period as a function of the price used by j ’s rival on the first period. Consider, for example, the “*tit for tat*” algorithm depicted on the right panel of figure 3.1. It starts by setting the low price on the first period, and mimics its rival’s first period price on the second period.

Before the game begins, firms set algorithms simultaneously and independently. Each period before the game is played, with probability $\mu \in (0, 1/2)$, firm 1 decodes (observes) firm 2’s current algorithm and is able to revise its own in response. With the same probability, it is firm 2 who decodes firm 1’s algorithm and can revise its own. With the remaining probability $1 - 2\mu$, none of the firms is able to revise its current algorithm.

I will argue that, if $\mu < 1/4$, there is *no* (sub-game perfect) equilibrium for the prototype model in which both firms set the low price in both periods. Consequently, under these parameter conditions, the firms’ joint profits in *any* equilibrium need be strictly greater than the competitive profits.

To see this, consider the history in which firm 2 has a revision opportunity on the first period and observes that firm 1’s current algorithm is “*tit for that*”. Firm 2 could choose an algorithm that sets p^L on the first period. In that case, firm 2’s first period profits would be 1. As for the second period, with probability $(1 - \mu)$ firm 1 would not be able to revise its “*tit for tat*” algorithm and would set p^L on the second period as well. In that case, 2’s profits on the second period would be no greater than 1. Even if firm 1 has a revision, firm 2’s profits on the second period cannot exceed 3. Hence, firm 2’s total expected profits from using any algorithm that sets p^L on the first period are bounded above by

$$1 + (1 - \mu)1 + \mu 3 = 2 + 2\mu < 5/2.$$

By an analogous argument, firm 2's expected profits from choosing an algorithm that sets p^H in the first period and p^L on the second period are bounded below by

$$0 + (1 - \mu)3 + \mu 1 = 3 - 2\mu > 5/2.$$

Hence, in any equilibrium, if firm 2 has a revision opportunity on the first period and observes that firm 1 is using “*tit for that*”, it will choose an algorithm that sets p^H on the first period.

Now suppose firm 1 chooses “*tit for that*” at the beginning of the game. If 1's rival sets p^L on the first period, firm 1's algorithm would set p^L on both periods generating profits no less than 2. Otherwise, firm 1's total profits would be weakly greater than 3. We know that, in any equilibrium, firm 2 would set the p^H on the first period if it had a revision opportunity. Hence, choosing “*tit for that*” at the beginning of the game guarantees an expected profit strictly greater than 2. Consequently, there cannot be any equilibrium in which both firms set the low price in both periods, as this would result in each firm's profit being exactly 2.

Note that the firm's ability to commit and to decode their rival's algorithms is not sufficient to preclude low-profit equilibria. Suppose that firms were restricted to choosing prices instead of algorithms. In this case, the strategy profile according to which both firms always choose p^L after any history constitutes an equilibrium of the dynamic game. This is because, if future play does not depend on the firms' current choices, they have no reason to deviate from their dominant price. Hence, it is crucial that firms can choose sophisticated algorithms that react to market conditions. Now suppose that $\mu = 0$, so that firms choose algorithms at the beginning of the game and cannot revise them ever after. If firm 1 expects firm 2 to choose an algorithm that always sets p^L , then choosing the same algorithm is a best response. Hence, it is also crucial that firms can revise their algorithms over time (and thus commitment is imperfect).

3.2. Dynamic price competition with pring algorithms

I consider a symmetric dynamic model of price competition with two firms $j \in \{1, 2\}$. Time is continuous and is indexed by $t \in [0, \infty)$. Consumers arrive randomly following a Poisson process with parameter $\lambda > 0$. Let $y = (y_n)_{n \in \mathbb{N}}$ denote the

sequence of (random) consumer-arrival times.

3.2.1. Price competition

When a consumer arrives in the market, firms simultaneously offer prices p_1 and p_2 , respectively, each belonging to a compact set $P := [0, \bar{p}]$ with $\bar{p} > 0$ an arbitrary exogenous parameter. The consumer observes both prices and decides whether to buy a single unit from one of the firms or to not buy at all. Firm j 's expected profits are given by $\pi_j : P^2 \rightarrow \mathbb{R}_+$. Profit functions are continuously differentiable, non-decreasing in rival's price, sub-modular, and symmetric in that $\pi_1(p, p') = \pi_2(p', p)$.

I use the notation $p = (p_1, p_2)$ and $\pi(p) = (\pi_1(p), \pi_2(p))$. A “tilde” accent denotes joint profits $\tilde{\pi} = \pi_1 + \pi_2$, and the same convention applies to joint expected discounted profits and continuation values defined ahead. Also, let $\Pi = \{\pi(p) \mid p \in P\}$ denote the set of feasible profits. Let $p^M \in P$ be the monopolistic price that maximizes joint profits, and let $\pi^M = \pi_j(p^M, p^M)$ and $\tilde{\pi}^M = 2\pi^M$. The minimax profit for each firm is $\underline{\pi} = \max_{p_j \in P} \pi_j(p_j, 0)$. The static best response profit for firm j is $\bar{\pi}_j(p_{-j}) = \max\{\pi_j(p_j, p_{-j}) \mid p_j \in P\}$, and maximum profit for any firm is $\bar{\pi} = \max\{\pi_j \mid \pi \in \Pi\}$. For tractability, I impose the following assumption.

Assumption 3.1 The set of feasible profits is convex and all Pareto efficient outcomes yield the same joint profits, i.e., $\tilde{\pi} = \tilde{\pi}^M$ for every π in the Pareto frontier of Π .

Assuming that Π is convex guarantees that all feasible expected discounted profits in the dynamic game (to be described ahead) can be achieved by pure stationary strategies. Assuming that the Pareto frontier is a line is more restrictive. Assumption 3.1 is satisfied by different standard models of price competition. For instance, it can be satisfied by linear demand systems functions. It is also satisfied by models with random utility shock demand systems—such as the one in [Berry et al. \(1995\)](#)—taking limits as firm specific shocks are perfectly correlated, and thus their products become perfect substitutes.

3.2.2. Algorithmic pricing

I assume perfect monitoring—each firm observes all past past consumer arrivals, prices and sales. Firms use *pricing algorithms* that automatically set prices as a

function of the history of market outcomes. Pricing algorithms are modelled as finite automata characterized by tuples $a = (\Omega, \omega_0, g, \alpha) \in A$. $\Omega \subseteq \mathbb{N}$ is a set of states and $\omega_0 \in \Omega$ is the initial state. When a consumer arrives to the market and the current state is ω , the algorithm sets price $\alpha(\omega) \in P$ and transitions to a new current state $g(\omega, p)$, where p denotes the price offered by the firm's rival. I restrict attention to algorithms such that g is Borel-measurable and the sets $\{p \in P \mid g(\omega, p) = \omega'\}$ are closed for all $\omega' \in \Omega$ such that $\omega' \neq 1$. These restrictions play the single role of helping to establish existence of equilibria.

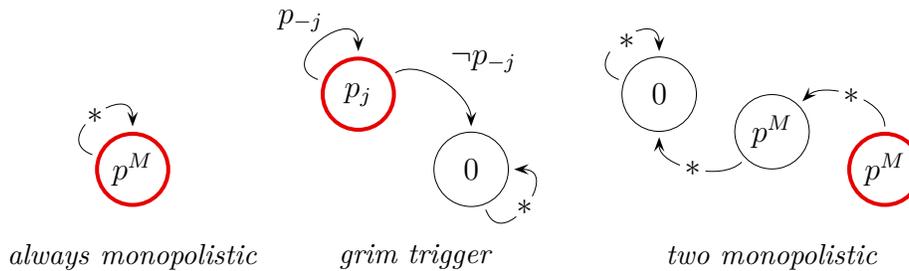


Figure 3.2 – Examples of pricing algorithms.

Figure 3.2 illustrates some examples of pricing algorithms. Each circle represents a state ω , and the price inside the circle corresponds to $\alpha(\omega)$. The initial state is denoted by a bold contour. The transitions are depicted by arrows. The algorithm “*always monopolistic*” always sets the price to p^M . The algorithm “*grim trigger*” sets the price p_j for as long as $-j$ offers p_{-j} , and switches to setting the price to 0 forever after the first time $-j$ sets any price other than p_{-j} . The algorithm “*two monopolistic*” sets the monopolistic price p^M for the first two consumers, and then offers the product for free to every subsequent consumer, regardless of market outcomes.

3.2.3. Dynamic game

The dynamic game begins with firms choosing pricing algorithms simultaneously and independently at time 0. Although firms can perfectly observe each other's prices and sales, they cannot observe each other's algorithms. However, over time, they either infer or are able successfully decode these. The time taken to decode is also random and independent across firms. So, for instance, after the initial choice of algorithms, firm 1 may be the first to decode firm 2's algorithm at time t_1 . At the time firm 1 decodes its rival's current algorithm, it can choose to revise its own algorithm if it so wishes.

Following Kamada and Kandori (2009), I call such an event a *revision opportunity* at time t_1 . Revision opportunities for firm j arise stochastically according to a Poisson process with parameter $\mu_j > 0$. I only consider the symmetric case with $\mu_1 = \mu_2 =: \mu$. Not only is the arrival of revision opportunities independent across firms, it is also independent of the arrival of consumers.

Feasible *paths of play* are fully characterized by sequences $\theta = (a_n, i_n, z_n)_{n=1}^\infty$. The n -th component of the sequence corresponds the n -th pivotal event, be it a consumer arrival or a revision opportunity. $z = (z_n) \in \mathbb{R}_+^\mathbb{N}$ specifies the time of the n -th occurrence.⁸ $i = (i_n) \in \{0, 1, 2\}^\mathbb{N}$ indicates the nature of the event. It takes the value 0 if it is a customer arrival, and the value $j \in \{1, 2\}$ if it is a revision opportunity for firm j . Finally, $a = (a_n) \in (A \times A)^\mathbb{N}$ indicates the profile of algorithms operating right before the event. Given a path θ , $t(\theta)$ denotes the first date after which each firm has had at least one revision opportunity. The *state of the game* at date t consists of the current algorithm profile $\hat{a}_t(\theta)$, and the time $\hat{z}_t(\theta)$ of the last pivotal event before t .

A *history* of length m is a truncated path $h^m = (a_n, i_n, z_n)_{n=1}^m$. Let H^m denote the set of all histories of length m , and $H = \cup_{m \in \mathbb{N}} H^m$. Each firm j must choose a pricing algorithm at time 0, and at each time it has a revision opportunity. Let $H_j^m = \{h \in H^m \mid i_m = j\}$, and let $H_{j0} = \{\emptyset\}$ denote the initial history at time 0. The set of *decision histories* for firm j is $H_j = \cup_{m \in \mathbb{N}} H_j^m$. A (behavior) *strategy* for firm j is a function $s_j : H_j \rightarrow \Delta(A)$. Let S_j denote the set of strategies for firm j , and S the set of strategy profiles $s = (s_j, s_{-j})$.

3.2.4. Subgame perfect equilibria

Firms discount future profits with a common discount rate $r > 0$, and seek to maximize their average expected discounted profits $v_j : S \rightarrow \mathbb{R}$ given by

$$v(s) = (v_1(s), v_2(s)) := \frac{r}{\lambda} \times \mathbb{E}_s \left[\sum_{n=1}^{\infty} \exp(-ry_n) \pi(p_n) \right],$$

where p_n is the price vector offered to the n -th consumer to arrive. The factor r/λ guarantees that expected discounted profits are expressed in the same units as the stage-game profits, so that $v(s) \in \Pi$ for all $s \in S$. Similarly, let $v_t(s, \theta)$ denote continuation expected discounted profits from date t onwards.

⁸I restrict attention to paths in which different events don't occur at the same time. That is, such that, $z_n < z_{n+1}$ for all $n \in \mathbb{N}$. The set of such paths occurs with probability 1.

Recall that there is perfect monitoring of market outcomes, and perfect observability of algorithms at every revision opportunity. Thus, the dynamic game is a game of complete perfect information, except for the first period when initial algorithms are chosen simultaneously. Consequently, it suffices to consider subgame perfect equilibria. Let S^* denote the set of equilibria, and $V^* = \{v(s) \mid s \in S^*\}$ the set of feasible equilibrium profits.

By a similar logic as in the example from Section 3.1, unconditional repetition of a static Nash equilibrium of the stage game might not constitute an equilibrium of the dynamic game. Hence, establishing existence of equilibria is not straightforward. Appendix C.1 includes a proof of existence based on the work of Levy (2015), which builds upon Mertens and Parthasarathy (2003).

Proposition 3.1 *The dynamic game always admits a symmetric equilibrium.*

3.3. Inevitability of collusion

The example in section 3.1 illustrates how pricing algorithms might serve as a renegotiation channel that precludes low-profit equilibria. The main result of the paper is that, in some limiting parameter regions when consumers arrive frequently and revision opportunities are infrequent, the power of renegotiation is such that *all* equilibria of the game lead to joint continuation profits that are *arbitrarily* close to monopolistic profits in the long run. The main result reads as follows.

Theorem 3.2 (Inevitability of collusion) *For every $\epsilon > 0$ there exists $\bar{\mu} : \mathbb{R}_{++} \rightarrow \mathbb{R}_{++}$ and $\underline{\lambda} > 0$ such that, if $\lambda > r\underline{\lambda}$ and $\mu < r\bar{\mu}(\lambda)$, then, at any date after each firm has had at least one revision opportunity, the joint expected discounted continuation profits are closer than ϵ from the joint monopolistic profits with probability greater than $(1 - \epsilon)$, i.e., for all $t_0 > 0$,*

$$\inf_{s \in S^*} \Pr \left(\tilde{v}_{t_0}(s, \theta) > \tilde{\pi}^M - \epsilon \mid t(\theta) < t_0 \right) > 1 - \epsilon.$$

Theorem 3.2 make a sharp prediction about the long run behavior of firms. Namely, it asserts that, in any equilibrium, the game will reach a point after which prices will remain high most of the time, and per-period joint profits will be close to those of

a monopolist. This strong prediction differs sharply from the plethora of equilibria that arise in many standard models of dynamic competition. However, the result is *not* formally an anti-folk-theorem because it is a statement about behavior in the long run. In contrast, folk theorems are statements about the set of date-0 expected discounted equilibrium payoffs.

Establishing Theorem 3.2 requires a number of steps. First, recursive formulation for expected discounted profits (Lemma 3.3). Then, in the limit when revisions are infrequent and consumers arrive frequently, some collusive proposals need to be accepted (Lemma 3.4). Then, I show that, after having a revision opportunity, a firm can include such proposals into its algorithm to guarantee high continuation values. As a consequence, the firms continuation values after subsequent revisions need to be close to Pareto frontier (Lemma 3.5). Finally, I show that such continuation values are *sufficiently* close to the Pareto frontier to guarantee that joint profits will remain high with probability approaching one. The lemmas are formally stated and explained with more detail below. The formal proof of the lemmas and the main result are in Appendix C.2.

3.3.1. Recursive representation

Since the arrivals of revisions and consumers is memoryless, the game has a recursive structure. At any moment in time, the set of continuation strategies that constitute SPNE of the corresponding subgame only depend on the current state of the game. Hence, it is possible to define $V^*(a)$ to be the set of expected discounted continuation equilibrium profits at any moment in which the current algorithm profile is a and there is neither a consumer arrival nor a revision opportunity. Similarly, $W^*(a, j)$ denotes the set of continuation values at a moment where firm j has a revision opportunity and the current algorithm profile is a . The first step of the proof is to obtain tractable bounds for $V^*(a)$.

Consider a moment in which the current algorithm profile is a and there is neither a consumer arrival nor a revision opportunity. For each natural number $n \in \mathbb{N}$, let $\pi^n(a)$ denote the profits that would be generated from the n -th consumer to arrive on the market if no new revisions occur before such arrival. Also, let $\bar{w}_j^n(a)$ and $\underline{w}_j^n(a)$ be the supremum and infimum of the set of possible equilibrium continuation values at the moment of the next revision, conditional on that revision arriving exactly after $(n-1)$ consumers have arrived to the market. Let $\pi(a)$ denote the discounted average

profits

$$\pi_j(a) = (1 - \beta) \sum_{n=0}^{\infty} \beta^n \pi_j^n(a),$$

where β is the *effective discount factor* defined as

$$\beta := \mathbb{E}[\exp(-rz_1) \times \Pr(i_1 = 0)] = \frac{\lambda}{\lambda + 2\mu + r}. \quad (3.1)$$

Analogously, let $\bar{w}_j(a)$ and $\underline{w}_j(a)$ be the expected discounted averages of the sequences $(\bar{w}^n(a))$ and $(\underline{w}^n(a))$, respectively.

Lemma 3.3 *For every firm $j = 1, 2$, and profile of algorithms $a \in A \times A$, and profile of equilibrium continuation values $v \in V^*(a)$,*

$$\frac{r}{\mu + r} \pi_j(a) + \frac{\mu}{\mu + r} \underline{w}_j(a) \leq v_j \leq \frac{r}{\mu + r} \pi_j(a) + \frac{\mu}{\mu + r} \bar{w}_j(a). \quad (3.2)$$

The bounds on expected utility are a convex combination of the profits that would arise if no new revisions were ever to occur and those that would be obtained after the next revision. The relative weight of these two terms depends on the ratio μ/r . When this ratio is very small, firms are very committed to their current algorithms in that the expected amount of time to the next revision opportunity is very long. In that case, firms have the incentive to play a best response the current algorithm ignoring the possibility of future revisions.

3.3.2. An offer that cannot be refused

Recall the the prototype model from section 3.1. In that model, low-profit equilibria are precluded because each firm's unique best response is to set a high price whenever it learns that its rival is using the “*tit-for-tat*” algorithm. A similar principle applies in the general model. However, the relevant algorithms are “*grim trigger*” algorithms corresponding to a special class of price profiles described ahead.

Given a price vector $p^0 \in P$, the corresponding “*grim trigger*” algorithm for firm j is the algorithm that plays p_j^0 as long as it sees p_{-j}^0 and plays 0 forever after any other history, as depicted in Figure 3.2. Say that p^0 is *strictly individually rational* if

$\Delta(p^0, \beta) > 0$, where

$$\Delta(p^0, \beta) := \min_{j=1,2} \left\{ \pi_j(p^0) - \left[(1 - \beta)\bar{\pi}(p_0) + \beta\underline{\pi} \right] \right\}.$$

Note that p^0 is strictly individually rational if and only if the corresponding profile of “grim trigger” algorithms constitutes a strict equilibrium of the discrete-time repeated game without revisions and with discount factor β . A sufficient condition for p^0 to be strictly individually rational is that $\pi_j(p^0) > \underline{\pi}$ for $j = 1, 2$, and β is sufficiently close to 1. In particular, it is straightforward to verify that the monopolistic price vector (p^M, p^M) is strictly individually rational whenever $\mu < r$ and $\lambda > r\underline{\lambda}$ where:

$$\underline{\lambda}(r) := 3 \frac{\bar{\pi}(p^M) - \pi^M}{\pi^M - \underline{\pi}}.$$

Lemma 3.4 (An offer that cannot be refused) *Fix any r , $\mu < r$, and $\lambda > r\underline{\lambda}$. Let p^0 be any strictly individually rational price vector such that $\pi(p^0)$ belongs to the Pareto frontier of the set of feasible profits, and let a^0 be the corresponding “grim trigger” algorithm profile. Every profile of equilibrium continuation values $w \in W^*(a^0_{-j}, j)$ after firm j has a revision opportunity and observes a^0_{-j} satisfies*

$$\bar{\pi}_{j'}(p^0) - w_{j'} < \epsilon(p^0; \lambda, \mu, r) := c_1 \exp \left(-c_2(r, \lambda) \frac{\Delta(p^0, \beta)}{\mu} \right), \quad (3.3)$$

for $j' = 1, 2$, where $c_1, c_2(r, \lambda) \in \mathbb{R}$ are constants that do not depend on μ , nor p^0 .

The intuition behind Lemma (3.4) is as follows. When firm j observes a^0_{-j} it has two alternatives. It can choose an algorithm that plays p^0_j for a “long period of time”, thus guaranteeing the profit of $\pi_j(p^0)$ until $-j$ ’s next revision. Otherwise, it can choose a different algorithm that might yield a high profit the first time it deviates from p^0_j , but, after that, it would activate $-j$ ’s grim trigger and result in profits of at most $\underline{\pi}$ until $-j$ ’s next revision. If p^0 is strictly individually rational and μ is close to 0, it is very unlikely that $-j$ ’s next revision happens soon and, therefore, the first option is preferable. More precisely, in any equilibrium j will choose an automaton that offers p^0_j to the next n consumers to arrive as long as it continues to observe p^0_{-j} , where n is the largest integer satisfying

$$n \leq \underline{N}(p^0; \lambda, \mu, r) := \frac{r}{\mu \bar{\pi}} \Delta(p^0, \beta) - c_3, \quad (3.4)$$

where $c_3 \in \mathbb{R}$ is a constant that does not depend on λ , μ , r nor p^0 . Equation (3.3) follows directly from this fact.

3.3.3. Efficient renegotiation

The previous lemma shows that firms can guarantee high continuation values by using algorithms that admit a unique best response. An important aspect of the model is that firms can choose complex algorithms that have such offers encoded while, at the same time, playing a best response to the current algorithm of their opponents. This guarantees that subsequent revisions must lead to continuation payoffs that are close to being Pareto efficient. If this were not the case, a firm could deviate by encoding a proposal for a Pareto improvement in its algorithm.

Lemma 3.5 (Efficient renegotiation) *Fix any $r > 0$ and $\mu \leq r$ and $\lambda < r\underline{\lambda}$. Consider any equilibrium and a history at which firm j has a revision opportunity, and let \tilde{w} denote firms joint expected discounted continuation profits at the time of firm $-j$'s next revision. For any constant $m > 3$*

$$\Pr\left(\tilde{w} < \tilde{\pi}^M - m\delta(\lambda, \mu, r)\right) \leq \frac{1}{m}, \quad (3.5)$$

with

$$\delta(\lambda, \mu, r) := c_1 \exp\left(-c_4(r, \lambda)\frac{1}{\mu}\right),$$

where c_1 is the constant from Lemma 3.4, and $c_4(r, \lambda)$ is a constant that does not depend on μ .

The main theorem does not follow trivially from Lemma 3.5, because Lemma 3.5 is a statement about continuation values at the moment of each revision, while Theorem 3.2 is about continuation values *at any moment of time* after both firms had at least one revision opportunity each. An additional step is required to show that the bound $m\delta(\lambda, \mu, r)$ is small enough as to guarantee that continuation profits will remain high with high probability before until the time a new revision opportunity arises.

Appendix A

Omitted proofs for Chapter 1

Proof of Proposition 1.1. Fix an action $a \in A$ and a state $x \in X$, and let u and v be compatible vNM utility functions such that

$$\{\alpha \in \Delta(A) \mid u(\alpha, x) \geq u(a, x)\} \subseteq \{\alpha \in \Delta(A) \mid v(\alpha, x) \geq v(a, x)\}. \quad (\text{A.1})$$

We want to show that $\tau_u(a, x) \geq \tau_v(a, x)$. If a is either \succsim_x -maximum or \succsim_x -minimum, then $\tau_u(a, x) = +\infty$ and $\tau_v(a, x) = +\infty$ by definition, and the result is trivial. Hence, we assume for the rest of the proof that $\underline{u}(x) < u(a, x) < \bar{u}(x)$.

By Assumption 1.1, there exists an action $b \in A$ such that $u(b, x) = u^-(a, x)$, and, consequently, $v(b, x) = v^-(a, x)$. Let (a_m) be a sequence of actions such that $a_m \succ_x a$ for all m , $\lim_{m \rightarrow \infty} u(a_m, x) = \bar{u}(x)$, and $\lim_{m \rightarrow \infty} v(a_m, x) = \bar{v}(x)$. Also, for each $\theta \in [0, 1]$ and each $m \in \mathbb{N}$, let $\alpha_{m, \theta}$ be the mixed action that plays a_m with probability θ , and b with probability $1 - \theta$. For all such m we have that $u(\alpha_{m, 1}, x) > u(a, x) > u(\alpha_{m, 0}, x)$. Hence, since expected utility is continuous in the mixing probabilities, there exists $\theta(m) \in (0, 1)$ such that $u(\alpha_{m, \theta(m)}, x) = u(a, x)$. After some simple algebra this implies that:

$$\frac{u(a, x) - u^-(a, x)}{u(a_m, x) - u(a, x)} = \frac{\theta(m)}{1 - \theta(m)}. \quad (\text{A.2})$$

By (A.1), we have that $v(\alpha_{m, \theta(m)}, x) \geq v(a, x)$, which implies that:

$$\frac{v(a, x) - v^-(a, x)}{v(a_m, x) - v(a, x)} \leq \frac{\theta(m)}{1 - \theta(m)}. \quad (\text{A.3})$$

Using (A.2) and (A.3) and taking limits as m goes to infinity thus yields the desired

result

$$\tau_v(a, x) = \lim_{m \rightarrow \infty} \frac{v(a, x) - v^-(a, x)}{v(a_m, x) - v(a, x)} \leq \lim_{m \rightarrow \infty} \frac{u(a, x) - u^-(a, x)}{u(a_m, x) - u(a, x)} = \tau_u(a, x). \quad (\text{A.4})$$

■

Proof of Lemma 1.2. Let $\beta = \alpha(W_x(a))$. Being that $u(b, x) \leq u^-(a, x)$ for $b \in W_x(a)$, and $u(b, x) \leq \bar{u}(x)$ for $b \in B$, it follows that:

$$\begin{aligned} u(\alpha, x) - u(a, x) &\leq \beta(u^-(a, x) - u(a, x)) + (1 - \beta)(\bar{u}(x) - u(a, x)) \\ &= -\beta \left(\frac{u^-(a, x) - u(a, x)}{\bar{u}(x) - u(a, x)} \right) \times \dots \\ &\quad \dots (\bar{u}(x) - u(a, x)) + (1 - \beta)(\bar{u}(x) - u(a, x)) \\ &= (1 - \beta \cdot (\tau_u(a, x) + 1))(\bar{u}(x) - u(a, x)) \leq 0. \end{aligned} \quad (\text{A.5})$$

■

Proof of Proposition 1.3. Fix a set $B \subseteq A$, an action $a \in A \setminus P(B)$, and a mixture α with $\alpha(B \setminus \{a\}) = 1$. There exists some $Y \subseteq X$ conditional on which a is not weakly dominated in B . Assume without loss of generality that for all $b \in B \setminus \{a\}$ there exists some $x \in Y$ such that $b \not\prec_x a$. This implies that for all $b \in B \setminus \{a\}$ there also exists some $x \in Y$ such that $a \succ_x b$, i.e., $B \setminus \{a\} \subseteq \cup_{x \in Y} W_x(a)$. Since $K = \min\{\|A\|, \|X\|\} < +\infty$, there exist a finite subset $Z = \{x_1, \dots, x_k\} \subseteq Y$ with cardinality $k \leq K$, and such that $B \setminus \{a\} \subseteq \cup_{x \in Z} W_x(a)$. Therefore:

$$\sum_{x \in Z} \alpha(W_x(a)) \geq \alpha(B \setminus \{a\}) = 1 \geq \frac{k}{K} = \sum_{x \in Z} \frac{1}{K} \geq \sum_{x \in Z} \frac{1}{\tau_u(a, x) + 1}. \quad (\text{A.6})$$

This implies that there exists a state x such that $(\tau_u(a, x) + 1)\alpha(W_x(a)) \geq 1$, and the result thus follows from Lemma 1.2. ■

Proposition 1.5. Let a, B, α and Y , be as in the proof of Proposition 1.3. As before, we know that $B \setminus \{a\} \subseteq \cup_{x \in Y} C_x(a, B)$, and thus:

$$\sum_{x \in Y} \alpha(C_x(a, B)) \geq 1 \geq \sum_{x \in X} \frac{1}{1 + \tau_u(a, x)} \geq \sum_{x \in Y} \frac{1}{1 + \tau_u(a, x)}. \quad (\text{A.7})$$

Hence, there exists $x \in Y$ such that $(\tau_u(a, x) + 1)\alpha(C_x(a, B)) \geq 1$, and the result follows from Lemma 1.2. ■

Proof of Proposition 1.7. Since $u(\cdot, x)$ is a monotone transformation of $n(\cdot, x)$, u^* is compatible. For dominance equivalence, fix a set $B \subseteq A$, a P -undominated action $a \in A \setminus P(B)$, and a mixed action α with $\alpha(B \setminus \{a\}) = 1$. As, in the proof of proposition 1.3, the fact that $a \notin P(B)$ implies that $B \setminus \{a\} \subseteq \cup_{x \in X} W_x(a)$. Therefore:

$$\begin{aligned} \int_X \alpha(W_x(a)) d\lambda &= \int_X \sum_{a \in W_x(a)} \alpha(a) d\lambda = \int_X \sum_{a \in A} \mathbf{1}(a \in W_x(a)) \alpha(a) d\lambda \\ &= \sum_{a \in A} \alpha(a) \int_{Z(a,b)} d\lambda = \sum_{a \in A} \alpha(a) \lambda(Z(b, a)) \geq \sum_{a \in A} \alpha(a) \delta = \alpha(A) \delta = \delta, \end{aligned} \tag{A.8}$$

where we used the fact that $x \in Z(a, b)$ if and only if $b \in W_x(a)$. Now, let $g_x^*(m) = -\exp(-h^*(x)m)$ so that $u(a, x) = g_x^*(n(a, x))$. Since $g_x^*(m-1) = e^{h^*(x)} g_x^*(m)$ and $g^*(m) < 0$ for all x and $m \in N(x)$, it follows that:

$$\int_X \frac{1}{1 + \pi_u(a, x)} d\lambda < \int_X \exp(-h^*(x)) d\lambda = \delta \int_X f(x) d\lambda \leq \delta. \tag{A.9}$$

From (A.8) and (A.9), that there exists x such that $(\tau_u(a, x) + 1)\alpha(W_x(a)) \geq 1$, and the result thus follows from Lemma 1.2. ■

Appendix B |

Omitted details from Chapter 2

B.1. Alternative definition of extensive form mechanisms

This section provides alternative definitions that slightly generalize the notions of EFM (Definition 2.1) and implementation (Definition 2.2). I will use the term EFM' to distinguish the alternative class of mechanisms defined ahead from the one in the main text. The class of EFM' is a strict superset of the class of EFM. However, all the results in the paper remain true under either definition.

The first requirement for an extensive form game to be an EFM' is that it must preserve the outcome and preference structure of the environment. That is, there must be a preference-preserving map from terminal nodes (outcomes of the game) to action profiles (outcomes of the environment).

Definition B.1 An *outcome homomorphism* is a function τ from terminal nodes *onto* action profiles preserving preferences, i.e. such that $v(z) = u(\tau(z))$ for every terminal node z . G is *outcome equivalent* to E if it admits an outcome homomorphism.

The next requirement is that each player should *freely* choose his own action at some point in the game. Formalizing this idea requires a form of identifying moves (choices in the game) with actions (choices in the environment). For the remainder of this section, let G be outcome equivalent to E and fix an outcome homeomorphism τ . For every player i and every corresponding decision node y , τ induces a *representation* relationship \approx_y from the set of moves available at y in the game to the set of i 's actions in the environment. A move m represents action a_i at y , if and only if choosing m at y in the game has the same effect in (payoff-relevant) outcomes as choosing a_i in the

environment. This idea is formalized by the following definition.

Definition B.2 Given a player $i \in I$ and a decision node $y \in Y_i$, a move $m \in M(y)$ represents an action $a_i \in A_i$ at y if and only if:

- (i) $\tau_i(z) = a_i$ for every $z \in Z(y, m)$
- (ii) There exist $m' \in M(y)$ and $z \in Z(y, m')$ such that $\tau_i(z) \neq a_i$

The representation relationship is denoted by $m \approx_y a_i$, and $M^{a_i}(y)$ denotes the set of moves that represent a_i at y . A move is *pivotal* at y if and only if it represents some action.

The first requirement for $m \approx_y a_i$ is that, if i chooses m at y , then the game will end at a terminal node which is equivalent to a_i according to τ_i . This is regardless of any previous or future moves by either i or his opponents. The second requirement is that, after the game reaches y , i could still choose a different move m' after which the game remains open to the possibility of ending at a terminal node that is *not* equivalent to a_i .

Definition B.3 A decision node $y \in Y_i$ is *pivotal* for player $i \in I$ if and only if $M^{a_i}(y) \neq \emptyset$ for every $a_i \in A_i$. $D_i \subseteq Y$ denotes the set of pivotal nodes for i .

In words, a decision node y is pivotal for player i if for every action $a_i \in A_i$ there exists a pivotal move which represents it at y . Using this language, the next requirement is that *every player makes a pivotal move at a pivotal node* along every possible play of the game. A final technical condition is that a player should always know when he is making a pivotal move representing some action.

Definition B.4 (G, τ) satisfies *full disclosure of consequences* if and only if $\approx_y = \approx_{y'}$ whenever y and y' belong to the same information set.

Finally, the alternative versions of definitions 2.1 and 2.2 are as follows:

Definition 2.1' A *extensive form mechanism* is a tuple (G, τ) satisfying full disclosure of consequences and such that for every terminal node z and every player i , there exists a pivotal node $y \in D_i$ and a pivotal move $m \in M^{\tau_i(z)}$ such that $z \in Z(y, m)$.

Definition 2.2' $\alpha \in \Delta(A)$ is (Nash, sequentially, ...) *implementable* if and only if there exist a mechanism (G, τ) and a (Nash, sequential, ...) equilibrium $s^* \in S$ such that for every $a \in A$:

$$\alpha(a) = \zeta^* \left(\tau^{-i}(a) \right) = \sum_{z \in Z} \zeta(z | s^*, s_0) \cdot \mathbf{1}(\tau(z) = a).$$

B.2. Omitted proofs

B.2.1. Nash implementation

Proof of Proposition 2.3. I will proof the Proposition using the alternative definitions from section B.1. Because every EFM according to Definition 2.1 is also an EFM according to the alternative definition, the proof implies that the result holds true using either of the two definitions. Consider an EMF' mechanism (G, τ) , a NE s^* and let α be the induced distribution. I will show that α is an ICE.

Fix any two of actions $a_i^*, a_i' \in A_i$ with $\alpha_i(a_i^*) > 0$ and $a_i' \neq a_i^*$. For each information set $H \in \mathcal{H}_i$, let $M^*(H)$ be the set of moves that represent a_i^* at H and are chosen with positive probability. Also, let \mathcal{H}_i^* be the set of information sets *along the equilibrium path* in which i chooses a move representing a_i^* with positive probability according to s^* . Finally, let ζ^* be distribution over nodes induced by s^* . All the expectations and conditional distributions in this proof are with respect to ζ^* .

Every $H \in \mathcal{H}_i^*$ must be pivotal, and thus admits a move $m' \in M^{a_i'}(H)$ representing a_i' . Since s^* is a NE, and H is along the equilibrium path, for each $m^* \in M^*(H)$:

$$\mathbb{E} \left[u_i(a_i^*, a_{-i}) \mid H, m^* \right] \geq \mathbb{E} \left[u_i(a_i', a_{-i}) \mid H, m' \right], \quad (\text{B.1})$$

where H, m denotes the set of nodes $H \times \{m\}$ for $m \in \{m^*, m'\}$.

Let $\Phi^H \subseteq H$ denote the event that τ_{-i} is already determined at H , i.e.:

$$\Phi^H = \left\{ y \in H \mid \left(\forall z, z' \in Z(y) \right) \left(\tau_{-i}(z) = \tau_{-i}(z') \right) \right\}, \quad (\text{B.2})$$

and let $\bar{\Phi}^H = H \setminus \Phi^H$ be its complement. Notice that the probability of Φ^H and the distribution of $\tau_{-i}^{-1}(a_{-i})$ conditional on Φ^H , are independent from i 's choice at H .

Hence, by Bayes' rule:

$$\begin{aligned}
\mathbb{E} \left[u_i(a'_i, a_{-i}) \mid H, m' \right] &= \zeta^* \left(\Phi^H \mid H, m' \right) \mathbb{E} \left[u_i(a'_i, a_{-i}) \mid H, m', \Phi^H \right] \\
&\quad + \zeta^* \left(\bar{\Phi}^H \mid H, m' \right) \mathbb{E} \left[u_i(a'_i, a_{-i}) \mid H, m', \bar{\Phi}^H \right] \\
&= \zeta^* \left(\Phi^H \mid H, m^* \right) \mathbb{E} \left[u_i(a'_i, a_{-i}) \mid H, m^* \right] \\
&\quad + \zeta^* \left(\bar{\Phi}^H \mid H, m^* \right) \mathbb{E} \left[u_i(a'_i, a_{-i}) \mid H, m', \bar{\Phi}^H \right] \\
&\geq \zeta^* \left(\Phi^H \mid H, m^* \right) \mathbb{E} \left[u_i(a'_i, a_{-i}) \mid H, m^* \right] + \zeta^* \left(\bar{\Phi}^H \mid H, m^* \right) \underline{w}_i(a'_i).
\end{aligned}$$

Together with (B.1), this yields the following inequality which *does not depend on m'* :

$$\mathbb{E} \left[u_i(a_i^*, a_{-i}) \mid H, m^* \right] \geq \zeta^* \left(\Phi^H \mid H, m^* \right) \mathbb{E} \left[u_i(a'_i, a_{-i}) \mid H, m^* \right] + \zeta^* \left(\bar{\Phi}^H \mid H, m^* \right) \underline{w}_i(a'_i).$$

The last inequality holds for each point in the game where i chooses a_i^* with positive probability. Integrating over them yields:

$$\sum_{a_{-i} \in A_{-i}} \zeta^*(a_i^*, a_{-i}) u_i(a_i^*, a_{-i}) \geq \sum_{a_{-i} \in A_{-i}} \left[\zeta^*(-i, a_i^*, a_{-i}) u_i(a'_i, a_{-i}) + \zeta^*(i, a_i^*, a_{-i}) \underline{w}_i(a'_i) \right].$$

After rearranging terms, using Bayes to write $\zeta^*(i, a_i^*, a_{-i}) = \zeta^*(i \mid a_i^*, a_{-i}) \zeta^*(a_i^*, a_{-i})$, and factorizing $\zeta^*(a_i^*, a_{-i})$, this equation corresponds to condition (2.1). Since i , a_i^* and a'_i were arbitrary, it follows that α is an ICE. \blacksquare

B.2.2. Sequential implementation in 2×2 environments

Proof of Proposition 2.5. First, consider the case in which some player i has an absolutely dominated actions, say a_i , and let a_{-i} be $-i$'s *unique* best response to a_i . In this case (a_i, a_{-i}) is the *unique* ICE with respect to A^{ICR} , and it is a (sequential) equilibrium of the simultaneous move game.

Now, suppose that there are no absolutely dominated actions and let α^* be any *strict* ICE. By Lemma B.1, there exists a distribution α^0 with full support that can arise from a pure-strategy sequential equilibrium s^0 of an EFM G^0 . For each $\mu \in (0, 1)$ let $\alpha^\mu = \mu \alpha^0 + (1 - \mu) \alpha^*$, and let G^μ be the EFM constructed as follows. At the initial node, chance chooses G^0 with probability μ , and the mediated mechanism G^* (see section 2.5) corresponding to α^* with probability $(1 - \mu)$. Then, a new information partition is created, by combining all the information sets in which the mediator recommends some a_i in G^* , with all the information sets in which i plays a_i in G^0

according to s_i^0 . Finally, let s^μ be the strategy profile that follows recommendations in G^* , and mimics s^0 in G^0 .

Since A is finite (which means there are finitely many incentive constraints), α^* is a strict ICE and s^0 is a sequential equilibrium of G^0 , s^μ is a Nash equilibrium for μ sufficiently close to 0. Also, since α^0 has full support, every information set in G^μ is reached with positive probability, and thus s^μ is a sequential equilibrium. Moreover, $\alpha^\mu \rightarrow \alpha^*$ as $\mu \rightarrow 0$. Hence, α^* can be approximated by sequentially implementable distributions. ■

Lemma B.1 *If a 2×2 environment has no absolutely dominated actions, then there is an sequentially pure-strategy implementable distribution α such that $\alpha(a) > 0$ for every $a \in A$.*

Proof. Let $A_i = \{a_i, b_i\}$ for $i = 1, 2$, and suppose that there are no repeated payoffs nor absolutely dominated actions. The result is straightforward if there are no strictly dominated strategies, because then there exists a completely mixed (sequential) equilibrium. As usual, the randomization can be delegated to chance, so that players use pure strategies in the implementation. The interesting cases are when there are no absolutely dominated strategies, but at least one player has a strictly dominated strategy.

Let $\lambda_i, \lambda'_i \in \Lambda_i$ denote the counterfactual beliefs:

$$\lambda_i(a_{-i}|a_i) = 1 \wedge \lambda_i(b_{-i}|b_{-i}) = 1, \quad \text{and} \quad \lambda'_i(b_{-i}|a_i) = 1 \wedge \lambda'_i(a_{-i}|b_{-i}) = 1. \quad (\text{B.3})$$

If b_i is not absolutely dominated but it is strictly dominated by a_i , then it must be a best response to either λ_i or λ'_i . Furthermore, since there are no repeated payoffs, it must be a *strict* best response. There are two cases to consider depending on whether one or two players have dominated strategies.

First suppose that player 2 has no dominated strategies but b_1 is dominated by a_1 . Further assume (without loss of generality) that a_2 is a best response to a_1 . This implies that b_2 is the unique best response to b_1 , and that (a_1, a_2) is a *strict* NE of the simultaneous move game. If b_1 is a best response to λ_1 , then it suffices to have player 1 move first and make his choice public. By backward induction, in the unique SPNE, player 2 will choose a_2 if he chooses a_1 and b_2 if he chooses b_1 . Hence, 1's counterfactual beliefs are λ_1 and b_1 is the unique best response.

Otherwise, if a'_1 is a best response to λ'_1 , then it can be implemented as an equilibrium of the mechanism in Figure (B.1), with $\epsilon > 0$ small enough. The equilibrium

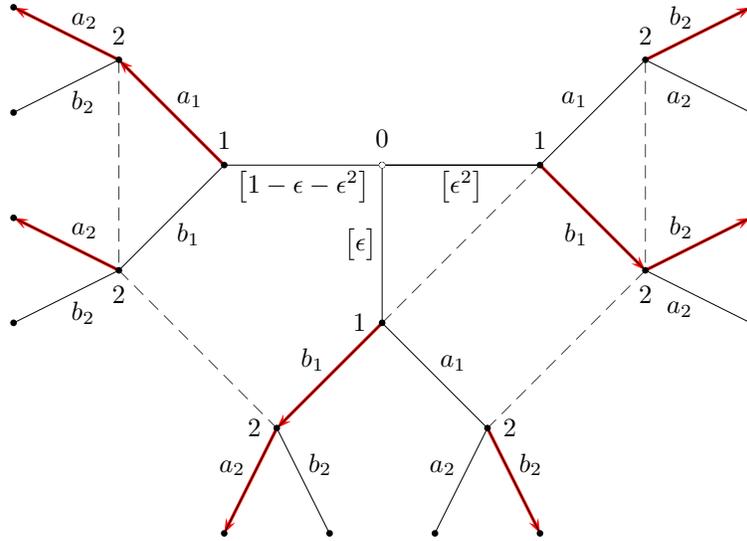


Figure B.1 – Implementation of b_1 when it is the only dominated action.

strategies are represented with arrows. Player's are willing to choose a_i because (a_1, a_2) is a strict Nash equilibrium. Player 2 is willing to choose b_2 because it is a best response to b_1 . Player 1 is willing to choose b_1 because his conjectures at that moment are close enough to λ'_1 . Since all the information sets are on the equilibrium path, the equilibrium is sequential.

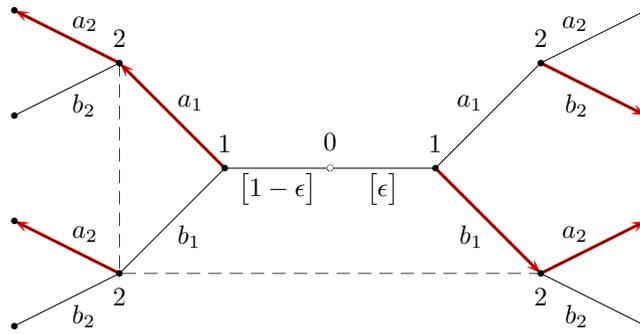


Figure B.2 – Implementation of b_1 when b_2 is also dominated.

Finally, suppose that both players have strictly dominated strategies, say b_1 and b_2 . In this case (a_1, a_2) is a *strict* NE. If b_i is a best response to λ_i , then it can be implemented as a NE of the mechanism where i moves first and $-i$ chooses b_{-i} along the equilibrium path and punishes deviations with a_{-i} . Otherwise, if b_i is a best response to λ'_i , then it can be played with positive probability in a NE of the

mechanism depicted in figure B.2, with $\epsilon > 0$ small enough. Hence there always exists EFMs G^1 and G^2 with NE in which b_1 and b_2 are played with positive probability.

The proof is not complete because the equilibria are not subgame perfect. For that purpose, one can construct a third mechanism in which nature randomizes between G^1 and G^2 and the simultaneous move game, and every action is played with positive probability along the equilibrium path. Information sets can be connected so that, whenever a player is supposed to choose b_i he believes that he is in G^i . Doing so guarantees that the equilibrium is sequential. ■

B.2.3. Properties of ICE, ICR, and FICR

Proof of Proposition 2.1. Note that $u_i(a'_i, a_{-i}) \geq \underline{w}_i(a'_i|B^*)$ by definition. Hence, condition (2.1) becomes tighter for higher values of $\theta(i|a)$. Setting $\theta(i|a) \equiv 0$ in condition (2.1) yields the definition of correlated equilibrium. Hence, α is a correlated equilibrium, then condition (2.1) is satisfied for any θ , which implies that α is also an ICE with respect to any B .

Similarly, if α is an ICE with respect to some B , then condition (2.1) is satisfied for some θ . Consequently, it is also satisfied setting $\theta(i|a) \equiv 1$. This implies that for every $a'_i \in A_i$ we have

$$\sum_{a_{-i} \in A_{-i}} \alpha(a) u_i(a) \geq \underline{w}_i(a'_i|B^*) = \min_{a'_{-i} \in B^*} u_i(a') \geq \min_{a'_{-i} \in A_{-i}} u_i(a')$$

In particular, this is true for whichever a'_i achieves i 's minimax. Therefore,

$$\sum_{a_{-i} \in A_{-i}} \alpha(a) u_i(a) \geq \max_{a'_i \in A_i} \min_{a'_{-i} \in A_{-i}} u_i(a').$$

That is, α is individually rational.

Let $\Gamma \subseteq \Delta(A \times I)$ be the set of distributions satisfying

$$\sum_{a_{-i} \in A_{-i}} \gamma(a) u_i(a) - \gamma(a, -i) u_i(a'_i, a_{-i}) - \gamma(a, i) \underline{w}_i(a'_i|A) \geq 0.$$

Since Γ is defined by a finite set of affine inequalities, it is a closed and convex polytope. The set of ICE with respect to A is the projection of Γ on $\Delta(A)$. ■

Proof of Proposition 2.2. ICR actions are clearly not absolutely dominated. For the opposite direction, fix an action $a_i^* \in A_i$ that is not absolutely dominated in A' . Let

$a_{-i}^* \in \arg \max_{a_{-i} \in A_{-i}} u_i(a_i^*, a_{-i})$. Since a_i^* is not dominated in A' , for every $a'_i \in A_i$ there exists some $a_{-i}(a'_i) \in A'_{-i}$ such that $u_i(a_i^*) \geq u_i(a'_i, a_{-i}(a'_i))$. Hence a_i^* is a best response to $\lambda_i \in \Lambda_i(A')$, with $\lambda_i(a_{-i}^* | a_i^*) = 1$ and $\lambda_i(a_{-i}(a'_i) | a_i^*) = 1$ for $a'_i \neq a_i^*$.

An elimination procedure can be described by a function $K : \mathcal{A} \rightarrow \mathcal{A}$, describing kept actions, such that for $A' \in \mathcal{A}$: (i) never adds new actions, i.e. $K(A') \subseteq A'$; (ii) never eliminates undominated actions, i.e. $\text{CR}(A') \subseteq K(A')$; and (iii) if there are dominated actions then it always eliminates at least one, i.e. $\text{CR}(A') \neq A'$ implies $K(A') \neq A'$.

Now consider the corresponding sequence of surviving actions $(A^n) \in \mathcal{A}^{\mathbb{N}}$ defined recursively by $A^1 = A$ and $A^{n+1} = K(A^n)$. For $n \in \mathbb{N}$ with $A^n \neq \emptyset$, there exists some action profile $a^0 \in A^n$. Ans, since the game is finite, for each player i there exists a best response a_i^* to a_{-i}^0 . By (ii) this implies that $a_i^* \in \text{CR}(A^n) \subseteq A^{n+1}$. Thus, by induction, (A^n) is weakly decreasing sequence of *nonempty* sets. Therefore, since \mathcal{A} is finite, (A^n) converges in finite iterations to a nonempty limit A^* . Since $A^{\text{ICR}} \subseteq \text{CR}(A^{\text{ICR}})$ and $\text{CR}(\cdot)$ is \subseteq -monotone, (ii) implies that $A^{\text{ICR}} \subseteq A^n$ for all $n \in \mathbb{N}$, and thus $A^{\text{ICR}} \subseteq A^*$. Finally, (iii) implies that $A^* \subseteq \text{CR}(A^*)$ and thus $A^* \subseteq A^{\text{ICR}}$. \blacksquare

Proof of Proposition 2.13. $a_i^* \in \text{FCR}_i(A')$ if and only if it is a best response to *some* $\lambda_i = \mu \lambda_i^0 + (1 - \mu) \lambda_i^1$, with $\lambda_i^0 \in \Delta(A_{-i} \setminus A'_{-i})$, $\lambda_i^1 \in \Lambda(A'_{-i})$ and $\mu \in [0, 1]$. Which holds if and only if it is a best response to those beliefs which are more favorable for a_i^* , i.e. beliefs with:

$$\lambda_i^1 \left(\arg \max_{a_{-i} \in A_{-i}} \left\{ u_i(a_i^*, a_{-i}) \right\} \middle| a_i^* \right) = 1, \quad \text{and} \quad \lambda_i^1 \left(\arg \min_{a_{-i} \in A_{-i}} \left\{ u_i(a_i, a_{-i}) \right\} \middle| a_i \neq a_i^* \right) = 1.$$

Hence, after some simple algebra, $a_i^* \in \text{FCR}_i(A')$ if and only if for every $a'_i \in A_i$:

$$(1 - \mu) \left[\bar{w}_i(a_i^*) - \underline{w}_i(a'_i) \right] + \sum_{a_{-i} \notin A'_{-i}} \mu \lambda_i^0(a_{-i}) \left[u_i(a_i^*, a_{-i}) - u_i(a'_i, a_{-i}) \right] \geq 0,$$

where $\bar{w}_i(a_i^*, A') \equiv \max_{a_{-i} \in A'_{-i}} \left\{ u_i(a_i^*, a_{-i}) \right\}$ That is, if and only if it is a best response to some (non-counterfactual) belief in the auxiliary strategic form game $(I, \tilde{A}, \tilde{u})$ with $\tilde{A}_i = A_i$, $\tilde{A}_{-i} = (A_{-i} \setminus A'_{-i}) \cup \{a_{-i}^0\}$, and $\tilde{u}_i : \tilde{A} \rightarrow \mathbb{R}$ given by:

$$\tilde{u}_i(a_i, a_{-i}) = \begin{cases} u_i(a_i, a_{-i}) & \text{if } a_{-i} \notin A'_{-i} \\ \bar{w}_i(a_i^*, A') & \text{if } a_{-i} \in A'_{-i} \wedge a_i = a_i^* \\ \underline{w}_i(a'_i, A') & \text{if } a_{-i} \in A'_{-i} \wedge a_i \neq a_i^* \end{cases} .$$

The result then follows from the well known equivalence between never best responses and dominated actions, cf. Lemma 3 in Pearce (1984). ■

B.2.4. Prisoners' dilemma

Proof of Proposition 2.6. Suppose that $\alpha \in \Delta(A)$ is supported as an ICE by θ . Since D is a dominant strategy, the incentive constraints for it are automatically satisfied. After some simple algebra, the incentive constraints when player 1 is asked to play C can be written as:

$$\theta(1|C, C) \geq \frac{l\alpha(C, D) + g\alpha(C, C)}{(1 + g)\alpha(C, C)}, \quad (\text{B.4})$$

and the corresponding constraint for player 2 can be written as:

$$\theta(1|C, C) = 1 - \theta(2|C, C) \leq \frac{\alpha(C, C) - l\alpha(D, C)}{(1 + g)\alpha(C, C)}. \quad (\text{B.5})$$

Hence α is an ICE if and only if there exists a number $\theta_0 \in [0, 1]$ such that $\theta(1|C, C) = \theta_0$ satisfies both (B.4) and (B.5). This happens if and only if:

$$\frac{\alpha(C, C) - l\alpha(D, C)}{(1 + g)\alpha(C, C)} \geq \frac{l\alpha(C, D) + g\alpha(C, C)}{(1 + g)\alpha(C, C)},$$

which is equivalent to condition (2.2). Note that $\theta(1|C, C) + \theta(2|C, C) = 1$ is the constraint that restricts the set of ICE. This constraint implies that in order to increase $\theta(1|C, C)$ (as to relax (B.4)), one has to decrease $\theta(2|C, C)$ (which tightens (B.5)). ■

Proof of proposition 2.8. The proof is divided into four cases, depending on the properties of the environment induced by μ . First, if (C, C) is an ICE, the total per capita sentence in equilibrium cannot be greater than $-u_i(C, C) = \mu_0$. Second, suppose that (D, D) is a dominant strategy but, $u_i(D, D) \geq u_i(C, C)$. In this case, once again, the per capita sentence cannot be greater than $-u_i(D, D) \leq -u_i(C, C) = \mu_0$. Third, suppose that (C, C) is *not* an ICE, but (D, D) is *not* dominant. Since (C, C) is not an ICE, it must be the case that

$$u_i(C, C) < \frac{1}{2}u_i(D, C) + \frac{1}{2}u_i(D, D) \quad \Rightarrow \quad \mu_2 < 2\mu_0 - \mu_1^+ \quad (\text{B.6})$$

Since (C, C) is not a Nash equilibrium and (D, D) is not dominant, it follows that

$$u_i(D, D) \leq u_i(C, D). \quad (\text{B.7})$$

This implies that (D, C) is a Nash equilibrium and thus an ICE. Consequently, in this case, the per capita sentence cannot be greater than

$$-\frac{1}{2}u_i(D, C) - \frac{1}{2}u_i(C, D) \leq -\frac{1}{2}u_i(D, C) - \frac{1}{2}u_i(D, D) < \frac{1}{2}\mu_i^+ + \mu_0 - \frac{1}{2}\mu_2 = \mu_0,$$

where the first inequality follows from (B.7) and the second one from (B.6). In all these cases, the maximum possible per capita sentence is bounded above by $\mu_0 \leq \bar{\mu}_0 < \min\{\bar{\mu}_1, 2\bar{\mu}_0\}$.

The actual optimal policy belongs to the last remaining case, with (D, D) being strictly dominant, $u_i(C, C) > u_i(D, D)$ and (C, C) *not* being an ICE. This case corresponds to a prisoner's dilemma in which cooperation is not an ICE. Hence, by corollary 2.7, $g > 1$ and the only ICE is (D, D). Which means that the optimal policy is obtained by maximizing μ_2 subject to the following constraints

$$0 \leq \mu_0 \leq \bar{\mu}_0 \quad \wedge \quad 0 \leq \mu_1^+, \mu_1^- \leq \bar{\mu}_1 \quad \wedge \quad 0 \leq \mu_2 \leq \bar{\mu}_2 \quad (\text{B.8})$$

$$\mu_1^+ < \mu_0 \quad \wedge \quad \mu_2 < \mu_1^- \quad \wedge \quad \mu_0 < \mu_2 \quad (\text{B.9})$$

$$2\mu_0 > \mu_2 + \mu_1^+ \quad (\text{B.10})$$

Condition (B.8) requires the policy to be feasible. Condition (B.9) requires the induced environment to be a prisoners' dilemma. Condition (B.10) is equivalent to $g > 1$.

Note that μ_1^+ and μ_1^- only appear as lower and upper bounds of μ_2 , respectively. Hence, it is optimal to set $\mu_1^+ = 0$ and $\mu_1^- = \bar{\mu}_1$. The program thus reduces to maximizing μ_2 subject to

$$2\mu_0 > \mu_2 > \mu_0 \quad \wedge \quad 0 \leq \mu_2 \leq \bar{\mu}_1 \quad \wedge \quad 0 \leq \mu_0 \leq \bar{\mu}_0.$$

Since $\bar{\mu}_0 < \bar{\mu}_1$, the constraint $\mu_0 \leq \bar{\mu}_0$ is binding, which means that the program reduces to

$$\max \left\{ \mu_2 \mid \mu_2 \leq 2\bar{\mu}_0 \wedge 0 \leq \mu_2 \leq \bar{\mu}_1 \right\} = \max\{\bar{\mu}_1, 2\bar{\mu}_0\}.$$

■

B.2.5. Quasi-sequential equilibrium

The proof of Proposition 2.10 is divided in two parts regarding necessity and sufficiency. To establish necessity it suffices to show that given a QSE of an EFM, every action played with positive probability (on or off the equilibrium path) is in A^{FICR} . Then the proof of Proposition 2.3 applies simply replacing $\underline{w}_i(a'_i, A)$ with $\underline{w}_i(a_i, A^{\text{FICR}})$. This fact is established in Lemma B.2. Given an EFM and an QSE (s^*, ψ^*) , let $A_i^* \subseteq A_i$ denote the set of actions that i plays with positive probability is some information set, i.e.:

$$A_i^* = \left\{ a_i \in A_i \mid \left(\exists H \in \mathcal{H} \right) \left(\exists m \in M^{a_i}(H) \right) \left([s_i^*(H)](m) > 0 \right) \right\}.$$

Lemma B.2 *Every quasi-sequential equilibrium s^* of an extensive form mechanism satisfies $A^* \subseteq A^{\text{FICR}}$.*

Proof. Fix some $a_i^* \in A_i^*$ chosen with positive probability in some $H \in \mathcal{H}_i$, and a move $m^{a_i^*} \in M^{a_i^*}(H)$ that represents a_i^* and is chosen with positive probability. For each other action $a'_i \neq a_i^*$, pick a move $m^{a'_i} \in M^{a'_i}(H)$ representing a'_i at H . Now let $\mu = \psi_i^*(\Phi^H \mid H) \in [0, 1]$, where Φ^H is the event that τ_{-i} is already determined at H , as defined in (B.2). Finally, let $\lambda_i^0 \in \Delta(A_{-i})$ and $\lambda_i^1 \in \Lambda_i(A^*)$ be the given by:

$$\lambda_i^0(a_{-i}) = \zeta_i^*(\tau_{-i}^{-1}(a_{-i}) \mid H, \Phi^H), \quad \text{and} \quad \lambda_i^1(a_{-i} \mid a_i) = \zeta_i^*(\tau_{-i}^{-1}(a_{-i}) \mid H, m^{a_i}, \bar{\Phi}^H),$$

and let $\lambda_i = \mu \lambda_i^0 + (1 - \mu) \lambda_i^1$.

Being that $\zeta_i^*(\Phi^H \mid H, m)$ and $\zeta_i^*(\tau_{-i}^{-1}(a_{-i}) \mid H, m, \Phi^H)$ are independent from m , sequential rationality implies that for every deviation a'_i :

$$\begin{aligned} \sum_{a_{-i} \in A_{-i}} \lambda_i(a_{-i} \mid a_i^*) u_i(a_i^*, a_{-i}) &= \sum_{a_{-i} \in A_{-i}} \zeta_i^*(\tau_{-i}^{-1}(a_{-i}) \mid H, m^{a_i^*}) u_i(a_i^*, a_{-i}) \\ &\geq \sum_{a_{-i} \in A_{-i}} \zeta_i^*(\tau_{-i}^{-1}(a_{-i}) \mid H, m^{a'_i}) u_i(a'_i, a_{-i}) \\ &= \sum_{a_{-i} \in A_{-i}} \lambda_i(a_{-i} \mid a'_i) u_i(a'_i, a_{-i}). \end{aligned}$$

Hence a_i^* is a best response to $\lambda_i^* \in \Lambda_i(A^*)$, and thus $a_i^* \in \text{FCR}(A^*)$. This holds for all i and $a_i^* \in A_i^*$. Hence, $A^* \subseteq \text{FCR}(A^*)$ and thus $A^* \subseteq A^{\text{FICR}}$. \blacksquare

The sufficiency proof is constructive, and the mechanics behind the construction are as follows. Every action $a_i^0 \in A_i^{\text{FICR}}$ can be rationalized by some beliefs about

future choices in A_{-i}^{FICR} and about arbitrary equilibrium or arbitrary past choices. Off path beliefs are assigned in such a way that, whenever i is asked to choose a_i^0 , he naively believes that doing so is in his best interest. Since weak consistency does not imply any consistency requirements *across* players, this can always be done even if it implies that i *must be certain that his opponent is or will be mistaken*.

Proof of sufficiency for Proposition 2.10. Fix an ICE α with respect to A^{FICR} . I will construct an extensive form mechanism (G, τ) and a QSE (s^*, ψ) implementing it. As an intermediate step, let G^0 denote the mediated game which implements α as an ICE with respect to A^{FICR} . I will add additional off-path histories to guarantee that the equilibrium becomes QS. Since equilibrium path remains unchanged, it is sufficient to ensure that sequential rationality off the equilibrium path, and that the off-path beliefs are weakly consistent.

In the construction, all the players' information sets are pivotal and have a unique pivotal move representing each action, and all the moves in each pivotal information set are pivotal, i.e. $M(H) = \cup_{a_i \in A_i} M^{a_i}(H)$ and $\#M(H) = \#A_i$ for $H \in \mathcal{H}_i$. Furthermore, the only information that a player has at the moment of making his choice is the action that he is supposed to choose. Hence, it is possible to specify equilibrium strategies by labelling each information set with the distribution of actions that the corresponding player is supposed to follow. For instance H^{a_i} represents a pivotal information set in which, according to s^* , i chooses the only move which represents a_i in H^{a_i} .

Fix a player i and some action $a_i^0 \in A_i^{\text{FICR}} \setminus \text{supp}(\alpha_i)$. Since A^{FICR} is self-FC-rationalizable, a_i^0 is a best response to some counterfactual belief $\lambda_i = (1 - \mu)\lambda_i^0 + \mu\lambda_i^3$, with $\mu \in [0, 1]$, $\lambda_i^0 \in \Delta(A_{-i})$ and $\lambda_i^3 \in \Lambda_i(A^{\text{FICR}})$. $(1 - \mu)\lambda_i^0$ can be further decomposed as $(1 - \mu)\lambda_i^0 = \gamma\lambda_i^1 + \eta\lambda_i^2$ with $\gamma, \eta \in [0, 1]$, $\lambda_i^1 \in \Delta(A_{-i} \setminus A_{-i}^{\text{FICR}})$ and $\lambda_i^2 \in \Delta(A_{-i}^{\text{FICR}})$. Assume without loss of generality that $\lambda_i^3(\bar{a}_{-i}|a_i^0) = 1$ and $\lambda_i^3(\underline{a}_{-i}(a_i^0)|a_i^0) = 1$ for every $a_i^0 \neq a_i^0$, where $\bar{a}_{-i} \in \arg \max_{a_{-i} \in A_{-i}^{\text{FICR}}} \{u_i(a_i^0, a_{-i})\}$ and $\underline{a}_{-i}(a_{-i}) \in \arg \min_{a_{-i} \in A_{-i}^{\text{FICR}}} \{u_i(a_i^0, a_{-i})\}$.

The entire mechanism starts from an initial node where Nature chooses between G^0 and other additional paths. For each action $a_i^0 \in \text{FCR}_i^\infty \setminus \text{supp}(\alpha_i)$, $G^{a_i^0}$ denotes a set of paths on which player i is willing to choose a_i and believe that the future choices of his opponents will be restricted to FCR_{-i}^∞ . The set of paths $G^{a_i^0}$ is depicted in Figure B.3. The nodes are labelled with circled numbers, and the player moving at each node can be inferred from the subindexes of the information sets.

The numbers within brackets, specify the sequence of mixed strategies that converges to the equilibrium assessment. (ϵ_n) denotes an arbitrary sequence of suffi-

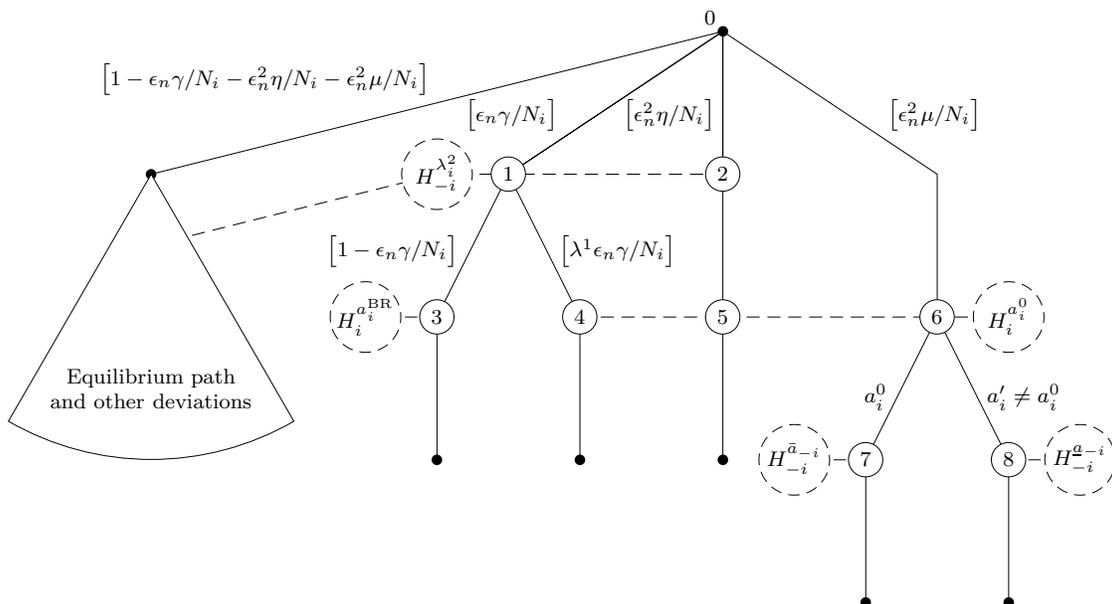


Figure B.3 – Incentives for $a_i^0 \in A_i^{\text{FICR}} \setminus A_i^*$.

ciently small positive numbers converging to 0, and $N_i = \#A_i^{\text{FICR}}$ is the number of FC-rationalizable actions. The sequence is not strictly mixed, but reach all the relevant information sets with positive probability.¹ The limit of this sequence generates weakly consistent beliefs. Hence, it only remains to verify the incentive constraints:

- At nodes (1) and (2), player $-i$ is willing to make choices according to λ_i^2 because he believes that he is on the equilibrium path.
- At nodes (7) and (8), \bar{a}_{-i} and \underline{a}_{-i} may not be best responses to a_i^0 or a_i' . However, they are in FCR_{-i}^∞ and thus $-i$ is willing to play them either along the equilibrium path, or on $G^{\bar{a}_{-i}}$ and $G^{\underline{a}_{-i}}$. Since $-i$ will consider the deviations to and in $G^{a_i^0}$ to be unlikely (of order at most ϵ^3), the incentives for these actions are independent from what happens in this figure.
- First suppose that the information sets for i are fully contained in the figure:
 - At (3) player i is supposed to choose an action which is a best response to λ_i^2 . And therefore his choice is trivially incentive compatible.
 - It is straightforward to see that equilibrium beliefs for player i would generate a conjecture λ_i at nodes (4)–(6), and thus he would be willing to choose a_i^0 .

¹One could use a strictly mixed sequences by assigning probabilities of order ϵ_n^3 or less to other strategies, but this would only complicate the exposition unnecessarily.

- Now suppose that either $H_i^{a_0}$ or $H_i^{a_{BR}}$ appear in other parts of the game. There are only two possibilities:
 - They could appear as punishments in the position analogous to (7) or (8) in some $G^{a_{-i}}$. From i 's perspective, this has probability of order ϵ^3 or lower, and hence it is irrelevant for i .
 - They could appear in the equilibrium path, or in some $G^{a_{-i}}$ in the positions of (3) - (8). In such cases, it will also be a best response to the conditional beliefs and thus to the average beliefs.

■

Appendix C

Omitted proofs for Chapter 3

C.1. Existence of equilibria

I establish existence of equilibria using the main result from Levy (2015), which is based on the work of Mertens and Parthasarathy (2003). Levy's result is for stochastic games in discrete time with compact metrizable action spaces, hence some steps are necessary in order to apply the result to my model. First, I endow the set of algorithms A with a metric ρ , such that (A, ρ) is compact. Then, I consider a discrete version of my model in which Levy's result can be directly applied to prove existence of a SPNE. Finally, I use the corresponding strategies to build a SPNE for my continuous-time model.

C.1.1. Compactness of the set of algorithms

The construction consists in writing each algorithm as a countable product of compact metric spaces, and then endow the set of algorithms with the metric corresponding to the product topology. First, let \mathcal{P} denote the set of closed subsets of $P = [0, \bar{p}]$, and let $\rho^H : \mathcal{P}^2 \rightarrow \mathbb{R}_+$ denote the Hausdorff distance on \mathcal{P} , with the convention that $\rho^H(P', \emptyset) = 2\bar{p}$ for any $P' \in \mathcal{P} \setminus \{\emptyset\}$. Since P is compact, it follows that (\mathcal{P}, ρ^H) is also compact. Also, let $\rho_{\mathbb{N}}$ be any metric on \mathbb{N} , such that $(\mathbb{N}, \rho_{\mathbb{N}})$ is compact, and the corresponding Borel σ -algebra is the power set of \mathbb{N} .

Moreover, given an algorithm $a = (\Omega, \omega_0, \alpha, \theta)$ define the following objects. First, let $g(a) = (g_n(a))_{n \in \mathbb{N}} \in (P \cup \{2\bar{p}\})^{\mathbb{N}}$ be defined by $g_n(a) = \alpha(n)$ if $n \in \Omega$, and $g_n(a) = 2\bar{p}$ otherwise. Second, let $G(a) = (G_{nm}(a)) \in \mathcal{P}^{\mathbb{N} \times \mathbb{N}}$ be defined by $G_{nm}(a) = \{p \in P \mid \theta(n, p) = m\}$ if $n \in \Omega$ and $n \neq m$, and $G_{nm}(a) = \emptyset$ otherwise. In words, if

$n \in \Omega$ and $m \neq n$, $g_n(a)$ is the price set by the algorithm in state n , and $G_{nm}(a)$ is the set of prices p such that, if the original state is n and the rival firm sets price p , the algorithm will transition to state m .

With this notation in mind, let $\rho : A^2 \rightarrow \mathbb{R}_+$ be defined by

$$\begin{aligned} \rho(a, a') &= \rho^{\mathbb{N}}(\omega_0, \omega'_0) + \sum_{n=1}^{\infty} 2^{-n} \frac{|g_n(a) - g_n(a')|}{1 + |g_n(a) - g_n(a')|} \\ &\dots + \sum_{n=1}^{\infty} \sum_{m=1}^{\infty} 2^{-n-m} \frac{\rho^{\mathbb{H}}(G_{nm}(a), G_{nm}(a'))}{1 + \rho^{\mathbb{H}}(G_{nm}(a), G_{nm}(a'))} \end{aligned} \quad (\text{C.1})$$

Lemma C.1 (A, ρ) is a compact metric space.

Proof. Say that two algorithms $a, a' \in A$ are machine-equivalent if they differ only on their initial state. Denote this by $a \sim a'$. Clearly \sim is an equivalence relation on A . Note that the map described above to define $g(a)$ and $G(a)$, is a bijection between $(P \cup \{2\bar{p}\})^{\mathbb{N}} \times \mathcal{P}^{\mathbb{N} \times \mathbb{N}}$ and the quotient set A / \sim .

Since $(\mathcal{P}, \rho^{\mathbb{H}})$ is compact, $\mathcal{P}^{\mathbb{N} \times \mathbb{N}}$ endowed with the product topology is also compact. The product topology is induced by the metric ρ^1 defined by

$$\rho^1(G, G') = \sum_{n=1}^{\infty} \sum_{m=1}^{\infty} 2^{-n-m} \frac{\rho^{\mathbb{H}}(G_{nm}, G'_{nm})}{1 + \rho^{\mathbb{H}}(G_{nm}, G'_{nm})}.$$

Consequently, the function $\rho^1(a, a') = \tilde{\rho}^0(G(a), G(a'))$ is a pseudo-metric on A , and (A, ρ^1) is a compact pseudo-metric space. Similarly, $(P \cup \{2\bar{p}\})^{\mathbb{N}}$ endowed with the (Euclidean) product topology is also compact, and the corresponding metric is ρ^2 given by

$$\rho^2(g, g') = \sum_{n=1}^{\infty} 2^{-n} \frac{|g_n - g'_n|}{1 + |g_n - g'_n|},$$

and we could define a corresponding pseudo-metric ρ^2 on A , and (A, ρ^2) is a compact pseudo-metric space. Finally, $\rho^0(a, a') = \rho^{\mathbb{N}}(\omega_0, \omega'_0)$ is also a pseudo-metric on A such that (A, ρ^0) is compact. Note that we can write $\rho = \rho^0 + \rho^1 + \rho^2$. Since $\rho(a, a') \neq 0$ whenever $a \neq a'$, it follows that ρ is a metric and (A, ρ) is compact. \blacksquare

C.1.2. Discrete-time model

The next step to establish existence is to introduce a discrete-time version of the environment. The discrete-time model actually corresponds to the infinite horizon version of the two-period prototype model from section 3.1. I use the superindex D to denote things related to the discrete model. Time is discrete and indexed by $t^D \in \mathbb{Z}_+$. Firms have a common discount factor

$$\delta^D = \frac{\lambda + 2\mu}{\lambda + 2\mu + r} \in (0, 1).$$

Each period, one and only one of three things happens: (i) with probability $\mu_0^D = \lambda/(\lambda + 2\mu)$, a consumer arrives to the market, (ii) with $\mu_1^D = \mu/(\lambda + 2\mu)$, firm 1 has a revision opportunity, and, (iii) with $\mu_2^D = \mu/(\lambda + 2\mu)$, firm 2 has a revision opportunity. When a consumer arrives, firms play the same stage game as in the continuous-time model simultaneously choosing $p_j \in P$ and realizing profits $\pi_j(p)$.

Algorithms are modelled the same as in the continuous-time model. In particular, just like in the continuous-time model, transitions between the states of the algorithm only take place when a consumer arrives to the market. A strategy chooses an algorithm at the beginning of the game and each time a revision opportunity arises. Formally, a period m history for the discrete-time environment is a tuple $h^{mD} = (a_n^D, i_n^D)_{n=1}^m$. $i_n^D \in \{0, 1, 2\}$ and $a_n^D \in A^2$ specify the type of event that took place at period n and the current algorithms at the beginning of the period, respectively. The decision histories for firm j are the initial history, and those histories h^{mD} of length m such that $i_m^D = j$. Let H^D and H_j^D denote the set of histories and j 's decision histories, respectively. A strategy is a mapping $s_j^D : H_j^D \rightarrow \Delta(A)$.

Lemma C.2 *The discrete-time environment admits a subgame perfect equilibrium.*

Proof. The discrete-time model can be mapped into a discrete-time stochastic game satisfying the assumptions in Levy (2015). Define the state of the game in period t (not to be confused with the states of the algorithms) as the tuple $(i, a_1, a_2) \in \{0, 1, 2\} \times A^2$, where i indicates whether a consumer arrived to the market or a firm had a revision opportunity, and (a_1, a_2) are the firm's current algorithms coming into the period. From Lemma C.1, the state space is a standard Borel space. The set of actions for each firm is simply the set of algorithms which, once again by Lemma C.1, is a compact metric space.

Levy (2015) requires payoff and the transitions between states of the game to

be Borel measurable functions of the state of the game and the actions taken on the current period, and to be continuous in actions. The transition function is continuous in actions because firms directly choose their algorithm, and the identity function is continuous. It is measurable with respect to the state of the game because when a consumer arrives to the market, the only thing that changes regarding the state of the game is the state of the algorithms, and the corresponding Borel algebra is the entire power set of \mathbb{N} .

As for payoffs, the price offered to consumers when they arrive on the market as a function of the state of the game is Borel-measurable, because algorithms have countably many states. Hence, since π is continuous, the profit function is Borel-measurable with respect to the state. Moreover, firms only realize non-null profits when a consumer comes into the market, and only take relevant actions when they have revision opportunities (and hence there are no consumers on the market). Hence, the payoff function is constant with respect to actions, and thus continuous and measurable.

The result thus follows from Theorem 2.1 in [Levy \(2015\)](#). ■

C.1.3. Existence proof

The only remaining task is to use a SPNE of the discrete-time model—which exists by the previous lemma—to construct a SPNE for the continuous-time environment. The main difference between the continuous-time and discrete-time models is that the amount of time that elapses between events in the continuous-time model is random. Say that a strategy for the continuous-time model is said to be *clock-homogeneous* (CH) if it doesn't condition on such random times, the formal definition is included ahead as part of the proof. There SPNE strategy profile of the discrete-time model can be mapped to a profile of CH strategies for the continuous time model, and it turns out that these strategies constitute a SPNE.

Proof of Proposition 3.1. Say that two histories $h, h' \in H$ differ only by clock times if $h_n = h'_n$, and $h_{ak} = h'_{ak}$ and $h_{ik} = h'_{ik}$ for $k \leq n$. Denote this by $h \sim h'$. A strategy σ is said to be CH if $\sigma(h) = \sigma(h')$ whenever $h \sim h'$. Let $\Sigma^{\text{CH}} \subset \Sigma$ denote the set of CH strategies. An equivalence class $[h] \in H / \sim$ is characterized by the common elements (h_n, h_i, h_a) , while a history for the discrete time model $h^{\text{D}} \in H^{\text{D}}$ is a tuple $(h_n^{\text{D}}, h_i^{\text{D}}, h_a^{\text{D}})$. Hence, there is a natural isomorphism between the quotient set H / \sim and H^{D} , and a between the set of CH strategies Σ^{CH} and the set of strategies of the

discrete-time model Σ^D .

Let σ^{D*} denote a SPNE of the discrete time model, which exists by Lemma C.2, and let $\sigma^* \in \Sigma^{CH}$ denote the corresponding CH strategy for the continuous-time environment. I will show that σ^* is a SPNE.

If *both* firms use CH strategies in the continuous-time game, conditional on i , the sequence of prices p is stochastically independent from the arrival time of the events z . Hence, for any profile of CH strategies $\sigma^0 \in \Sigma^{CH}$ and any length- n history $h \in H$, expected discounted continuation profits can be written as

$$\begin{aligned} v(\sigma^0; h) &= \mathbb{E} \left[\sum_{k=1}^{\infty} \exp(-r(y_{n+k} - y_n)) \pi_j(p_{n+k}) \right] \\ &= \sum_{k=1}^{\infty} \mathbb{E}[\exp(-ry_k)] \mathbb{E}[\pi_j(p_{n+k})] = \sum_{k=1}^{\infty} (\delta^D)^k \mathbb{E}[\pi_j(p_{n+k})]. \end{aligned}$$

This expression corresponds exactly to the expected discounted profits generated by the corresponding strategies in the discrete-time model. Hence, the fact that σ^{D*} is a SPNE implies that σ_j^* is a best response to σ_{-j}^* *within the class of CH strategies* after any history, that is

$$v_j(\sigma_j^*, \sigma_{-j}^*; h) \geq v_j(\sigma_j', \sigma_{-j}^*; h) \quad (\text{C.2})$$

for all $h \in H$ and all $\sigma' \in \Sigma^{CH}$.

Now, suppose towards a contradiction that σ^* is not a SPNE. That is, suppose that there is some history h^0 and some strategy $\sigma_j \in \Sigma$ for firm j such that

$$v_j(\sigma_j', \sigma_{-j}^*, h^0) > v_j(\sigma_j^*, \sigma_{-j}^*, h^0). \quad (\text{C.3})$$

Standard arguments can be used to show that the one-shot deviation principle applies in this setting, in that, if there exists a profitable deviation from a proposed strategy, there exists a profitable deviation that differs from the original strategy at a single history. See, for instance, section 2.2 in Mailath and Samuelson (2006). Hence, we can assume without loss of generality that $\sigma_j'(h) = \sigma_j^*(h)$ for all $h \neq h^0$. However, this would imply that the continuation strategy derived from σ_j' conditional on h is CH. Therefore, equations (C.2) and (C.3) would contradict each other. ■

C.2. Inevitability of collusion

C.2.1. Proof of Lemma 3.3

I will only establish the second inequality, the first one can be established using a completely analogous argument. First note that, each time a consumer arrives, the state of each algorithm will either remain the same or transition to a different state. This transition of states can be interpreted as a transition to a different algorithm a' that differs from a only on which state is considered to be the initial one. With this convention, let (a^n) denote the sequence of algorithm profiles that would result starting from the algorithm profile a if no new revisions were ever to arise. Also, let $\bar{v}_j(a) = \sup V_j^*(a)$.

The first event that will happen is either a consumer arrival or a revision opportunity. If it is an arrival there will be some normalized profit and a continuation value w^0 . If it is a revision opportunity, there will be a continuation w^i for $i = 1, 2$. Hence, we can decompose expected discounted profits as

$$\begin{aligned}
 v_j &= \mathbb{E} \left[\exp(-rz_1) \times \left(\mathbf{1}_0 \cdot \left(\frac{r}{\lambda} \pi^1(a) + w_j^0 \right) + \mathbf{1}_1 \cdot w_j^1 + \mathbf{1}_2 \cdot w_j^2 \right) \right] \\
 &\leq \mathbb{E} \left[\exp(-rz_1) \times \left(\mathbf{1}_0 \cdot \left(\frac{r}{\lambda} \pi^1(a) + \bar{v}(a^1) \right) + (\mathbf{1}_1 + \mathbf{1}_2) \cdot \bar{w}^1(a) \right) \right] \\
 &= \frac{\lambda + 2\mu}{\lambda + 2\mu + r} \times \left(\Pr(i_1 = 0) \cdot \left(\frac{r}{\lambda} \pi^1(a) + \bar{v}(a^1) \right) + \Pr(i_1 \neq 0) \cdot \bar{w}^1(a) \right) \\
 &= \frac{r}{\lambda + 2\mu + r} \pi^1(a) + \frac{\lambda}{\lambda + 2\mu + r} \bar{v}(a^1) + \frac{2\mu}{\lambda + 2\mu + r} \bar{w}^1(a) \tag{C.4}
 \end{aligned}$$

Here, I used the fact that z is a Poisson process with arrival rate $\lambda + 2\mu$, and it is independent of i . These are standard properties of the superposition of Poisson processes.

Using a similar argument, and the fact that this is for any s , we obtain the following recursive equation that is satisfied for all $n \in \mathbb{N}$

$$\begin{aligned}
 \bar{v}_j(a^{n-1}) &\leq \frac{r}{\lambda + 2\mu + r} \pi^n(a) + \frac{\lambda}{\lambda + 2\mu + r} \bar{v}(a^n) + \frac{2\mu}{\lambda + 2\mu + r} \bar{w}^n(a) \\
 &= \frac{r}{\lambda + 2\mu + r} \pi^n(a) + \frac{2\mu}{\lambda + 2\mu + r} \bar{w}^n(a) + \beta \bar{v}(a^n). \tag{C.5}
 \end{aligned}$$

Applying (C.5) recursively, it can be shown by induction that

$$\bar{v}_j(a) \leq \frac{r}{\lambda + 2\mu + r} \sum_{k=1}^n \beta^k \pi^k(a) + \frac{2\mu}{\lambda + 2\mu + r} \sum_{k=1}^n \beta^k \bar{w}^k(a) + \beta^n \bar{v}(a^n)$$

for all $n \in \mathbb{N}$. Hence, taking limits as $n \rightarrow \infty$ and using the fact that $\bar{v}(a^n)$ is uniformly bounded, it follows that

$$\begin{aligned} \bar{v}_j(a) &\leq \frac{r}{\lambda + 2\mu + r} \beta^{-1} \sum_{k=0}^n \beta^k \pi^{k+1}(a) + \frac{2\mu}{\lambda + 2\mu + r} \beta^{-1} \sum_{k=0}^n \beta^k \bar{w}^{k+1}(a) \\ &= \frac{r}{2\mu + r} \pi(a) + \frac{2\mu}{2\mu + r} \bar{w}(a), \end{aligned}$$

thus completing the proof. ■

C.2.2. Proof of Lemma 3.4

The following discussion applies at any history in which firm j has a revision opportunity and observes a_{-j}^0 . Say that j accepts $-j$'s offer of a_j^0 for n consumers, if it follows a strategy that sets automata which offer p_j^0 as long as it continues to observe p_{-j}^0 to at least the next n consumers to arrive. Let \underline{N}_j be the minimum amount of consumers for which j must accept $-j$'s offer in any equilibrium. Define \underline{N}_{-j} similarly and let $\underline{N} = \min\{\underline{N}_1, \underline{N}_2\}$.

From Lemma 3.3, it follows that, if j accepts the offer forever, its expected discounted continuation profits are weakly greater than $\underline{v}_j^{(\infty)}$ defined by

$$\underline{v}_j^{(\infty)} := \frac{r}{\mu + r} \pi_j^0 + \frac{\mu}{\mu + r} \underline{w}_j^{(\infty)}, \quad (\text{C.6})$$

where $\underline{w}_j^{(\infty)} = \inf W_j(a_j^0, -j)$ is the worst possible equilibrium continuation value after $-j$'s next revision. Similarly, the best possible continuation value for j if it accepts the offer for exactly n periods and deviates afterwards is bounded above by

$$\bar{v}_j^{(n)} := \frac{r}{\mu + r} \left[\pi_j^0 - \beta^n \Delta_j(p^0, \beta) \right] + \frac{\mu}{\mu + r} \bar{w}_j^{(n)}, \quad (\text{C.7})$$

where $\bar{w}_j^{(n)}$ is the expected discounted averages of the best possible continuation values after $-j$'s next revision. In equilibrium, j must choose an automaton that plays p_j^0 for at least n periods whenever $\underline{v}_j^{(\infty)} \geq \bar{v}_j^{(n)}$. Using (C.6) and (C.7), this condition

can be expressed as

$$\frac{r}{\mu+r}\pi_j^0 + \frac{\mu}{\mu+r}\underline{w}_j^{(\infty)} \geq \frac{r}{\mu+r}\left[\pi_j^0 - \beta^n \Delta_j(p^0, \beta)\right] + \frac{\mu}{\mu+r}\bar{w}_j^{(n)}.$$

After some simple algebra, this yields the following bound for \underline{N}_j

$$\beta^{\underline{N}_j} \leq \frac{\mu}{r\Delta_j(p^0, \beta)} \left(\bar{w}_j^{(\underline{N}_j)} - \underline{w}_j^{(\infty)}\right). \quad (\text{C.8})$$

The next step is to use (C.8) to obtain a bound for the difference in continuation values $\bar{w}_j^{(n)} - \underline{w}_j^{(\infty)}$. Then, substituting with this bound in (C.8) will result in a tighter bound for \underline{N}_j .

To obtain a lower bound for $\underline{w}_j^{(\infty)}$, note that we could repeat the same analysis to conclude that $-j$ will accept the offer a_j^0 for at least \underline{N}_{-j} periods, defined as in (C.8). Using an argument analogous to that of Lemma 3.3, this implies that, by always choosing a_j^0 , firm j can guarantee a continuation value of

$$\underline{w}_j^{(\infty)} \geq \left(1 - \beta^{\underline{N}_{-j}}\right) \pi_j^0. \quad (\text{C.9})$$

Obtaining an upper bound for $\bar{w}_j^{(n)}$ is more complicated. First, Assumption 3.1 and the fact that π^0 is Pareto efficient imply that that all profiles of continuation values w satisfy $\tilde{w} \leq \tilde{\pi}^0$. In particular, this allows to relate the best continuation value $\bar{w}_j^{(n)}$ for firm j , to the worst continuation value under the same circumstances $\underline{w}_{-j}^{(n)}$ for firm $-j$ as follows

$$\bar{w}_j^{(n)} \leq \pi_j^0 + \pi_{-j}^0 - \underline{w}_{-j}^{(n)}. \quad (\text{C.10})$$

This condition allows to find an upper bound for the best continuation value for j , by finding a lower bound for the worst continuation value for $-j$.

The lower bound that I will use is as follows. Suppose that $-j$ has a revision opportunity and j 's current algorithm mimics a_j^0 for m periods with $0 \leq m \leq \underline{N}_{-j}$. By choosing a_{-j}^0 , firm $-j$ guarantees a continuation profit of

$$\begin{aligned} \underline{w}_{-j}^{(m)} &\geq \frac{r}{\mu+r}(1 - \beta^m)\pi_{-j}^0 + \frac{\mu}{\mu+r}(1 - \beta^{\underline{N}_{-j}})\pi_{-j}^0 \\ &= \left(1 - \frac{r}{\mu+r}\beta^m - \frac{\mu}{\mu+r}\beta^{\underline{N}_{-j}}\right)\pi_{-j}^0 \geq (1 - \beta^m)\pi_{-j}^0. \end{aligned} \quad (\text{C.11})$$

As before, the first term of the right-hand side corresponds to the profits that $-j$

would secure if j never had a new revision opportunity. The second term corresponds to the profits after j has a revision. The first inequality follows because, after a revision, j will mimic a_j^0 for at least \underline{N}_j customers. Hence, it follows that

$$\begin{aligned}
\bar{w}_j^{(\underline{N}_j)} &= (1 - \beta) \sum_{k=0}^{\underline{N}_j} \beta^k \bar{w}_j^{(\underline{N}_j - k)} + (1 - \beta) \sum_{k=\underline{N}_j}^{\infty} \beta^k \bar{w} \\
&\leq \pi_j^0 + (1 - \beta) \sum_{k=0}^{\underline{N}_j} \beta^k \left(\pi_{-j}^0 - \underline{w}_{-j}^{(\underline{N}_j - k)} \right) + (1 - \beta) \sum_{k=\underline{N}_j}^{\infty} \beta^k \pi_{-j}^0 \\
&\leq \pi_j^0 + (1 - \beta) \sum_{k=0}^{\underline{N}_j} \beta^{k + \underline{N}_j - k} \pi_{-j}^0 + (1 - \beta) \sum_{k=\underline{N}_j}^{\infty} \beta^k \pi_{-j}^0 \\
&= \pi_j^0 + \beta^{\underline{N}_j} \pi_{-j}^0 \left[(1 - \beta) \underline{N}_j + 1 \right] < \pi_j^0 + (\underline{N}_j + 1) \beta^{\underline{N}_j} \pi_{-j}^0, \tag{C.12}
\end{aligned}$$

the first inequality follows from (C.10), and the second from (C.11).

Combining (C.9) and (C.12) yields the following bound

$$\bar{w}_j^{\underline{N}_j} - \underline{w}_j^{(\infty)} < (\underline{N}_j + 1) \underline{N}_j \beta^{\underline{N}_j} \pi_{-j}^0 + \beta^{\underline{N}_j - j} \pi_j^0 \leq (\underline{N} + 2) \beta^{\underline{N}} \bar{\pi}$$

Suppose without loss of generality that $\underline{N}_j = \underline{N}$, then, substituting in (C.8) yields

$$\beta^{\underline{N}} < \frac{\mu}{r \Delta(p^0, \beta)} (\underline{N} + 2) \beta^{\underline{N}} \bar{\pi}.$$

This implies that

$$\underline{N} > \frac{r}{\mu \bar{\pi}} \Delta(p^0, \beta) - 2$$

which is precisely equation (3.4).

Since offers must be accepted for at least \underline{N} periods, the desired result follows from Lemma (3.4) which implies that

$$w_{j'} > \left(1 - \beta^{\underline{N}(p^0; \lambda, \mu, r)} \right) \pi_j^0$$

Hence, the difference between π_j^0 and j 's continuation value after a history in which firm $-j$ observes a_j^0 is bounded by

$$\pi_j^0 - w_j < \beta^{\underline{N}} \pi_j^0 < \exp \left[(\underline{N} + 2) \log \beta \right] \pi_j^0 < \pi_j^0 \exp \left(\frac{r}{2\mu\bar{\pi}} \Delta(p^0; \beta) \log \beta \right)$$

$$< \bar{\pi} \exp\left(\frac{r[\log(\lambda/(\lambda+r))]}{2\bar{\pi}} \frac{\Delta(p^0; \beta)}{\mu}\right) = c_1 \exp\left(-c_2(\lambda, r) \frac{\Delta(p^0; \beta)}{\mu}\right),$$

where $c_1 = \bar{\pi}$ and $c_2(\lambda, r) = -r \log(\lambda/(\lambda+r))/2\hat{\pi}$. ■

Proof of Lemma 3.5

Fix any equilibrium strategy $s \in S^*$ and any constant $m > 3$, and start from a history in which firm j has a revision opportunity and observers a_{-j}^0 . For tractability, suppose that firm j never has another revision opportunity, a similar argument applies to the general case. Let \hat{n} be the number of consumers before $-j$'s next revision, and, let w^n be the expected equilibrium continuation profits if $-j$ gets a revision exactly after n consumers I will show that

$$\Pr\left(\tilde{w}_{\hat{n}} \geq \tilde{\pi}^M - m\delta(\lambda, \mu, r)\right) \geq \frac{m}{m+1}.$$

For that purpose, note that \hat{n} is distributed according to

$$\Pr(\hat{n} = n) = \left(\frac{\lambda}{\mu + \lambda}\right)^n \left(\frac{\mu}{\mu + \lambda}\right).$$

Also, let $\mathcal{N}(\epsilon)$ be the values of n for which the joint continuation profits are ϵ -close to the joint monopolistic profits, i.e.,

$$\mathcal{N}(\epsilon) = \left\{n \in \mathbb{Z}_+ \mid \tilde{w}^n \geq \tilde{\pi}^M - \epsilon\right\}.$$

Note that the we can write

$$\begin{aligned} \Pr\left(\tilde{w}_{\hat{n}} \geq \tilde{\pi}^M - \epsilon\right) &= \Pr(\hat{n} \in \mathcal{N}(\epsilon)) \\ &= \frac{\mu}{\mu + \lambda} \sum_{n \in \mathcal{N}(\epsilon)} \left(\frac{\lambda}{\mu + \lambda}\right)^n = (1 - \gamma) \sum_{n \in \mathcal{N}(\epsilon)} \gamma^n, \end{aligned} \quad (\text{C.13})$$

where $\gamma := \lambda/(\lambda + \mu)$.

To obtain the desired bound for this probability, consider the following potential deviation. Let a_j^0 be an algorithm that firm j chooses with positive probability at the moment of the current revision. Instead, firm j could choose the alternative algorithm a'_j described as follows. First, let (p^n) be the sequence of prices induced by a^0 . The algorithm a'_j mimics a_j^0 along this sequence. For every n , let \hat{p}^n be unique Pareto

Since $0 < \gamma < \beta < 1$, combining (C.13) and (C.14) yields the desired result. \blacksquare

Proof of Theorem 3.2

We are now in position to prove the main theorem. Fix any $\epsilon > 0$ and $\lambda \geq r\underline{\lambda}$. To simplify the exposition, it is convenient to reparametrize as follows. Let $\epsilon' > 0$ be such that $(1 - \epsilon')^2 = (1 - \epsilon)$, and let $\rho = \epsilon/\tilde{\pi}^M$, I want to show that, for sufficiently small μ , we have that

$$\inf_{s \in S^*} \Pr \left(\tilde{v}_{t_0}(s, \theta) > \tilde{\pi}^M(1 - \rho) \mid t(\theta) < t_0 \right) > (1 - \epsilon')^2.$$

First, from lemma 3.5, $\tilde{w} > \tilde{\pi}^M(1 - m\delta(\lambda, \mu, r))$ at the time of any subsequent renegotiation with probability at least $(m + 1)/m$, for any $m > 3$. Fix some m such that $(m + 1)/m > (1 - \epsilon')$.

Now, I will show that, after n consumers have arrived to the market

$$\tilde{v}^n \geq \tilde{\pi}^M \left(1 - \frac{m\delta(\lambda, \mu, r)}{\beta^n} \right) \quad (\text{C.15})$$

The proof is by induction. It is trivially true for $n = 0$. Now suppose that it is true for some n . Note that:

$$\begin{aligned} \tilde{\pi}^M \left(1 - \frac{m\delta(\lambda, \mu, r)}{\beta^n} \right) &\leq \tilde{v}^n = \frac{r}{\lambda + 2\mu + r} \tilde{\pi}^n + \frac{\lambda}{\lambda + 2\mu + r} \tilde{v}^{n+1} + \frac{2\mu}{\lambda + 2\mu + r} \tilde{w} \\ &\leq \frac{\lambda}{\lambda + 2\mu + r} \tilde{v}^{n+1} + \frac{2\mu + r}{\lambda + 2\mu + r} \tilde{\pi}^M \\ &= \beta \tilde{v}^{n+1} + (1 - \beta) \tilde{\pi}^M, \end{aligned}$$

where the first inequality is the induction hypothesis, the next equality is analogous to (C.4), and the second inequality follows from feasibility. Solving for \tilde{v}^{n+1} yields:

$$\tilde{v}^{n+1} \geq \frac{1}{\beta} \left[\tilde{\pi}^M \left(1 - \frac{m\delta(\lambda, \mu, r)}{\beta^n} \right) - (1 - \beta) \tilde{\pi}^M \right] = \tilde{\pi}^M \left(1 - \frac{m\delta(\lambda, \mu, r)}{\beta^{n+1}} \right).$$

Hence, by the induction principle, (C.15) holds for all $n \in \mathbb{Z}_+$.

A sufficient condition for $\tilde{v}_{t_0}(s, \theta) > \tilde{\pi}^M(1 - \rho)$ is that (i) at the moment of the last revision $\tilde{w} > \tilde{\pi}^M(1 - m\delta(\lambda, \mu, r))$, and (ii) the (random) number of consumers

since the last revision \hat{n} satisfies

$$\frac{m\delta(\lambda, \mu, r)}{\beta^{\hat{n}}} \leq \rho \quad \Leftrightarrow \quad \hat{n} \leq \bar{n}(\rho, \mu) := \frac{\log \rho - \log(m\delta(\lambda, \mu, r))}{-\log \beta}. \quad (\text{C.16})$$

This last condition is valid as long as $m\delta(\lambda, \mu, r) < \rho$, which is true for sufficiently small μ . The probability of this event is thus bounded by

$$\begin{aligned} \Pr\left(\tilde{w} > \tilde{\pi}^M(1 - m\delta(\lambda, \mu, r))\right) &\times \Pr\left(\hat{n} \leq \frac{\log \rho - \log(m\delta(\lambda, \mu, r))}{-\log \beta}\right) \\ &\geq \frac{m}{m+1} \left(1 - \left(\frac{\lambda}{\lambda + 2\mu}\right)^{\bar{n}(\rho, \mu)}\right), \end{aligned}$$

where the first factor follows from Lemma 3.5, and the second one from the fact that the probability of having more than n consumers between revisions is one minus the probability that the first n events are consumer arrivals. Hence, the desired condition obtains as long as

$$\left(\frac{\lambda}{\lambda + 2\mu}\right)^{\bar{n}(\rho, \mu)} \geq \epsilon'$$

The next steps are to substitute with (C.16) and the definition of $\delta(\lambda, \mu, r)$ from Lemma 3.5. Then, doing some simple algebraic manipulations, and substituting ρ and ϵ' with $\epsilon/\tilde{\pi}^M$ and ϵ yields the following sufficient condition:

$$c_1 \exp\left(-c_4(r, \lambda)\frac{1}{\mu}\right) \leq \frac{\epsilon}{m\tilde{\pi}^M} \exp\left(-\frac{\log \beta \log(\epsilon)}{\log((\lambda + 2\mu)/\lambda)}\right).$$

Both the left-hand and the right-hand sides of this condition converge to zero. However, since the left-hand side converges at a faster rate, it follows that the condition is satisfied for μ sufficiently small. ■

Bibliography

- Abreu, D., Pearce, D., and Stacchetti, E. (1990). Towards a theory of discounted repeated games with imperfect monitoring. *Econometrica*, 58(5):1041–1063.
- Abreu, D. and Rubinstein, A. (1988). The structure of nash equilibrium in repeated games with finite automata. *Econometrica*, 56(6):1259–1281.
- Ambrus, A. and Ishii, Y. (2015). On asynchronicity of moves and coordination. Working Paper 185, Economic Research Initiatives at Duke (ERID).
- Athey, S. and Bagwell, K. (2001). Optimal collusion with private information. *The RAND Journal of Economics*, 32(3):428–465.
- Aumann, R. and Brandenburger, A. (1995). Epistemic conditions for nash equilibrium. *Econometrica*, 63(5):1161–1180.
- Aumann, R. J. (1987). Correlated equilibrium as an expression of Bayesian rationality. *Econometrica*, 55(12):1–18.
- Aumann, R. J. and Sorin, S. (1989). Cooperation and bounded recall. *Games and Economic Behavior*, 1(1):5–39.
- Awaya, Y. and Krishna, V. (2015). On communication and collusion. *The American Economic Review*. Forthcoming.
- Bade, S., Guillaume, H., and Renou, L. (2009). Bilateral commitment. *Journal of Economic Theory*, 144(4):1817–1831.
- Bergemann, D. and Morris, S. (2013). Robust predictions in games with incomplete information. *Econometrica*, 81(4):1251–1308.
- Bernheim, B. D. (1984). Rationalizable strategic behavior. *Econometrica*, 52(4):1007–1028.
- Bernheim, B. D. and Ray, D. (1989). Collective dynamic consistency in repeated games. *Games and Economic Behavior*, 1(4):295–326.
- Berry, S., Levinsohn, J., and Pakes, A. (1995). Automobile prices in market equilibrium. *Econometrica*, 63(4):841–890.

- Block, J. I. (2013). Timing and codes of conduct. Mimeo, Washington University in St Louis.
- Block, J. I. and Levine, D. K. (2015). Codes of conduct, private information, and repeated games. *International Journal of Game Theory*. Forthcoming.
- Boehmer, E., Li, D., and Saar, G. (2015). Correlated high-frequency trading. Mimeo.
- Bohner, M. and Gelles, G. M. (2012). Risk aversion and risk vulnerability in the continuous and discrete case. *Decisions in Economics and Finance*, 35(1):1–28.
- Bonanno, G. (2008). A syntactic approach to rationality in games with ordinal payoffs. In Bonanno, G., Van Der Hoek, W., and Wooldridge, M., editors, *Logic and the Foundations of Game and Decision Theory (LOFT 7)*, volume 3 of *Texts in Logic and Games*, pages 59–85. Amsterdam University Press.
- Borenstein, S. (2004). Rapid price communication and coordination: The airline tariff publishing case (1994). In Kwoka, Jr., J. E. and Finnegan, N. F., editors, *The Antitrust Revolution: Economics, Competition, and Policy*, pages 233–251. Oxford University Press.
- Börgers, T. (1993). Pure strategy dominance. *Econometrica*, 61(2):423–430.
- Calcagno, R., Kamada, Y., Lovo, S., and Sugaya, T. (2014). Asynchronicity and coordination in common and opposing interest games. *Theoretical Economics*, 9(2):409–434.
- Calvo, G. A. (1983). Staggered prices in a utility-maximizing framework. *Journal of Monetary Economics*, 12(3):383–398.
- Chambers, C. P., Echenique, F., and Shmaya, E. (2010). General revealed preference theory. Social Science Working Paper 1332, California Institute of Technology.
- Chen, Y. and Luo, X. (2012). An indistinguishability result on rationalizability under general preferences. *Economic Theory*, 51(1):1–12.
- Eaton, J. and Engers, M. (1990). Intertemporal price competition. *Econometrica*, 58(3):637–659.
- Eisert, J., Wilkens, M., and Lewenstein, M. (1999). Quantum games and quantum strategies. *Physical Review Letters*, 83(15):3077–3080.
- Epstein, L. G. (1997). Preference, rationalizability and equilibrium. *Journal of Economic Theory*, 73(1):1–29.
- Ezrachi, A. and Stucke, M. E. (2015). Artificial intelligence & collusion: When computers inhibit competition. Oxford Legal Studies Research Paper 18/2015.
- Farrell, J. and Maskin, E. (1989). Renegotiation in repeated games. *Games and Economic Behavior*, 1(4):327–360.

- Forges, F. (1986). An approach to communication equilibria. *Econometrica*, 54(6):1375–1385.
- Forgó, F. (2010). A generalization of correlated equilibrium: a new protocol. *Mathematical Social Sciences*, 60(1):186–190.
- Friedman, J. W. (1971). A non-cooperative equilibrium for supergames. *The Review of Economic Studies*, 38(1):1–12.
- Fudenberg, D. and Levine, D. K. (1993). Self-confirming equilibrium. *Econometrica*, 61(3):523–545.
- Gafarov, B. and Salcedo, B. (2015). Ordinal dominance and risk aversion. *Economic Theory Bulletin*, 3(2):287–298.
- Gibbard, A. and Harper, W. L. (1980). Counterfactuals and two kinds of expected utility. In Harper, W. L., Stalnaker, R., and Pearce, G., editors, *Ifs, conditionals, beliefs, decision, chance and time*. Springer.
- Green, E. J., Marshall, R. C., and Marx, L. M. (2014). Tacit collusion in oligopoly. In Blair, R. D. and Sokol, D. D., editors, *The Oxford Handbook of International Antitrust Economics*, volume 2, pages 464–497. Oxford University Press.
- Halpern, J. Y. and Pass, R. (2012). Game theory with translucent players. Cornell University, mimeo.
- Halpern, J. Y. and Rong, N. (2010). Cooperative equilibrium. In *Proceedings of the 9th International Conference on Autonomous Agents and Multiagent Systems*.
- Harrington, J. E. (2015). A theory of collusion with partial mutual understanding. Mimeo.
- Howard, N. (1971). *Paradoxes of rationality: theory of metagames and political behavior*. MIT University Press.
- Iijima, R. and Kasahara, A. (2015). Gradual adjustment and equilibrium uniqueness under noisy monitoring. Mimeo.
- Ivaldi, M., Jullien, B., Rey, P., Seabright, P., and Tirole, J. (2003). The economics of tacit collusion. Final report for DG competition, European Commission.
- Jakobsen, S. K., Sørensen, T. B., and Conitzer, V. (2016). Timeability of extensive-form games. In *Proceedings of the 2016 ACM Conference on Innovations in Theoretical Computer Science*, pages 191–199. ACM.
- Kalai, A. T., Kalai, E., Lehrer, E., and Samet, D. (2010). A commitment folk theorem. *Games and Economic Behavior*, 69(1):123–137.
- Kalai, E. (1981). Preplay negotiations and the prisoner’s dilemma. *Mathematical Social Sciences*, 1(4):375–379.

- Kalai, E. and Stanford, W. G. (1985). Conjectural variations strategies in accelerated cournot games. *International Journal of Industrial Organization*, 3(2):133–152.
- Kamada, Y. and Kandori, M. (2009). Revision games. Mimeo.
- Kamada, Y. and Kandori, M. (2011). Asynchronous revision games. Mimeo.
- Klibanoff, P. (2001). Characterizing uncertainty aversion through preference for mixtures. *Social Choice and Welfare*, 18(2):289–301.
- Kreps, D. M. and Wilson, R. (1982). Sequential equilibria. *Econometrica*, 50(4):863–894.
- Lagunoff, R. and Matsui, A. (1997). Asynchronous choice in repeated coordination games. *Econometrica*, 65(6):1467–1477.
- Ledyard, J. O. (1986). The scope of the hypothesis of Bayesian equilibrium. *Journal of Economic Theory*, 39(1):59–82.
- Levine, D. K. and Pesendorfer, W. (2007). The evolution of cooperation through imitation. *Games and Economic Behavior*, 58(2):293–315.
- Levy, Y. (2015). Existence of spe in discounted stochastic games; revisited and simplified. Department of Economics Discussion Paper Series 739, University of Oxford.
- Lewis, D. (1979). Prisoner’s dilemma is a Newcomb problem. *Philosophy and Public Affairs*, 8(3):235–240.
- Lo, K. C. (2000). Rationalizability and the Savage axioms. *Economic Theory*, 15(3):727–733.
- Mailath, G. J. and Samuelson, L. (2006). *Repeated Games and Reputations: Long-run Relationships*. Oxford University Press.
- Maskin, E. and Tirole, J. (1988a). A theory of dynamic oligopoly, I: Overview and quantity competition with large fixed costs. *Econometrica*, 56(2):549–569.
- Maskin, E. and Tirole, J. (1988b). A theory of dynamic oligopoly, II: Price competition, kinked demand curves, and edgeworth cycles. *Econometrica*, 56(2):571–599.
- Mehra, S. K. (2015). Antitrust and the robo-seller: Competition in the time of algorithms. *Minnesota Law Review*, 100. Forthcoming.
- Mertens, J.-F. and Parthasarathy, T. (2003). Equilibria for discounted stochastic games. In Neyman, A. and Sorin, S., editors, *Stochastic Games and Applications*, pages 131–172. Springer.
- Miller, D. A. and Watson, J. (2013). A theory of disagreement in repeated games with bargaining. *Econometrica*, 81(6):2303–2350.

- Monderer, D. and Tennenholtz, M. (2009). Strong mediated equilibrium. *Artificial Intelligence*, 173(1):180–195.
- Moulin, H. and Vial, J. (1978). Strategically zero-sum games: the class of games whose completely mixed equilibria cannot be improved upon. *International Journal of Game Theory*, 7(3-4):201–221.
- Myerson, R. B. (1986). Multistage games with communication. *Econometrica*, 54(2):323–358.
- Nishihara, K. (1997). A resolution of n-person prisoners’ dilemma. *Economic Theory*, 10(3):531–540.
- Nishihara, K. (1999). Stability of the cooperative equilibrium in n-person prisoners’ dilemma with sequential moves. *Economic Theory*, 13(2):483–494.
- Osborne, M. J. and Rubinstein, A. (1994). *A course in game theory*. MIT University Press.
- Pearce, D. (1984). Rationalizable strategic behavior and the problem of perfection. *Econometrica*, 52(4):1029–1050.
- Pearce, D. (1987). Renegotiation-proof equilibria: Collective rationality and intertemporal cooperation. Discussion Paper 855, Yale Cowles Foundation.
- Perea, A. (2013). Belief in the opponents’ future rationality. Maastricht University, mimeo.
- Rahman, D. (2014). The power of communication. *The American Economic Review*, 104(11):3737–3751.
- Rapoport, A. (1965). *Prisoner’s dilemma*. University of Michigan Press.
- Renou, L. (2009). Commitment games. *Games and Economic Behavior*, 66(1):488–505.
- Rockafellar, R. T. (1996). *Convex analysis*, volume 28 of *Princeton Landmarks in Mathematics and Physics*. Princeton University Press.
- Rubinstein, A. (1986). Finite automata play the repeated prisoner’s dilemma. *Journal of Economic Theory*, 39(1):83–96.
- Safronov, M. and Strulovici, B. (2014). Explicit renegotiation in repeated games. Mimeo.
- Tennenholtz, M. (2004). Program equilibrium. *Games and Economic Behavior*, 49(2):363–373.
- Wald, A. (1947). An essentially complete class of admissible decision functions. *The Annals of Mathematical Statistics*, 18(4):549–555.

- Weinstein, J. (2014). The dependence of solution concepts on risk attitude. Mimeo.
- Yoon, K. (2001). A folk theorem for asynchronously repeated games. *Econometrica*, 69(1):191–200.

Bruno Salcedo

<http://brunosalcedo.com>

Education

- *Ph.D. in Economics* — Pennsylvania State University, 2016
Fields: Economic theory, Econometrics, Industrial Organization
Committee: Edward Green, Vijay Krishna, Ruilin Zhou, Lisa Posney
- *B.S. in Applied Mathematics and Economics* — Instituto Tecnológico Autonomo de México, 2010, special mention
- *Partial studies in Electrical Engineering* — University of Houston, 2003–2004

Publications

- *Ordinal dominance and risk aversion* (2015) Joint with Bulat Gafarov. *Economic Theory Bulletin*. 3(2):287–298

Conferences & Seminars

- *Pricing algorithms and tacit collusion* — Presented at Pennsylvania Economic Theory conference (poster, PSU, 2016), Warwick University, Collegio Carlo Alberto, Cornell University, CU-PSU Macro workshop (Cornell, 2015)
- *Identification of solution concepts for discrete semi-parametric games of complete information* — Joint work with Nail Kashaev, Presented at International Conference on Game Theory (SBU, 2015), World Congress of the Econometric Society (Montreal, 2015), NAMES (Minnesota, 2014)
- *Ordinal dominance and risk aversion* — Joint work with Bulat Gafarov, Presented at International Conference on Game Theory (SBU, 2014), Midwest Economic Theory Meetings (IUPUI, 2014)
- *Interdependent choice* — Presented at PERC conference (PSU, 2012), Midwest Theory Meetings (MSU, 2013), International Conference on Game Theory (SBU, 2013), AMES (NUS, 2013), LAMES (Colmex, 2013)
- *Computation of SPNE for repeated games with perfect monitoring and public randomization* — Joint work with Bruno Sultanum, Poster presentation at the Algorithmic Game Theory Summer School (CMU, 2012)

Teaching experience

- *Instructor with full teaching responsibilities* — Decision making and strategy in Economics (PSU, 2012), Introduction to Econometrics (PSU, 2013, 2014, 2015), Knowledge, belief & rationality (NES, 2013)
- *Teaching assistant* — Economics of the corporation (PSU, 2013), Introduction to Macroeconomics (PSU, 2011, 2012), Intermediate Microeconomics (ITAM, 2010)
- *Teaching assistant for Ph.D. core courses* — Microeconomic theory (PSU, 2013), Advanced microeconomic theory (PSU, 2014, 2015)

Fellowships & Awards

- Rosenberg Liberal Arts Centennial Graduate Endowment in Economics (2014), Bates White Fellowship (2013)