

Accepted Manuscript

Equilibrium Selection in Sequential Games with Imperfect Information

Jon X. Eguia, Aniol Llorente-Saguer, Rebecca Morton, Antonio Nicolò


PII: S0899-8256(18)30016-2
DOI: <https://doi.org/10.1016/j.geb.2018.01.008>
Reference: YGAME 2823

To appear in: *Games and Economic Behavior*

Received date: 14 October 2015

Please cite this article in press as: Eguia, J.X., et al. Equilibrium Selection in Sequential Games with Imperfect Information. *Games Econ. Behav.* (2018), <https://doi.org/10.1016/j.geb.2018.01.008>

This is a PDF file of an unedited manuscript that has been accepted for publication. As a service to our customers we are providing this early version of the manuscript. The manuscript will undergo copyediting, typesetting, and review of the resulting proof before it is published in its final form. Please note that during the production process errors may be discovered which could affect the content, and all legal disclaimers that apply to the journal pertain.



Highlights

- A new criterion to choose equilibria in sequential games with imperfect information.
- It selects the action profile that holds for the largest set of possible beliefs.
- It performs better than alternatives in an experiment on vertical contracting.

Equilibrium Selection in Sequential Games with Imperfect Information*

Jon X. Eguia[†]

Aniol Llorente-Saguer[‡]

Rebecca Morton[§]

Antonio Nicolò[¶]

February 5, 2018

Abstract

Games with imperfect information often feature multiple equilibria, which depend on beliefs off the equilibrium path. Standard selection criteria such as passive, symmetric or wary beliefs rest on ad hoc restrictions on beliefs. We propose a new selection criterion that imposes no restrictions on beliefs: we select the action profile that is supported in equilibrium by the largest set of beliefs. We conduct an experiment to compare the predictive power of the existing and our novel selection criteria in an application on vertical multi-lateral contracting. We find that our criterion outperforms the other selection criteria.

Keywords: Equilibrium selection, passive beliefs, symmetric beliefs, vertical contracting, multiple equilibria, imperfect information.

*We thank Antonio Cabrales, Laurent Denant-Boemont, Christoph Engel, Ignacio Esponda, Kristoffel Grechenig, Sergiu Hart, PJ Healy, Asen Ivanov, Matthew Jackson, Isabel Marcin, Marco Ottaviani, Erkut Ozbay, Patrick Rey, Nicolas Roux, Andrew Schotter, Chris Tyson, Piercarlo Zanchettin, seminar attendants at NYU-CESS, Carlos III, Venice, DICE, Hebrew University of Jerusalem, Michigan State, Michigan, Lund, Glasgow, ASSET 2013, Grenoble-Alpes, and workshops at Porto Conte (2013) and Rennes (2014) for their helpful suggestions. We would also like to thank Dominic Land, Nicolas Meier and Maximilian Pohl for excellent assistance in running the experiments. We acknowledge financial support from the Max Planck Society, NYU-NYC, NYU-AD and Università di Padova.

[†]Michigan State U. Email (corresponding author): eguia@msu.edu

[‡]School of Econ. and Finance, Queen Mary U. London, and CEPR. Email: a.llorente-saguer@qmul.ac.uk

[§]Dept. of Politics, NYU-NYC and NYU-Abu Dhabi. Email: rebecca.morton@nyu.edu

[¶]Dept. of Economics and Management, Università degli Studi di Padova, and School of Social Sciences, U. of Manchester. Email: antonio.nicolo@unipd.it

JEL Codes: C72, C92, D86.

1 Introduction

We propose a solution to the problem of multiplicity of equilibria in a class of two-stage games in which players who move at the second stage (receivers) are imperfectly informed about the actions played by those who move at the first stage (proposers). The most prominent application of this class of games is multi-lateral vertical contracting. These are games in which one or more upstream firms make private offers to each of two or more downstream firms. Each contract signed by a downstream firm affects all downstream firms (contracts generate externalities), but at the time a downstream firm decides whether or not to accept the offer, it does not know what offers other firms have received (downstream firms operate under imperfect information).¹

Other applications of this class of games can be found in the literature of consumer search; international trade; markets with network effects; markets with intermediaries; financial and health care markets; and in games of electoral competition.²

These games typically feature multiple Sequential equilibria. Equilibria depend on how, after observing a deviation, the second movers update their beliefs about the first movers' actions. Refinements that are useful for signaling games such as the intuitive criterion (Cho and Kreps, 1987), D1 or universal divinity (Banks and Sobel, 1987) have no bite in this context, because the lack of information is about

¹See theories by Hart and Tirole (1990); O'Brien and Shaffer (1992); Segal (1999); de Fontenay and Gans (2005); Caprice (2006); Rey and Tirole (2007); Nocke and Rey (2014); Miklos-Thal and Shaffer (2016), to cite only a few.

²See Bar-Isaac, Caruana and Cuñat (2012), Buehler and Schuett (2014), or Janssen and Shelegia (2014) on consumer search; Bernard and Dhingra (2015) on trade; de Fontenay and Gans (2014) or Hagiú and Halaburda (2014) and Mayzlin and Yoganarasimhan (2012) on network effects; Inderst and Ottaviani (2012) on markets with intermediaries; Ho (2009) on medical care markets; Brunnermeier and Oehmke (2013) on financial markets; Gavazza and Lizzeri (2009) on electoral competition.

the first movers' actions, not about their type.

The literature has dealt with this multiplicity of equilibria by imposing particular beliefs off the equilibrium path, selecting equilibria that can be supported by these beliefs, and discarding all other equilibria. In their seminal paper, McAfee and Schwartz (1994) propose three possible beliefs to consider: passive beliefs, symmetric beliefs, and wary beliefs.

Passive beliefs, sometimes called “passive conjectures” (Rey and Tirole, 2007), are such that a downstream firm that receives an out of equilibrium offer does not update its beliefs on the offers received by all other players; rather, it believes that all the other unobserved actions remain as in equilibrium. The selection criterion based on singling out equilibria that can be supported by such passive beliefs is the one most frequently used in the literature.³

However, “*in many circumstances the ad hoc restriction to passive beliefs may not be compelling*” (Segal and Whinston, 2003). Indeed, while defending the assumption of passive beliefs in the particular game in which they use it, Rey and Tirole (2007) concede that assuming passive beliefs “*is much less appealing in the case of Bertrand competition, and indeed in many games of contracting with externalities.*”

An alternative criterion to solve the multiplicity problem is to select equilibria that can be supported by symmetric beliefs. These beliefs are such that a downstream player who receives an out of equilibrium offer believes that all other downstream players receive this same offer as well. If the equilibrium offer to each downstream player i is x_i , a downstream player who receives an offer of $y \neq x_i$ believes that the offer to every downstream player is also y .⁴

A third suggestion is to consider equilibria supported by wary beliefs. These be-

³See Hart and Tirole, 1990; O'Brien and Shaffer, 1992; McAfee and Schwartz, 1994; Rey and Tirole, 2007; Ho, 2009; Arya and Mittendorf, 2011; Brunnermeier and Oehmke, 2013; Caprice and Rey, 2015; Reisinger and Tarantino, 2015, among many others.

⁴Symmetric beliefs are used, among others, by Pagnozzi and Piccolo (2011). However, most of the literature seems to agree with McAfee and Schwartz's (1994): “[*Symmetric*] beliefs are not very compelling.”

liefs are such that a downstream player who observes a deviation believes that the upstream player must have deviated to a strategy that is optimal given the action that the downstream player observes. Despite their greater conceptual merit, unfortunately wary beliefs have had scant following in the literature of vertical contracting (Rey and Vergé, 2004; Avenel, 2012a; Reisinger and Tarantino, 2015). Hagiu and Hałaburda (2014) apply it in a setup with markets with two-sided network effects.⁵

The problem common to all these criteria is the lack of a convincing argument for why only one particular set of beliefs should be admissible off the equilibrium path. The existence of several alternative choices of specific beliefs that have received consideration in the literature underscores that none of these beliefs are an obvious choice for all possible games of imperfect information. Eliminating all equilibria that rest on different beliefs is not warranted. We argue that passive beliefs, wary beliefs, or symmetric beliefs, may be plausible in a given particular application, but not in others. A sharp restriction on the set of admissible beliefs to the exclusion of all others is often inappropriate. Martin, Normann and Snyder (2001) provide evidence of the weakness of this approach. Using a laboratory experiment that mimics a vertical industry structure with an upstream firm and two competing downstream firms, they find that no specific restriction on beliefs (passive or symmetric beliefs) is consistent with the data.⁶

A more cautious or modest approach is to accept that we cannot pin down the exact beliefs off the equilibrium path, beyond the standard consistency restriction on beliefs required in a Sequential equilibrium (Kreps and Wilson, 1982). Any assumption of specific equilibrium beliefs, even a plausible one, is ad hoc and it is difficult

⁵Passive and wary beliefs are justified by two very different rationales. Passive beliefs consider deviations from the equilibrium strategy as unintended mover's mistakes; therefore, receivers should not revise their beliefs. In contrast, wary beliefs implicitly interpret deviations as mover's intentional attempts to leave the equilibrium path. In that case, receivers should revise their beliefs accordingly. Still, wary beliefs cannot rationalize why in the first instance the mover deviated from the equilibrium strategy.

⁶See as well the contracting experiments by Boone, Müller and Suetens (2014), and by Möllers, Normann and Snyder (2016).

to justify as superior to all others.

We solve the multiplicity problem by suggesting that the equilibrium action profile most likely to emerge is the action profile that can be supported in equilibrium by the largest set of different beliefs. In application of Bernoulli's (1713) and Laplace's (1820) "Principle of Insufficient Reason," absent any motivation to consider some equilibrium beliefs as more likely to emerge than others, we should treat all equilibrium beliefs as a priori equally probable to emerge.⁷ If so, the strategy profile and action profile most likely to be played in equilibrium are the ones that can be sustained by the largest collection of different equilibrium beliefs. We predict that this action profile will be the one played in equilibrium.

We do not identify the specific beliefs that support the predicted action profile in equilibrium: we make a prediction only about players' equilibrium actions. Since actions (unlike beliefs) are directly observable, our predictions are directly testable.

Our selection criterion provides a "stress test" measure of the robustness of an equilibrium path of play to changes in the beliefs for which it is sustained. The criterion explores how likely is it that the particular path of play is still played in equilibrium if beliefs change, and then selects the equilibrium path of play most likely to be played in a new equilibrium after the change of beliefs.

If a profile of actions a is played in equilibrium only under some narrowly specific beliefs (so only one or few equilibria result in playing a), while on the other hand action profile a' can be supported regardless of agents' off-path beliefs, or for a wide array of possible beliefs (so that lots of different equilibria result in playing a'), then we predict that in equilibrium, players will play according to a' , not a . In Section 2, we offer a simple example to illustrate our approach.

More generally, to identify the equilibrium actions that are most robust in the sense that they can be sustained in equilibrium by the largest set of beliefs, we

⁷As defined by Wolfram *Mathworld*, the Principle of Insufficient Reason states that "if we are ignorant of the ways an event can occur (and therefore have no reason to believe that one way will occur preferentially compared to another), the event will occur equally likely in any way."

proceed in three steps. First, we define a measure over sets of out of equilibrium beliefs. In line with the Principle of Insufficient Reason, we use the standard Lebesgue measure. Second, for each strategy profile, we identify the set of beliefs that support this strategy profile in equilibrium. For each action profile, we calculate the measure of the set of beliefs that support at least one equilibrium strategy profile in which this action profile is played. Finally, we select the action profile with the largest set of beliefs that support it as an equilibrium action profile. In any finite game, our criterion yields a non-empty prediction, a property that selection by passive, symmetric or wary beliefs fails to meet.

We compare the predictive power of several selection criteria, including selection by the largest of beliefs, using a laboratory experiment on a vertical multilateral contracting game.

We find that selection according to the largest set of beliefs has better predictive power than any of the other selection criteria. Over all treatments, 74% of groups of agents play some equilibrium, and among these, 94% play the equilibrium selected by the largest set of supporting beliefs, while only 54% play the equilibrium selected by passive or wary beliefs and only 46% play the equilibrium supported by symmetric beliefs. We find strong and significant comparative statics across treatments. These differences across treatments correspond perfectly with the comparative static predictions of our selection criterion, and clash with the predictions of each of the other equilibrium selection criteria.

Our results show that in a multilateral contracting game, our criterion correctly predicts which action profile is more likely to emerge. An important question concerns the external robustness of our criterion's predictions. Our selection criterion was originally designed to resolve the multiplicity of equilibria in a game of electoral competition over multiple districts. We tested our predictions for this electoral competition game in laboratory experiments. The criterion performed reasonably

well in this application, for which it was designed.⁸ The experiments and results on vertical contracting that we detail in Section 4 are a first test of the external validity and robustness of our criterion. Whether our criterion is similarly useful in other applications (beyond multi-district electoral competition and multi-lateral vertical contracting) remains an open question. We conjecture that it will perform better in cases in which the difference in the size of equilibrium beliefs that support each equilibrium is very large, and not so well if this difference is very small. Tests across a wide range of games and applications would best settle this question as to whether our conjecture holds.

A second central limitation of our study is that we do not demonstrate that individuals behave in line with the refinement we offer for the reasons for which the refinement is prescribed. That is, we do not provide a theory of belief formation underlying our refinement and explain which beliefs players should hold. Our theoretical contribution predicts how agents will play, but it does not explain why nor do we investigate why within our experiment. In any application in which there are good reasons to assume specific beliefs, researchers should use these beliefs and not our approach. We recommend using our selection criterion only as a second best, in applications in which researchers do not have a clear basis on which to theorize how subjects form out of equilibrium beliefs.

In what follows we first offer a motivating example. We then define our selection criterion and we present our experiments on vertical contracting. Appendixes A1-A5 contain proofs, additional results, a microfoundation of the experimental parameters, a discussion of alternative selection criteria, and experimental instructions.

⁸The experimental design, procedures and data analysis of this electoral competition experiment are available in the working paper version Eguia, Llorente-Saguer, Morton and Nicolò (2014).

2 A Motivating Example

Consider a simplified vertical contracting game with one supplier and two retailers. A supplier makes independent offers to each of two retailers $\{1, 2\}$. Let i denote an arbitrary retailer and let $-i$ denote the other retailer. The supplier offers to sell a product either at an exogenously given high price $p^H \in (15, 20)$, or a low price $p^L = 7$. Offers are simultaneously and privately made, so that each retailer observes the price offered to her, but not the price the other retailer is offered. Each retailer chooses an integer amount of units to purchase, which must be between 0 and 3. Let $q_i \in \{0, 1, 2, 3\}$ be the quantity purchased by retailer i . Retailers sell their units of the product in the consumer market. The unitary price of the product in the consumer market is strictly decreasing in the aggregate supply $Q = q_1 + q_2$ and is summarized in the following table:

Quantity Q	1	2	3	4	5	6
Price in the cons. market	21	20	15	14	10	6

Notice that if $Q > 5$ retailers' profits are negative if they buy any quantity from the supplier. Assume the supplier's marginal costs are equal to zero. Let $\Pi^{prod}(p_i, q_i)$ with $i \in \{1, 2\}$ denote the supplier's profits when she sells q_i units at price p_i to retailer i . Then $\Pi^{prod}(p^L, 3) > \Pi^{prod}(p^H, 1) > \Pi^{prod}(p^L, 2) > 0$.

Because a retailer's profits depend on the number of units purchased by the other retailer, contracts between the supplier and one retailer exert an externality on the other retailer. And since contracts with one retailer are not observed by the other, each retailer faces uncertainty about the expected payoff of purchasing any positive number of units from the supplier. Therefore, a retailer's willingness to purchase depends on beliefs about the trades executed between the supplier and the other retailer. It is easy to check that there exist only two symmetric pure Sequential Equilibria in this game. In one equilibrium the supplier offers p^L to both retailers and each retailer buys 2 units if the price offered is p^L and 0 units if the price is p^H .

In the other equilibrium the supplier offers p^H to both retailers and each retailer buys 1 unit if the price offered is p^H and 2 units if the price is p^L .

Off-path beliefs are critical to sustain one or the other equilibrium.⁹ Whether an equilibrium holds or not hinges on the belief that the other retailer was offered a high price p^H , following a deviation. If, following a deviation, the probability assigned to a high price for the other retailer is low enough, the equilibrium holds. Otherwise, it fails.

For each retailer $i \in \{1, 2\}$, let an out of equilibrium belief $\omega_i(p^H|\tilde{p}_i) \in [0, 1]$ be the probability that i assigns to the event that retailer $-i$ is offered p^H , when i receives an off-equilibrium offer \tilde{p}_i . The equilibrium with supplier's strategy (p^L, p^L) is supported by beliefs such that, after observing offer $\tilde{p}_i = p^H$, retailer i chooses not to purchase any unit, which implies that buying either one or two units¹⁰ gives negative profits. That is, buying one unit is not profitable:

$$\begin{aligned} \omega_i(p^H|\tilde{p}_i = p^H)(21 - p^H) + (1 - \omega_i(p^H|\tilde{p}_i = p^H))(15 - p^H) &\leq 0, \text{ or} \\ \omega_i(p^H|\tilde{p}_i = p^H) &\leq \frac{p^H - 15}{6}, \end{aligned} \quad (1)$$

and buying two units is also not profitable:

$$\begin{aligned} \omega_i(p^H|\tilde{p}_i = p^H)2(20 - p^H) + (1 - \omega_i(p^H|\tilde{p}_i = p^H))2(14 - p^H) &\leq 0, \text{ or} \\ \omega_i(p^H|\tilde{p}_i = p^H) &\leq \frac{p^H - 14}{6}. \end{aligned} \quad (2)$$

Constraint (2) is implied by Constraint (1).

The equilibrium with the pair of offers (p^H, p^H) is supported by beliefs such that, observing an off-equilibrium offer $\tilde{p}_i = p^L$, retailer i chooses to purchase two units

⁹There are also asymmetric equilibria in pure strategies, but for expositional simplicity here we focus on the two pure symmetric equilibria.

¹⁰Since $p^H > 15$, a retailer who observes an off-equilibrium offer p^H never buys more than two units.

and not three units,¹¹ which implies that

$$\begin{aligned} & \omega_i(p^H|\tilde{p}_i = p^L)2(15 - 7) + (1 - \omega_i(p^H|\tilde{p}_i = p^L))2(14 - 7) \\ & \geq \omega_i(p^H|\tilde{p}_i = p^L)3(14 - 7) + (1 - \omega_i(p^H|\tilde{p}_i = p^L))3(10 - 7); \\ & \omega_i(p^H|\tilde{p}_i = p^L) \leq 0.5. \end{aligned} \quad (3)$$

Which prediction is more reliable? Passive beliefs support the equilibrium with (p^L, p^L) because in this equilibrium, passive beliefs imply $\omega_i(p^H|\tilde{p}_1 = p^H) = 0$, but do not support the equilibrium with (p^H, p^H) because in this equilibrium passive beliefs imply $\omega_i(p^H|\tilde{p}_i = p^L) = 1$. Wary beliefs deliver the same prediction as passive beliefs. On the other hand, symmetric beliefs do not support the equilibrium with (p^L, p^L) because $\omega_i(p^H|\tilde{p}_1 = p^H) = 1$ but support the equilibrium with (p^H, p^H) because in this equilibrium symmetric beliefs imply $\omega_i(p^H|\tilde{p}_i = p^L) = 0$. Both passive and symmetric beliefs are extreme beliefs, and they each support a given equilibrium irrespective of the value of parameter p^H .

We argue that the focal equilibrium must depend on the value of parameter p^H . We expect that the equilibrium most likely to be played is the one sustained by the largest set of beliefs. The equilibrium in which the supplier's strategy profile is (p^H, p^H) is supported in equilibrium only if $\omega_i(p^H|\tilde{p}_i = p^L) \in [0, 0.5]$ for both $i \in \{1, 2\}$. Suppose that $p^H = 15.3$. The strategy profile (p^L, p^L) can be supported as a sequential equilibrium only if a retailer i who observes an off-equilibrium offer $\tilde{p}_i = p^H$ is almost sure that the other retailer is still offered p^L (and therefore refuses to buy). In particular, this strategy profile is supported in equilibrium only if $\omega_i(p^H|\tilde{p}_1 = p^H) \in [0, 0.05]$ for both $i \in \{1, 2\}$. Since $[0, 0.5]$ is a larger interval than $[0, 0.05]$, we argue that the equilibrium in which the supplier proposes (p^H, p^H) is a better prediction given $p^H = 15.3$. Suppose instead that $p^H = 19.8$; the supplier's strategy

¹¹It is easy to check that for a retailer i who observes an off-equilibrium offer $\tilde{p}_i = p^L$, buying two units is more profitable than buying one or zero units, irrespectively of her beliefs.

profile (p^L, p^L) is supported in equilibrium only if $\omega_i(p^H | \tilde{p}_1 = p^H) \in [0, 0.8]$ for both $i \in \{1, 2\}$: in this case most retailer's beliefs support an equilibrium with (p^L, p^L) , and we argue that then the equilibrium in which the supplier proposes (p^L, p^L) is more robust.

In the next section we formalize this intuition for a larger class of games.

3 A New Selection Criterion

Consider a class of finite extensive form games \mathcal{G} with one upstream player (proposer), and two or more downstream players (receivers).¹² Let 0 denote the proposer, and let R denote a set of receivers. The proposer moves first: 0 chooses an action $a_0 \in A_0$, where A_0 is a finite set of feasible actions for the proposer. For each receiver $i \in R$, let $\Pi_i(A_0)$ denote a partition of A_0 , and for any $a_0 \in A_0$, let $\Pi_i(a_0)$ denote the element of partition $\Pi_i(A_0)$ that contains a_0 . We assume that if the proposer chooses action a_0 , receiver i observes $\Pi_i(a_0)$. At the second and last stage, all receivers take simultaneous actions in response to the information they have observed. We do not impose any restrictions or assumptions (beyond finiteness) on the action set of receivers.

For any game $\Gamma \in \mathcal{G}$, let s_0 denote a pure strategy for the proposer, s_R a receivers' pure strategy profile, and $s = (s_0, s_R)$ a pure strategy profile. Let S_0 , S_R and $S = S_0 \times S_R$ be the set of all proposer's pure strategies, the set of all receivers' strategy profiles, and the set of all strategy profiles. Since the proposer's set of pure strategies is the set of actions available to the proposer, the proposer's strategy profile and proposer's action profile coincide. Let I be an information set in game Γ , after the proposer moves. Let \mathcal{I} be the collection of all information sets, so that $|\mathcal{I}|$ denotes the number of information sets in this collection, and for each $I \in \mathcal{I}$, let $i(I)$ be the player who moves at information set I . For each information set $I \in \mathcal{I}$, let $a(I)$ denote an action taken at I . For each information set $I \in \mathcal{I}$ and each strategy profile s , let $s(I)$ be the

¹²One can easily generalize the setup to two or more proposers. See the working paper version Eguia et al. (2014) for details.

action taken by $i(I)$ at I according to strategy profile s .

Given any proposer's action a_0 , let $\mathcal{I}(a_0) \subset \mathcal{I}$ be the collection of all information sets contained in the branches of the game tree that follow after the proposer plays a_0 . Let $\mathcal{O}(a_0) \subset \mathcal{I}$ be the collection of every information set I such that mover $i(I)$ observes that the proposer has not taken action a_0 . That is, $\mathcal{O}(a_0)$ represents the information sets in which for each $I \in \mathcal{O}(a_0)$, receiver $i(I)$ finds herself off-path from a_0 . Since the game Γ is finite, the size of $\mathcal{O}(a_0)$, denoted $|\mathcal{O}(a_0)|$, is also finite for any a_0 . At each information set I , let beliefs μ_I be a probability distribution over the nodes contained in I . Let $\mu = (\mu_1, \dots, \mu_{|\mathcal{I}|})$ be a list of such beliefs, one per information set.

A pure Sequential Equilibrium of game Γ is a pair (s, μ) that satisfies sequential rationality and belief consistency (Kreps and Wilson, 1982; or Fudenberg and Tirole, 1991, for a textbook treatment).

Given any proposer's action $a_0 \in A_0$, let a receivers' action profile be a list specifying the action taken by each receiver at each $I \in \mathcal{I}(a_0)$, and let this receivers' action profile be denoted $a_R^{a_0}$. We then refer to the pair $a = (a_0, a_R^{a_0})$ as an *action profile*. Given any Sequential Equilibrium strategy profile s^* with associated action profile a^* , beliefs along the path of equilibrium play are pinned down by equilibrium strategies and Bayes rule. Beliefs about the actions of other receivers at information sets that follow an observed deviation by the proposer are also pinned down by receivers' equilibrium strategies. There is, however, some indeterminacy in the beliefs held by receivers about the actions taken by the proposer who is observed to have deviated from s^* . Each information set $I \in \mathcal{O}(a_0^*)$, corresponds to an element other than $\Pi_i(a_0^*)$ of the partition $\Pi_i(A_0)$. For each $I \in \mathcal{O}(a_0^*)$, denote Π_i^I this corresponding element of the partition of A_0 .

Let Δ_I be the set of all probability distributions over the actions in Π_i^I and let $\omega_I \in \Delta_I$ be a probability distribution over the actions in Π_i^I . Set Δ_I is the set of all beliefs that $i(I)$ might hold at I about what the proposer did to make the game

unexpectedly reach information set I . Hereafter, we refer to ω_I as the belief *by i at I* . The standard “belief” μ_I over the nodes in information set I can be calculated directly from the belief ω_I .

Let $\Delta_{\mathcal{O}(a_0)} = \prod_{I \in \mathcal{O}(a_0)} \Delta_I$ be the set of all possible belief profiles over the collection of information sets $\mathcal{O}(a_0)$, and let $\omega = (\omega_1, \dots, \omega_{|\mathcal{O}(a_0)|}) \in \Delta_{\mathcal{O}(a_0)}$ be a belief profile, which specifies beliefs at each information set in $\mathcal{O}(a_0)$.

For each equilibrium strategy profile $s = (s_0, s_R)$, define $\Delta_{\mathcal{O}(a_0)}^s \subseteq \Delta_{\mathcal{O}(a_0)}$ as the set of off-path belief profiles such that a Sequential Equilibrium with strategy profile s and with belief profile ω exists if and only if $\omega \in \Delta_{\mathcal{O}(a_0)}^s$. Each equilibrium strategy profile s is associated with one equilibrium action profile a , which specifies the actions taken along the equilibrium path. Let A be the set of all possible action profiles.

For any action profile $a = (a_0, a_R^{a_0}) \in A$, let $S^a = \{s \in S : s_0 = a_0 \text{ and } s_R(I) = a_R^{a_0}(I) \text{ for each } I \in \mathcal{I}(a_0)\}$ be the subset of strategy profiles such that according to any strategy profile in S^a , agents play action profile a along the equilibrium path.

Let $\Delta_{\mathcal{O}(a_0)}^a = \bigcup_{s \in S^a} \Delta_{\mathcal{O}(a_0)}^s$. Then $\Delta_{\mathcal{O}(a_0)}^a$ is the set of belief profiles such that a Sequential Equilibrium in which agents play a and hold belief profile ω exists if and only if $\omega \in \Delta_{\mathcal{O}(a_0)}^a$.

Let L be the Lebesgue measure over $\Delta_{\mathcal{O}(a_0)}$, where $\Delta_{\mathcal{O}(a_0)}$ is the Cartesian product of $|\mathcal{O}(a_0)|$ simplexes and hence is itself a subset of a unit hypercube. Then, $L(\Delta_{\mathcal{O}(a_0)})$ represents the size of the set of belief profiles over the collection of information sets $\mathcal{O}(a_0)$, and $L(\Delta_{\mathcal{O}(a_0)}^a)$ is the size of the subset of these that support action profile a in a Sequential Equilibrium. The fraction $\frac{L(\Delta_{\mathcal{O}(a_0)}^a)}{L(\Delta_{\mathcal{O}(a_0)})}$ captures the relative size of the off-path belief profiles that support action profile a in a Sequential Equilibrium, over all belief profiles. We refer to this size as the “Size of Supporting Beliefs of a ” or $SSB(a)$.

Definition 1 *The size of off-path beliefs that support an action profile a in equilibrium is $SSB(a) = \frac{L(\Delta_{\mathcal{O}(a_0)}^a)}{L(\Delta_{\mathcal{O}(a_0)})}$. Equilibrium action profile a^* has a larger set of supporting beliefs than equilibrium action profile \hat{a}^* if $SSB(a^*) \geq SSB(\hat{a}^*)$. An equilibrium*

action profile a^* has a largest set of supporting beliefs if $SSB(a^*) \geq SSB(a)$ for any $a \in A$.

Our criterion allows for the ranking of all the equilibrium action profiles according to the size of the sets of supporting beliefs, from the action profile with the largest set of supporting beliefs to the one with the smallest set. Therefore, differently from standard selection criteria we are able to compare any pair of action profiles of equilibria and assess which is the one that is most likely to be played according to our criterion.¹³ We can also take a further step, relying on a cardinal interpretation of our criterion to assess the robustness of the proposed selection, by comparing the size of the largest set of supporting beliefs and the size of the second largest set in the ranking. For instance, in our introductory example in Section 2 if $p^H = 15.3$, the proposer's action profile of equilibrium (p^L, p^L) had a size of supporting beliefs equal to 0.05 while the other proposer's action profile of equilibrium (p^H, p^H) had a size of supporting beliefs equal to 0.5, which is ten times larger than the former one, and therefore in this game there is a substantial difference between the size of the two sets.

Notice also, that we predict an equilibrium action profile. We do not specify the particular equilibrium that supports this equilibrium action profile. We seek to explain agents' behavior and choices, and an action profile contains a full prediction over these.

The motivation for this selection criterion is agnostic about beliefs off the equilibrium path. Selection criteria based on passive, symmetric or wary beliefs assume that off-path beliefs take a particular form, and discard any equilibria not supported by these particular beliefs. But the focality of these beliefs to the exclusion of all others is often difficult to justify, as we discussed in the introduction. We take a more

¹³Note that there can be more than one equilibrium action profile with a largest set of beliefs. In Appendix A1, however, we show that the equilibrium supported by a largest set of beliefs is generally unique over a collection of perturbed games. Hence, we refer to a^* as the equilibrium action profile with *the* largest set of beliefs.

open-minded approach toward off-path beliefs. Following traditional pure game theory, we conjecture that agents may have any beliefs off the equilibrium path. Unable to predict which off equilibrium beliefs agents hold, we use a uniform prior over all possible beliefs.

In some applications, we may have a-priori knowledge about which beliefs are more likely to emerge. In these cases, we can adapt our selection criterion, replacing the Lebesgue measure by a measure that assigns weights to beliefs according to the probability distribution that best fits the a-priori knowledge. Passive, wary and symmetric beliefs can be interpreted as an extreme case of this approach, in which the measure on beliefs assigns all weight to a specific belief. Absent good a-priori knowledge over off-path beliefs, the measure to use under the veil of ignorance should be the Lebesgue measure.

The purpose of any equilibrium refinement or selection criterion is to solve the problem of multiplicity under a preferred solution concept by imposing additional restrictions to yield a sharper, ideally unique, prediction. An essential property of a useful refinement or selection criterion is that the criterion must make a non-empty selection. Unfortunately, selection by passive or symmetric beliefs fail this basic requirement even for finite games: in some finite games with pure strategy equilibria, requiring beliefs to be passive, or to be symmetric, eliminates all pure strategy equilibria (see Rey and Tirole, 2007, for lack of existence of pure equilibria with passive beliefs).¹⁴ In fact, in applications with capacity constraints, passive or symmetric beliefs are untenable, as they would imply a belief in non-feasible actions (see Avenel, 2012a and 2012b).

In contrast, given a non-empty set of pure equilibria in a finite game, our selection

¹⁴For lack of existence with symmetric beliefs, consider a game in which the proposer privately makes an offer A or B to each of two receivers, and receivers accept or reject. Let the proposer's preferred outcome be that one receiver accepts A and the other B , the second best that both accept A , and the third best that both offers be rejected; and let receivers' preferred outcome be to both accept A , the second best to both accept B , and the third best to both reject their offers. The only pure equilibrium is for both to accept A , and it is not sustained by symmetric beliefs.

criterion always relies on feasible beliefs, and it always makes a non-empty selection.

Proposition 1 *Given any game $\Gamma \in \mathcal{G}$ with a pure Sequential equilibrium, there exists an equilibrium action profile with the largest set of supporting beliefs.*

Selecting the equilibrium action profile supported by the largest set of beliefs guarantees that we have a well-defined prediction, while selecting using passive or symmetric beliefs does not.

In addition to existence of a non-empty prediction, sharpness is a desirable property of a selection criterion. In most games, selection by the size of supporting beliefs yields a unique prediction; the subset of equilibrium action profiles with a largest set of supporting beliefs is a singleton. An exception are games in which two or more equilibrium action profiles are supported by all off-path beliefs: in this class of games, all the equilibria that are supported by any beliefs appear to be equally compelling, and they are both selected according to passive beliefs, symmetric beliefs, wary beliefs, the largest set of beliefs, or any other belief-based notion: there is no belief-based reason to favor an equilibrium action profile over the other. However, in a vast collection of games, each equilibrium holds for some beliefs and not for others. Within this class of games, selection by the largest set of beliefs obtains uniqueness generically (while selection using passive or symmetric beliefs does not attain generic uniqueness). We formalize and prove this claim in Appendix A1.

We also note that our selection criterion is robust (i.e. delivers an invariant prediction) if we relabel agents or actions, and it is also invariant to affine transformations of the payoff matrix. Being a refinement of Sequential Equilibrium, our selection criterion is sensitive to changes in the extensive form tree resulting from the addition of strategically superfluous moves, the elimination of dominated strategies, or the addition of cloned (payoff-equivalent) strategies.

Having established that our selection criterion better satisfies desirable theoretical properties (existence and uniqueness), in the next section we test how its predictive power compares to that of other alternatives in the main application of interest.

4 Experimental Evidence

4.1 Experimental Design and Procedures

We use controlled laboratory experiments to evaluate the predictive power of the different selection criteria. We implemented two different treatments, denoted by treatment L and treatment H , which constitute a variation of the motivating example presented in Section 2. Unlike in the motivating example, the games implemented in each treatment allow for only two pure equilibrium profiles.

All participants were given the role of either a supplier or a retailer, and kept that role throughout the experiment. In each period, suppliers made independent offers to each of two retailers. An offer took the form of one of two prices: a low price p^L or a high price p^H . Because we preferred a more parsimonious game, we constrained the strategy set of retailers. If they were offered p^H , retailers could buy either zero or one units; and if the price was p^L , they could purchase either zero, two or three units. Suppliers' payoffs were kept fixed across treatments, while retailers' payoffs varied across treatments. These payoffs are summarized in tables 1 and 2. All these payoffs are expressed in talers, the experimental currency.

Participants played the game for 50 rounds, being re-matched after every round within matching groups of 12 subjects. After each round, subjects received full feedback concerning the actions of all subjects in their subgroup and their payoffs for that round. To determine payment, the computer randomly selected five periods for the final payment. The total amount earned in these periods was transformed into euros through the conversion rate of 0.03 in Treatment L and 0.045 in Treatment H . In total, subjects earned an average of €12.87, including a show-up fee of €4. Each experimental session lasted approximately one hour.

Experiments were conducted at the BonnEconLab of the University of Bonn in March 2013. We ran a total of 6 sessions with 24 subjects each. No subject participated in more than one session. Therefore, we have six independent observations per

Treatment LQuantity bought
by the other retailer

	0	1	2	3
Quantity bought	33	33	33	33
1	40	39	3	2
2	116	44	42	28
3	64	61	40	25

Treatment HQuantity bought
by the other retailer

	0	1	2	3
Quantity bought	28	28	28	28
1	79	64	18	4
2	175	83	55	19
3	127	85	31	4

Table 1: Retailer's payoffs. Payoffs depend on prices offered by the proposer, but this is captured by the different action spaces.

		0	1	2	3
Price Charged by the supplier	High	0	28	–	–
	Low	0	–	20	40

Table 2: Supplier's payoffs for each transaction.

treatment. Students were recruited through the online recruitment system ORSEE (Greiner, 2004) and the experiment was programmed and conducted with the software z-Tree (Fischbacher, 2007).

All experimental sessions were organized using the same procedure. Subjects received detailed written instructions, which an experimenter read aloud (see Appendix A5). Subjects were asked to answer a quiz to check their full understanding of the experimental design before beginning the experiment. At the end of the experiment, subjects completed a short survey which we describe in Appendix A4.

Predictions. As we show in Claim 3 in Appendix A2, games in treatments L and H allow for two pure Sequential Equilibrium action profiles. In the first equilibrium action profile, the supplier offers p^L to both retailers, and each retailer purchases two units. We refer to this action profile $((p^L, p^L), (2, 2))$ as the L equilibrium action profile. The strategy profile that sustains this action profile in equilibrium is such that retailers would not purchase any unit if offered the high price, because they both mutually fear that the other retailer would still purchase two units at a low price. The supplier is thus stuck in a low price equilibrium (which gives him lower profits), being unable to credibly deviate to offer the high price to both retailers.

We denote by H the second equilibrium action profile $((p^H, p^H), (1, 1))$, in which the supplier offers p^H to both retailers, and each retailer purchases one unit. The strategy profile that sustains action profile H in equilibrium is such that retailers who observe a deviation to a low price do not believe that the offer is exclusive to them: they fear the other retailer received it as well, and thus that quantities in the market will be large, which limits their willingness to purchase to only two units, thwarting the supplier's incentive to deviate in the first place.

As we show in Appendices A2 and A3, the predictions by symmetric beliefs, passive beliefs and proper equilibrium are constant across treatments. Symmetric predict action profile H while passive beliefs and proper equilibrium predict action profile L . Unlike the previous selection criteria, the action profile with largest set of

supporting beliefs varies across treatments.

Our selection criterion predicts equilibrium *action profile L in Treatment L* and equilibrium *action profile H in Treatment H*.¹⁵

4.2 Experimental Results

We organize our discussion of the experimental results by focusing, in turn, on prices, quantities and goodness of fit. Appendix A4 has additional data on out-of-equilibrium beliefs.

4.2.1 Prices

Figure 1 plots the aggregate prices offered by suppliers in both treatments, aggregated by groups of five periods. Each bar is divided into three tones of gray: the dark gray part represents the frequency of suppliers who offered both retailers the high price, the medium gray part represents the frequency of suppliers who offered a high price for one retailer and a low price for the other, and the light gray part represents the frequency of suppliers who offered the low price to both retailers.

[FIGURE 1 AROUND HERE]

Caption: Aggregate prices in both treatments, aggregated by groups of five periods.

The figure makes two points quite clearly. First, there is a stark contrast between the pricing strategies used in both treatments: although in Treatment *L* prices are high only in 12% of the cases, in Treatment *H* they are high in 72.21% of the cases. This difference is clearly significant (Mann-Whitney test, $z = 2.882$, $p = 0.0039$).¹⁶ This difference across treatments is in line with the prediction of the largest set of

¹⁵In Treatment *L*, the size of the set of beliefs that support action profile *L* in equilibrium is 0.81, and the size of the set of beliefs that support *H* is 0.10. In Treatment *H*, this size is 0.16 for action profile *L* and 0.93 for action profile *H*.

¹⁶In all nonparametric tests we used a matching group as an independent observation, because from period 2 onwards, individual choices were affected by observing other group members. Unless otherwise noted, we aggregated the data across all periods in a matching group.

	Coef.	Std. Err.	Odds Ratio	z	Pr > $ z $
Treatment L	-1.66	0.58	0.19	-2.86	0.00
Period	0.07	0.01	1.07	13.57	0.00
Treatment L * Period	-0.10	0.01	0.90	-12.58	0.00
Cut 1	-0.39	0.40		-0.96	0.34
Cut 2	0.95	0.40		2.35	0.02
Group variance	0.49	0.38			
Subject variance	1.52	0.44			
Number of Observations			2400		
Number of Groups			12		
Number of Subjects			48		
Log Likelihood			-1450.58		
Wald χ^2			255.56		
Hausman test χ^2 *			0.00		

Table 3: Three-level (overall, group, and subject) mixed effects ordered logistic regression of high prices offered by supplies as a function of treatment, period, and treatment and period interacted. Treatment L is a dummy variable that takes value one if the treatment is L and zero otherwise. * The Hausman test is done on a linear mirror version of the model.

beliefs but not by the competing selection criteria. Second, although this difference is evident from the first periods, it increases over time due to opposite convergence processes. In the last five periods, the percentage of high prices in Treatment L is as low as 8.33%, while in Treatment H is as high as 91.67%.

Table 3 presents a general evaluation of our predictions in a three-level (overall, matching group, and subject levels) mixed-effects ordered logistic regression of the amount of high prices set by suppliers as a function of treatment and the period number, and both interacted.¹⁷ Table 3 indicates the strong treatment effect found above and the convergence pattern observed in Figure 1.

4.2.2 Quantities

Figure 2 displays the aggregate quantities bought conditional on the price in each treatment in groups of five periods. The graphs on the top (bottom) belong to

¹⁷We ran an ordered logit regression since the number of high prices is discrete (0, 1 or 2) and they are ordered. A linear probability and multinomial logit model yield similar results.

Treatment L (H). The graphs on the left (right) correspond to situations in which retailers were offered a high (low) price. Each bar is divided into different tones of gray which represent the amount bought: dark gray if the retailer bought zero units, dark/medium for one, light/medium for two and light gray for three.

[FIGURE 2 AROUND HERE]

Caption: Quantities bought conditional on the price in each treatment.

Figure 2 shows two clear patterns. First, retailers' behavior when receiving a low price is in line with the predictions of all equilibria and similar across treatments: they buy two units 89.96% of the time in Treatment L and 91.90% in Treatment H . The small difference across treatments is not significant (Mann-Whitney test, $z = 0.961$, $p = 0.3367$). Second, there is a substantial difference in retailers' behavior across treatments when receiving a high price. In Treatment L , retailers demand one unit 24.65% of the time. In contrast, this percentage is 93.48% in Treatment H . This difference in percentages is not surprisingly highly significant (Mann-Whitney test, $z = 2.882$, $p = 0.0039$). As we found with prices, this significant difference across treatments is in line with the predictions of behavior given by our largest set of beliefs criterion but not with the predictions from passive, symmetric or wary beliefs criteria.

4.2.3 Goodness of Fit

The main goal of the experimental part of this paper is to evaluate the predictive power of the different selection criteria described in the previous section. In this subsection we address this question by comparing a measure of goodness of fit. The measure of fit that we use is the percentage of observations in which the entire group behaved as predicted by a given equilibrium conditional on playing an equilibrium. Table 4 displays the results for the second half of the experiment as well as the

	All Data	L	H
Passive beliefs / Wary beliefs / Proper Eq.	54.10	99.57	10.43
Symmetric beliefs	45.90	00.43	89.56
Largest set of beliefs	94.46	99.57	89.56
Random	50.00	50.00	50.00
% Equilibrium Play	73.92	76.33	71.50

Table 4: Measure of fit of the different selection criteria. "Random" displays a measure of fit of a randomly picked equilibrium. The data is restricted to the second half of the experiment.

percentage of groups that played an equilibrium action profile.¹⁸

Overall, we find that our selection criterion outperforms both the common prediction of passive/wary beliefs and proper equilibrium, as well as the prediction of symmetric beliefs. While our selection criterion makes the right prediction in 94.46% of cases, passive/wary beliefs or proper equilibrium make the right prediction in 54.10% of cases, and symmetric beliefs in 45.90%. These differences are significant (Wilcoxon test, $z = 1.819$, $p = 0.0690$ and $z = 2.411$, $p = 0.0159$).

When we disaggregate by the different games, the predictive power of our selection criterion equalizes the best of the other criteria. Recall that in L , the prediction of our criterion coincides with the prediction of the equilibrium under passive and wary beliefs or proper equilibrium, while in H , the prediction of our criterion coincides with the prediction of the equilibrium under symmetric beliefs. Therefore, our selection criterion coincides with one of the others in each game by construction. The noticeable feature is that our criterion matches in each case the best performer of the other criteria.

¹⁸We restrict the measure to the second half of the experiment due to the convergence process observed in Section 4.2.1. Qualitatively similar results are obtained when considering the whole sample.

5 Discussion

In the absence of a theory that explains how agents form out of equilibrium beliefs, any “ad hoc” assumption on out of equilibrium beliefs is arbitrary. The out of equilibrium beliefs held by the players depend on the specific game they play: games with different characteristics lead to different patterns of beliefs. If we cannot predict with confidence which beliefs are salient in a given game, we recommend that we consider all out of equilibrium beliefs as equally likely.

Under a uniform prior over out of equilibrium beliefs, not all equilibrium action profiles are equally likely to emerge in the equilibrium actually played: those that require specific beliefs are less likely to be played than action profiles that hold in equilibrium for a large set of beliefs. The equilibrium action profile that is more likely to be played is the one which is supported by the largest set of beliefs. We select this equilibrium action profile.

To compute the size of beliefs that support a strategy profile s in equilibrium, we must construct a measure over sets of beliefs. In finite games, we use the standard Lebesgue measure defined over the set of all possible beliefs at each out of equilibrium information set that follows an individual deviation.

Our criterion selects the action profile that has the largest set of supporting beliefs. The confidence in our prediction is increasing in the ratio of the largest set of supporting beliefs, over the size of the set of beliefs that sustain other action profiles. If the largest set is little larger than others, we conjecture that the equilibrium action profile with the largest set of beliefs may not always be played, but that nevertheless the frequency of play of each equilibrium action profile that is played will be increasing in the size of beliefs that support it in equilibrium. If our reasoning is correct, we can construct an order or ranking of equilibrium action profiles from most to least likely to be played, where each action profile is ranked according to the size of the set of beliefs that support it in equilibrium.

In an experiment on vertical contracting games, Martin, Normann and Sny-

der (2001) find that no specific restriction on the set of beliefs fits the data. Beyond this negative finding, it is difficult to conduct an equilibrium analysis with their data, because in their experiment, no group of players plays according to any equilibrium. Since all their subject groups are out of equilibrium, comparisons of equilibrium selection predictions are problematic.¹⁹ Martin, Normann and Snyder (2001) suggest that subjects' behavior can be explained by a (non-equilibrium) model in which some subjects hold passive beliefs and others hold symmetric beliefs. Such an explanation lacks predictive power, because ex-ante we don't know which fraction of subjects holds each beliefs. Our selection criterion offers a sharp prediction, based solely on the characteristics of the game and payoff matrix, and requiring no information about subjects' beliefs.

Our theory can be directly generalized to larger classes of finite sequential games, including games with more than two stages or with a more complex information structure. Our intuition –select the equilibrium supported by the largest set of beliefs– also applies to continuous games. While a precise formalization poses some technical challenges to define a measure of beliefs over a continuous strategy space, a formal definition that overcomes these challenges is available in the online appendix.

We argue that in applications in which we cannot predict players' out of equilibrium beliefs, we should not assume an arbitrarily chosen set of beliefs. Rather, we should accept that different beliefs might emerge. Our experiments identify one application to vertical contracting in which assuming passive, symmetric or wary beliefs is unwarranted: the predictive power of any of these criteria is not better than selecting equilibria at random. In contrast, selection by the largest set of beliefs predicts the right equilibrium in 94% of cases in which an equilibrium is played.

¹⁹Their experiment is aimed at testing different theories of vertical foreclosure, not at testing different selection criteria.

References

- [1] Arya, Anil, and Brian Mittendorf. 2011. “Disclosure standards for vertical contracts.” *The RAND Journal of Economics* 42(3): 595-617.
- [2] Avenel, Eric. 2012a. “Upstream capacity constraint and the preservation of monopoly power in private bilateral contracting.” *Journal of Industrial Economics* 60(4): 578-598.
- [3] Avenel, Eric. 2012b. “Suppliers’ merger and consumers’ welfare.” *Ekonomia* 15(1): 1-21.
- [4] Banks, Jeffrey, and Joel Sobel. 1987. “Equilibrium selection in signalling games.” *Econometrica* 55: 647-662.
- [5] Bernoulli, Jakob. 1713. *Ars Conjectandi*. Thurneysen Brothers Press, Basel.
- [6] Bar-Isaac, Heski, Guillermo Caruana and Vicente Cuñat, 2012. “Search, design and market structure.” *American Economic Review* 102 (2): 1140-1160.
- [7] Bernard, Andrew B., and Swati Dhingra. 2015. “Contracting and the division of the gains from trade.” NBER wp 21691.
- [8] Boone, Jan, Wieland Müller, and Sigrid Suetens. 2014. “Naked exclusion in the lab: the case of sequential contracting.” *Journal of Industrial Economics* 62(1): 137-166.
- [9] Brunnermeier, Markus K., and Martin Oehmke. 2013. “The maturity rat race.” *Journal of Finance* 68(2): 483-521.
- [10] Buehler, Benno, and Florian Schuett. 2014. “Certification and minimum quality standards when some consumers are uninformed.” *European Economic Review* 70: 493-511.
- [11] Cabrales, Antonio, Walter Garcia-Fontes and Massimo Motta. 2000. “Risk dominance selects the leader: an experimental analysis.” *International Journal of Industrial Organization* 18: 137-162.
- [12] Caprice, Stéphane. 2006. “Multilateral vertical contracting with an alternative supply: the welfare effects of a ban on price discrimination.” *Review of Industrial Organization* 28(1): 63-80.
- [13] Caprice, Stéphane, and Patrick Rey. 2015. “Buyer power from joint listing decision.” *Economic Journal* 125: 1677-1704.
- [14] Cho, In-Koo, and David M. Kreps. 1987. “Signaling games and stable equilibria.” *Quarterly Journal of Economics* 102(2): 179-222.
- [15] Croson, Rachel T.A. 2000. “Thinking like a game theorist: factors affecting the frequency of equilibrium play.” *Journal of Economic Behavior & Organization* 41(3): 299-314.

- [16] Eguia, Jon X., Aniol Llorente-Saguer, Rebecca Morton and Antonio Nicolò. 2014. “Equilibrium selection in sequential games with imperfect information.” Max Planck Institute for Research on Collective Goods wp 2014-04.
- [17] Fischbacher, Urs. 2007. “z-Tree - Zurich toolbox for ready-made economic experiments.” *Experimental Economics* 10: 171–178.
- [18] de Fontenay, Catherine C., and Joshua S. Gans. 2005. “Vertical integration in the presence of upstream competition.” *RAND Journal of Economics* 36(3): 544–572.
- [19] de Fontenay, Catherine C., and Joshua S. Gans. 2014. “Bilateral bargaining with externalities.” *Journal of Industrial Economics* 62(4): 756–788.
- [20] Fudenberg, Drew, and Jean Tirole. 1991. *Game Theory*. MIT Press, Cambridge.
- [21] Gächter, Simon, and Elke Renner. 2010. “The effects of (incentivized) belief elicitation in public goods experiments.” *Experimental Economics* 13(3): 364–377.
- [22] Gavazza, Alessandro, and Alessandro Lizzeri. 2009. “Transparency and economic policy.” *Review of Economic Studies* 76(3): 1023–1048.
- [23] Govindan, Srihari, and Robert Wilson. 2009. “On forward induction.” *Econometrica* 77(1): 1–28.
- [24] Greiner, Ben. 2004. “An online recruitment system for economic experiments,” in *Forschung und Wissenschaftliches Rechnen 2003*, ed. by K. Kremer and V. Macho. GWDG Bericht 63. Gesellschaft für Wissenschaftliche Datenverarbeitung Göttingen: Datenverarbeitung, 79–93.
- [25] Hagiui, Andrei, and Hanna Halaburda. 2014. “Information and two-sided platform profits.” *International Journal of Industrial Organization* 34: 25–35.
- [26] Harsanyi, John C., and Reinhard Selten. 1988. *A General Theory of Equilibrium Selection in Games*. MIT Press.
- [27] Harsanyi, John C. 1995. “A new theory of equilibrium selection for games with complete information.” *Games and Economic Behavior* 8: 91–122.
- [28] Hart, Oliver, and Jean Tirole. 1990. “Vertical integration and market foreclosure.” *Brookings Papers on Economic Activity*, Special Issue, 205–276.
- [29] Ho, Katherine. 2009. “Insurer-provider networks in the medical care market.” *American Economic Review* 99(1): 393–430.
- [30] In, YOUNGHWAN, and Julian Wright. 2017. “Signaling private choices.” *Review of Economic Studies*, forthcoming.
- [31] Inderst, Roman, and Marco Ottaviani. 2012. “Competition through commissions and kickbacks.” *American Economic Review* 102(2): 780–809.
- [32] Janssen, Maarten, and Sandro Shelegia. 2014. “Beliefs, market size and consumer search.” U. Vienna Economics wp. 1501.

- [33] Kreps, David M., and Robert Wilson. 1982. "Sequential equilibria." *Econometrica* 50: 863–894.
- [34] Laplace, Pierre S. 1820. *Theorie analytique des probabilités*. Courcier, Paris.
- [35] Martin, Stephen, Hans-Theo Normann and Christopher M. Snyder. 2001. "Vertical foreclosure in experimental markets." *RAND Journal of Economics* 32(3): 466–496.
- [36] Mayzlin, Dina, and Hema Yoganarasimhan. 2012. "Link to success: How blogs build an audience by promoting rivals." *Management Science* 58(9): 1651–1668.
- [37] McAfee, R. Preston, and Marius Schwartz. 1994. "Opportunism in multilateral vertical contracting: nondiscrimination, exclusivity, and uniformity." *American Economic Review* 84: 210–230.
- [38] Miklós-Thal, Jeanine, and Greg Shaffer. 2016. "Naked exclusion with private offers." *American Economic Journal: Microeconomics* 8(4): 174–194.
- [39] Möllers, Claudia, Hans-Theo Normann and Christopher M. Snyder. 2017. "Communication in vertical markets: Experimental evidence." *International Journal of Industrial Organization* 50: 214–258.
- [40] Myerson, Roger B. 1978. "Refinements of the Nash equilibrium concept." *International Journal of Game Theory* 15:133–154.
- [41] Nocke, Volker, and Patrick Rey. 2014. "Exclusive dealing and vertical integration in interlocking relationships." *CEPR Discussion Paper No. DP10176*.
- [42] O'Brien, Daniel P., and Greg Shaffer. 1992. "Vertical control with bilateral contracts." *RAND Journal of Economics* 23(3): 299–308.
- [43] Pagnozzi, Marco, and Salvatore Piccolo. 2011. "Vertical separation with private contracts." *Economic Journal* 122: 173–207.
- [44] Peski, Marcin. 2010. "Generalized risk-dominance and asymmetric dynamics." *Journal of Economic Theory* 145: 216–248.
- [45] Reisinger, Markus, and Emmanuele Tarantino. 2015. "Vertical integration, foreclosure, and productive efficiency." *RAND Journal of Economics* 46: 461–479.
- [46] Rey, Patrick, and Jean Tirole. 2007. "A primer on foreclosure," in *Handbook of Industrial Organization* (eds. M. Armstrong and R. Porter), 3: 2145–2220.
- [47] Rey, Patrick, and Thibaud Vergé. 2004. "Bilateral control with vertical contracts." *RAND Journal of Economics* 35(4): 728–746.
- [48] Schmidt, David, Robert Shupp, James M. Walker and Elinor Ostrom. 2003. "Playing safe in coordination games: the roles of risk dominance, payoff dominance, and history of play." *Games and Economic Behavior* 42: 281–299.
- [49] Segal, Ilya. 1999. "Contracting with externalities." *Quarterly Journal of Economics* 114: 337–388.

- [50] Segal, Ilya, and Michael D. Whinston. 2003. "Robust predictions for bilateral contracting with externalities." *Econometrica* 71(3): 757–791.
- [51] Selten, Reinhard. 1995. "An axiomatic theory of a risk dominance measure for bipolar games with linear incentives." *Games and Economic Behavior* 8: 213–263.

Appendices

Appendix A1. Proofs

Proof of Proposition 1. Given a finite game, there are only finitely many action profiles, and thus, only finitely many of them are equilibrium action profiles. Let k be the number of different action profiles that can be supported in equilibrium. For an arbitrary action profile a that can be supported in equilibrium, $L(\Delta_{\mathcal{O}(a_P)}^a)$ and $L(\Delta_{\mathcal{O}(a_P)})$ are well defined, and thus $\frac{L(\Delta_{\mathcal{O}(a_P)}^a)}{L(\Delta_{\mathcal{O}(a_P)})}$ and $SSB(a)$ are well defined as well. There are at most k different values of size SSB , and thus, there exists a maximum among them. ■

Uniqueness of the prediction

In Section 3 we claim that selection by the largest set of supporting beliefs generically yields a unique prediction if no more than one equilibrium action profile is supported by every off-path belief. We provide a formal statement and proof of this claim.

Let G be a game form, which indicates a set of agents N with $n = 1 + n_R$ players (one proposer and n_R receivers), a set of feasible strategy profiles $S = S_1 \times \dots \times S_n$ of size $|S|$, and the timing and information structure of a collection of games in \mathcal{G} , without specifying the payoffs. Let $U \in \mathbb{R}^{|A| \times n}$ define the payoff for each agent, for each possible action profile. Then $\Gamma = (G, U) \in \mathcal{G}$ defines a specific game. For any $\varepsilon > 0$, let $\mathcal{P}_\varepsilon(U)$ be the set of possible perturbations of U such that the proposer's payoff stays unaltered, and the payoff of each receiver is altered by no more than ε . Formally, let $U_a \in \mathbb{R}^n$ denote the payoff vector for any action profile $a \in A$, and let U_{ak} denote the k -th component of vector U_a . Then, for any $\varepsilon \geq 0$, $\mathcal{P}_\varepsilon(U) = \{U' \in \mathbb{R}^{|A| \times n} : |U_{ak} - U'_{ak}| \leq \varepsilon \text{ for any } k \in \{2, \dots, n_R + 1\}, \text{ for any } a \in A\}$. Let $\mathcal{P}_\varepsilon(\Gamma) = (G, \mathcal{P}_\varepsilon(U))$ be the collection of games with the same set of pure strategy equilibria as game Γ that can be generated by perturbing game Γ according to perturbations in $\mathcal{P}_\varepsilon(U)$.

Proposition 2 *Assume action profiles a and a' can be supported by a largest set of beliefs in game $\Gamma = (G, U) \in \mathcal{G}$ and $SSB(a) = SSB(a') \in (0, 1)$. Then there exists $\varepsilon > 0$ such that the equilibrium action profile with a largest set of supporting beliefs is generically unique over the class of games $\mathcal{P}_\varepsilon(\Gamma)$.*

Proof. Assume $SSB(a) = SSB(a') \in (0, 1)$ in game $\Gamma = (G, U)$. That is, in game Γ , $\frac{L(\Delta_{\mathcal{O}(a_P)}^a)}{L(\Delta_{\mathcal{O}(a_P)})} \in (0, 1)$. Then there exists $\bar{\varepsilon} > 0$ such that for any $\varepsilon \in (0, \bar{\varepsilon}]$ and any $\tilde{U} \in \mathcal{P}_\varepsilon(U)$, in game $\tilde{\Gamma} = (G, \tilde{U})$ we also find that $\frac{L(\Delta_{\mathcal{O}(a_P)}^a)}{L(\Delta_{\mathcal{O}(a_P)})} \in (0, 1)$ and thus $SSB(a) \in (0, 1)$. By an analogous reasoning, in game $\tilde{\Gamma} = (G, \tilde{U})$, $SSB(a') \in (0, 1)$. Consider any perturbation $\tilde{U} \in \mathcal{P}_\varepsilon(U)$ and let \tilde{U}^λ be the family of perturbations such that for any λ , $\tilde{U}_{a''k}^\lambda = \tilde{U}_{a''k}$ for

any agent k and any action $a'' \notin \{a, a'\}$; $\tilde{U}_{ak}^\lambda = \tilde{U}_{ak} - \lambda$; and $\tilde{U}_{a'k}^\lambda = \tilde{U}_{a'k} + \lambda$. Then $\Delta_{\mathcal{O}(a_P)}^a$ strictly decreases with λ , and $\Delta_{\mathcal{O}(s'_P)}^{a'}$ strictly increases with λ , while $\Delta_{\mathcal{O}(a_P)}$ and $\Delta_{\mathcal{O}(s'_P)}$ do not depend on λ . Thus, there is at most one value of λ for which $\frac{L(\Delta_{\mathcal{O}(a_P)}^a)}{L(\Delta_{\mathcal{O}(a_P)})} = \frac{L(\Delta_{\mathcal{O}(s'_P)}^{a'})}{L(\Delta_{\mathcal{O}(s'_P)})}$ and hence $SSB(a) = SSB(a')$ is a non-generic event in the family of games perturbed by \tilde{U}^λ . ■

Appendix A2. Microfoundation of the Experimental Parameters

In this section we present a microfoundation for the games in the experimental section. This variation of the motivating example presented in Section 2 allows for only two pure equilibria. The game is adapted from Segal (1999) and Rey and Tirole (2007) (see as well Hart and Tirole, 1990; McAfee and Schwartz, 1994; or Rey and Vergé, 2004, among others).

An upstream firm (supplier) labeled 0, makes independent offers to each of two retailers $\{1, 2\}$. The supplier sells a good. An offer takes the form of a price. Let $p_1, p_2 \in \{p^H, p^L\}$ be the prices offered by the supplier to retailers 1 and 2, where $p^H \in \mathbb{R}_{++}$ and $p^L \in \mathbb{R}_{++}$ are an exogenously given high and low price, respectively. The strategy set of the supplier is $S_0 = \{(p^H, p^H), (p^H, p^L), (p^L, p^H), (p^L, p^L)\}$, where the first component indicates the offer made to retailer 1 and the second the offer made to retailer 2. Offers are simultaneously and privately made, so that each retailer observes the offer she receives, but not the offer the other retailer receives. Retailers choose how many units to purchase. Because we prefer a more parsimonious game, we constrain the strategy set of retailers. If the price is p^H , we assume that retailers can buy zero or one units; and if the price is p^L , zero, two or three units. Hence the strategy set for each retailer i is $S_i = \{0, 1\} \times \{0, 2, 3\}$, where each strategy $(s_{iH}, s_{iL}) \in S_i$ corresponds to how much to purchase following a p^H offer (first coordinate) and following a p^L offer (second coordinate). A strategy profile is an element of $\{p^H, p^L\}^2 \times (\{0, 1\} \times \{0, 2, 3\})^2$.

The supplier incurs a transaction cost $c_0 \in \mathbb{R}_+$ for each executed trade, and each retailer incurs a transaction cost $c_r \in \mathbb{R}_+$ if she accepts an offer. Let $q_i \in \{0, 1, 2, 3\}$ be the quantity purchased by retailer i . Retailers sell their units of the good in the consumer market. The aggregate supply is $Q = q_1 + q_2$. If $Q > 0$, the price of the good in the consumer market is a function $p : \{1, 2, 3, 4, 5, 6\} \rightarrow \mathbb{R}_+$ of the aggregate supply Q . We assume function p is strictly decreasing in Q . Profits for the supplier and retailer i are $\sum_{i=1}^2 (p_i q_i - c_0 \mathbf{1}[q_i > 0])$ and $(p(Q) - p_i) q_i - c_r \mathbf{1}[q_i > 0]$ respectively, where $\mathbf{1}[\cdot]$ is the indicator function; if $Q = 0$ profits are zero for all players.

The game form contains ten parameters. Let $\theta = (p^H, p^L, c_0, c_r, p(1), p(2), p(3), p(4), p(5), p(6))$ be a parameter vector. We are interested in parameter values that generate multiple (two) equilibria in which different selection criteria generate different predictions.

Definition 2 *Let Θ be the set of parameter vectors such that $\theta \in \Theta$ if and only if*

i) $3p^L > p^H > 2p^L > c_0$, and

ii) for each retailer i , the best response correspondences $BR_i(p_i, q_{-i})$ are given by the following table:

	Quantity q_{-i}			
	0	1	2	3
$BR_i(p^H, q_{-i})$	1	1	0	0
$BR_i(p^L, q_{-i})$	2	3	2	0

According to the first condition, selling one unit at a high price is better for the supplier than selling two at a low price, but not as good as selling three at a low price (and any transaction is better than no transaction). The second condition on retailers' best responses leads to the existence of two equilibria: one in which both retailers purchase one unit at price p^H , and another in which both retailers purchase two units at price p^L . No other outcome can be sustained in a pure equilibrium.

For any parameter vector θ , let Γ_θ be the game with the specified game form, and parameter values θ . Let an action profile sustained in equilibrium (or, simply, an equilibrium action profile), be the action played by the supplier according to the equilibrium strategy profile, and the actions played by the retailers along the equilibrium path, that is, in case the supplier does not deviate from her equilibrium strategy.

Claim 3 For any $\theta \in \Theta$, the set of action profiles that are sustained in a pure Sequential Equilibrium of game Γ_θ is $\{((p^L, p^L), 2, 2), ((p^H, p^H), 1, 1)\}$.

Proof. Consider an arbitrary $\theta \in \Theta$ and the corresponding game Γ_θ ; we first prove that there are two pure strategy profiles that are supported in a Sequential Equilibrium of the game. Consider first strategy profile $((L, L), (0, 2), (0, 2))$. For each retailer to buy two units is the unique best response when the price offered is low and the other retailer is buying two units. If the supplier deviates by proposing a high price to some retailer, then she decreases her payoff because the retailer buys zero units and the deviation is not observed by the other retailer. When the price offered is high, to buy zero units is a best response for a retailer who assigns probability one that the supplier is proposing a low price to the other retailer. Thus $((L, L), (0, 2), (0, 2))$ satisfies sequential rationality. A totally mixed strategy profile that assigns probability $1 - \varepsilon$ to suppliers' strategy $((L, L), (0, 2), (0, 2))$ and probability $\frac{\varepsilon}{2} - \varepsilon^2$ to $((L, H), (0, 2), (0, 2))$ and $((H, L), (0, 2), (0, 2))$ and ε^2 to all other strategies combined, with each of these strategies receiving equal weight, converges as $\varepsilon \rightarrow 0$ to profile $((L, L), (0, 2), (0, 2))$ and the beliefs obtained by Bayesian updating for any $\varepsilon > 0$ converge to the beliefs specified above that a retailer observing an offer of a high price believes the supplier offers a low price to the other retailer with probability one. Thus the equilibrium satisfies belief consistency.

Consider the strategy profile $((H, H), (1, 2), (1, 2))$. For each retailer to buy one unit is the unique best response when the price is high and the other retailer is buying one unit. If the supplier deviates proposing a low price to any retailer, she decreases her payoff

because the retailer buys two units and by assumption $p^H > 2p^L$. To buy two units when the price offered is low is a best response for a retailer who assigns probability one that the supplier is proposing a low price to the other retailer, too. Thus $((H, H), (0, 2), (0, 2))$ satisfies sequential rationality. A totally mixed strategy profile that assigns probability $1 - \varepsilon$ to suppliers' strategy $((H, H), (1, 2), (1, 2))$ and probability $\varepsilon - \varepsilon^2$ to $((L, L), (1, 2), (1, 2))$ and ε^2 to all other strategies combined, with each of these strategies receiving equal weight, converges as $\varepsilon \rightarrow 0$ to profile $((H, H), (1, 2), (1, 2))$ and the beliefs obtained by Bayesian updating for any $\varepsilon > 0$ converge to the beliefs specified above that a retailer observing an offer of a low price believes the supplier offers a low price to the other retailer with probability one. Thus the equilibrium satisfies belief consistency.

Consider now any other strategy profile. First, notice that if the supplier proposes a high price to retailer $i = 1, 2$, retailer i buys a positive amount if and only if retailer $j \neq i$ buys at most one unit. If retailer i buys at most one unit (as he is forced to do if the price offered to him is high), retailer j 's best response is to buy two or three units when the price offered to him is low. It follows that if, in equilibrium, the supplier offers a high price to retailer i and a low price to retailer j , then retailer i buys zero unit, and retailer j buys two units. However, if the supplier deviates and offers a low price to retailer i , then retailer i 's best response is to buy a positive amount, irrespective of his beliefs about the price offered by the supplier to retailer j . Therefore the deviation is profitable for the supplier. It follows that there are no asymmetric equilibria such that the supplier offers a high price to a retailer and a low price to the other one. If the supplier offers a high price to both retailers, to buy one unit is the best response for each retailer irrespective of the amount that the other retailer buys. If the supplier offers a low price to both retailers, if retailer i buys three units, retailer j 's best response is to buy zero units. However, to buy three units for retailer i is not the best response when retailer j buys zero units. Hence, there are no equilibria such that the supplier offers a low price to both retailers and some retailer buys a quantity different than two. ■

The first equilibrium action profile is $((p^L, p^L), 2, 2)$, that is, the supplier offers a low price to both retailers, and each retailer purchases two units. We refer to this action profile as the L equilibrium action profile. The strategy profile that sustains this action profile in equilibrium is such that retailers would not purchase any unit if offered the high price, because they both mutually fear that the other retailer would still purchase two units at a low price. The supplier is thus stuck in a low price equilibrium (which gives him lower profits), being unable to credibly deviate to offer the high price to both retailers.

In the second equilibrium action profile, denoted H , prices are high, quantities low. The supplier could be tempted to deviate by offering a low price to a retailer, who would then benefit from purchasing three units, which would increase the supplier's profit. However, the strategy profile that sustains action profile H in equilibrium is such that retailers who observe a deviation to a low price do not believe that the offer is exclusive to them: they fear the other retailer received it as well, and thus that quantities in the market will be

large, which limits their willingness to purchase to only two units, thwarting the supplier's incentive to deviate in the first place.

Claim 4 For any $\theta \in \Theta$, $((p^L, p^L), 2, 2)$ is the unique action profile supported by **passive beliefs** in a pure Sequential equilibrium of game Γ_θ , and is also the unique action profile supported by wary beliefs; whereas $((p^H, p^H), 1, 1)$ is the unique action profile supported by **symmetric beliefs**.

Proof. This claim follows immediately from the proof of Claim 3 and the definitions of passive beliefs, and symmetric beliefs. ■

The action profile $((p^L, p^L), 2, 2)$ is also the unique action profile sustained in a **proper equilibrium** (Claim 5 in Appendix A3).

Whether an equilibrium holds or not hinges on the belief that the other retailer was offered a high price p^H . Let $\omega_i(p^H|\tilde{p}_i)$, be the probability assigned by retailer i with beliefs ω_i to $p_{-i} = p^H$ after observing \tilde{p}_i .

Equilibrium action profile L is supported by beliefs such that, after observing offer $\tilde{p}_i = p^H$, retailer i chooses not to purchase any units, which implies that

$$\begin{aligned} \omega_i(p^H|\tilde{p}_i = p^H)p(1) + (1 - \omega_i(p^H|\tilde{p}_i = p^H))p(3) - p^H - c_r &\leq 0 \\ \omega_i(p^H|\tilde{p}_i = p^H) &\leq \frac{p^H + c_r - p(3)}{p(1) - p(3)}. \end{aligned} \quad (4)$$

Equilibrium action profile H is supported by beliefs such that, observing offer $\tilde{p}_i = p^L$, retailer i chooses to purchase two units, and not three, which implies that

$$\begin{aligned} 2[\omega_i(p^H|\tilde{p}_i = p^L)p(3) + (1 - \omega_i(p^H|\tilde{p}_i = p^L))p(4) - p^L] - c_r \\ \geq 3[\omega_i(p^H|\tilde{p}_i = p^L)p(4) + (1 - \omega_i(p^H|\tilde{p}_i = p^L))p(5) - p^L] - c_r \\ \omega_i(p^H|\tilde{p}_i = p^L) &\leq \frac{p^L - 3p(5) + 2p(4)}{5p(4) - 2p(3) - 3p(5)}. \end{aligned} \quad (5)$$

Let

$$\begin{aligned} \Theta_L &= \left\{ \theta \in \Theta : \frac{p^H + c_r - p(3)}{p(1) - p(3)} > \frac{p^L - 3p(5) + 2p(4)}{5p(4) - 2p(3) - 3p(5)} \right\} \text{ and} \\ \Theta_H &= \left\{ \theta \in \Theta : \frac{p^H + c_r - p(3)}{p(1) - p(3)} < \frac{p^L - 3p(5) + 2p(4)}{5p(4) - 2p(3) - 3p(5)} \right\}. \end{aligned}$$

The hyperplane $\frac{p^H + c_r - p(3)}{p(1) - p(3)} = \frac{p^L - 3p(5) + 2p(4)}{5p(4) - 2p(3) - 3p(5)}$ divides the set of parameter vectors Θ into two regions denoted by Θ_L and Θ_H . In the region Θ_L a larger set of beliefs sustains the equilibrium L and in the region Θ_H a larger set of beliefs sustains the equilibrium H . Notice that $\Theta - (\Theta_L \cup \Theta_H)$ has measure zero; the equilibrium supported by the largest set of beliefs is generically unique over the set of parameters. We obtain the following prediction:

Prediction *If $\theta \in \Theta_L$, we predict that subjects will play the L equilibrium, if $\theta \in \Theta_H$, we predict that subjects will play the H equilibrium, based on the size of the set of beliefs that support each equilibrium.*

The two treatments presented in Section 4.1 use different parameter constellations, with $\theta_L \in \Theta_L$ in treatment L and $\theta_H \in \Theta_H$ in treatment H . Unlike in the motivating example, the parameters related to supplier payoffs are kept fixed across the two games: supplier prices are $p^H = 36$, $p^L = 15$ and the supplier's transaction cost for any executed trade is $c_0 = 15$. The two treatments differ, however, in the variables that affect retailers' payoffs. In Treatment L , the retailer's transaction cost for any positive purchase is $c_r = 29$ and the vector of consumer market prices is $(72, 71, 35, 34, 27, 22)$, where coordinate k denotes the price if k units are sold in the market, for any number of units $k \in \{1, 2, 3, 4, 5, 6\}$. In Treatment H , the retailer's transaction cost for any positive purchase is $c_r = 33$, and the vector of consumer market prices is $(120, 105, 59, 45, 27, 18)$.

In Treatment L with parameter values θ_L , the size of the set of beliefs that support action profile L in equilibrium is 0.81, and the size of the set of beliefs that support H is 0.10. In Treatment H with parameter values θ_H , this size is 0.16 for action profile L and 0.93 for action profile H .

We made an affine transformation to the payoff function, which has no strategic consequences, but yields two advantages for experimental purposes: we avoid negative payoffs to subjects by adding a constant, and we equalize supplier and retailers' expected payoffs by multiplying the supplier's payoffs by $4/3$. Tables 1 and 2 respectively summarize retailers' payoffs and supplier's payoffs.

Appendix A3. Alternative Selection Criteria

For completeness, we have also considered a number of other refinements that, to our knowledge, have not been previously used in our application of interest, but are prominent in other game theoretic research. Harsanyi and Selten (1988) propose payoff dominance and risk-dominance. Payoff dominance selects a Pareto superior equilibrium (one that yields higher payoffs to every player) over a Pareto inferior one. This criterion may seem plausible when it is applicable, which is only seldom.²⁰ Risk dominance, extended by Harsanyi (1995) and Selten (1995) selects the equilibrium that is least risky in the sense that each player minimizes the potential losses if she cannot anticipate which equilibrium will be played by other agents. Risk dominance is defined for normal form games and while it has performed well in laboratory experiments on coordination games (Cabrales, Garcia-Fontes and Motta, 2000; Schmidt, Shupp, Walker and Ostrom, 2003), it is not obvious how to apply the intuition underlying this concept to games with imperfect information in a compelling manner. Peski (2010) introduces two variations of risk dominance that apply

²⁰Payoff dominance only offers a prediction in one of the four treatments in our experiments, and when it does, it coincides with our prediction.

to any finite normal form game, but the intuition behind these two refinements does not extend well to sequential games and in our application, both refinements deliver an empty prediction. Forward induction as defined by Govindan and Wilson (2009) is applicable and makes a non-empty prediction. Alas, this prediction does not solve our multiplicity problem: all equilibria in the vertical contracting game we study satisfy forward induction.²¹

Myerson's (1978) proper equilibrium refinement is the only one of these concepts that is applicable and offers a sharper prediction in our games, as the next claim shows.

Claim 5 *For any $\theta \in \Theta$, the unique proper equilibrium of the vertical contracting game with parameter vector θ is such that agents play strategy profile $((p^L, p^L), (0, 2), (0, 2))$.*

Proof. For any $\theta \in \Theta$, the set of pure strategy profiles that are sustained in a Sequential Equilibrium of game Γ_θ is $\{((p^L, p^L), (0, 2), (0, 2)), ((p^H, p^H), (1, 2), (1, 2))\}$. Consider first the equilibrium $((p^H, p^H), (1, 2), (1, 2))$. A Sequential Equilibrium in which $(q_1, q_2) = (1, 1)$ is supported by strategy profile $((p^H, p^H), (1, 2), (1, 2))$ and beliefs such that retailer i observing $p_i = p^L$ believes with sufficiently high probability that $p_{-i} = p^L$ and thus expects $q_{-i} = 2$. However the deviation (consistent with $p_i = p^L$) that provides the highest payoff to the supplier is such that the supplier is still playing the strategy $p_{-i} = p^H$ and therefore in any proper equilibrium, the retailer should assign a probability which gives a higher weight (of the order of $\frac{1}{\varepsilon}$ for $\varepsilon > 0$ infinitesimally small) to strategies according to which the deviating proposer plays $p_{-i} = p^H$. However, the best response according to these beliefs is such that when i observes $p_i = p^L$, i should buy three units and not two units as prescribed by the equilibrium, making the deviation profitable for the supplier. Therefore this equilibrium strategy profile cannot be supported as a proper equilibrium.

Consider the equilibrium $((p^L, p^L), (0, 2), (0, 2))$. A Sequential Equilibrium in which $(q_1, q_2) = (2, 2)$ is supported by strategy profile $((p^L, p^L), (0, 2), (0, 2))$ and beliefs such that retailer i observing $p_i = p^H$ believes with sufficiently high probability that $p_{-i} = p^L$ and thus expects $q_{-i} = 2$. Since the deviation (consistent with $p_i = p^H$) that provides the highest payoff to the supplier is such that the supplier is still playing the strategy $p_{-i} = p^L$, then this strategy profile is supported by beliefs that satisfy the ε -proper equilibrium refinement. ■

Despite the sharper prediction, the predictive power of proper equilibrium is underwhelming: its predictions coincide with those of selection by wary beliefs, and thus it underperforms relative to equilibrium selection based on identifying the strategy profile that can be supported in equilibrium by the largest set of beliefs. Finally, in a more recent research paper, In and Wright (2017) propose a new equilibrium concept to refine the set of Sequential equilibria in a class of senders-receivers games. They name this concept "RI-equilibrium." As shown by In and Wright (2017), Section 5.4., in games of vertical contracting with one proposer and two receivers—such as those in our treatments—the RI-equilibrium coincides with the equilibrium with wary beliefs and with the proper equilibrium. More generally, the RI-equilibria are a superset of the proper equilibria). Therefore,

²¹A proof of this statement is available in the online appendix.

		Other Retailer's price if offered a high price			
		Treatment L		Treatment H	
		High	Low	High	Low
Other Retailer's price if offered a low price	High	1	2	17	4
	Low	15 ^{<i>S</i>}	30 ^{<i>P</i>}	20 ^{<i>S</i>}	7 ^{<i>P</i>}

Table 5: Joint distribution of beliefs (in absolute numbers). ^{*P*} indicates the prediction by the passive beliefs and ^{*S*} indicates the prediction by symmetric beliefs.

the RI-equilibrium does not offer as good a prediction as our selection criterion in our experiment.

Appendix A4. Beliefs

Beliefs are a crucial element of equilibrium selection in the games presented. In order to assess whether game play was related to participants' beliefs, we elicited beliefs of retailers in a non-incentivized manner at the end of the experiment.²² Immediately after finishing the main part of the experiment, subjects completed a questionnaire. The two first questions related to their beliefs. In particular, we asked the following questions: "Suppose that you play an additional period as a retailer. If the supplier offers you a *low* / *high* price, which price do you think the supplier will offer to the other retailer?" They could either answer "Low Price" or "High Price."

Table 5 shows the joint distributions of beliefs. The matrix shows a substantial difference between treatments. As indicated in the table, according to the equilibrium with symmetric beliefs, a retailer i would expect the retailer j to be offered a low price when i is offered a low price, and a high price when i is offered a high price. In the equilibrium with passive beliefs, retailers expect the other retailer to be offered a low price regardless of the price offered to themselves.

Note that in Treatment L , most retailers' beliefs are in line with the beliefs predicted by passive beliefs although in Treatment H , instead, most retailers' beliefs are in line with the prediction of symmetric beliefs. These findings reinforce the results shown in the previous section: a selection criterion that imposes an invariant assumption on beliefs regardless of the nature of the game is inappropriately restrictive, and a poor fit of the data. Off-path beliefs depend on the particular game.

²²Because eliciting beliefs can produce significant changes to the actions played (Croson, 2000, and Gächter and Renner, 2010), we elicited them only at the end of the experiment, providing subjects with instructions about belief elicitation only after the contracting game experiment was over.

Appendix A5. Instructions for Treatment H

(The instructions for Treatment L were exactly the same except for the payoff matrix for retailers and the numbers in the explanation of the payoffs' table)

Thank you for taking part in this experiment. Please read these instructions very carefully. It is important that you do not talk to other participants during the entire experiment. In case you do not understand some parts of the experiment, please read through these instructions again. If you have further questions after hearing the instructions, please give us a sign by raising your hand out of your cubicle. We will then approach you in order to answer your questions personally. Please do not ask anything aloud.

During this experiment you will earn money. How much you earn depends partly on your own decisions, partly on the decisions of other participants, and partly on chance. Your personal earnings will be paid to you in cash as soon as the experiment is over. Your payoffs during the experiment will be indicated in talers. At the end of the instructions we are going to explain you how we are going to transform talers into euros.

After the experiment, we will ask you to complete a short questionnaire, which we need for the statistical analysis of the experimental data. The data of the questionnaire, as well as all your decisions during the experiments will be anonymous.

The experiment consists of 50 rounds. The rules are the same for all rounds. At the beginning of the experiment you will be assigned a role. One third of the participants will be assigned the role of suppliers and the other two thirds will be assigned the role of retailers. Given that there are 24 participants in the experiment, 8 participants will be assigned as suppliers and 16 will be assigned as retailers. Your role will be the same in all rounds.

At the beginning of the experiment, participants will be randomly and anonymously divided in two sets. Each set will consist of 4 suppliers and 8 retailers. These sets remain unaltered for the entire experiment. At no point are you told who is in your set. In each period, participants in each set are randomly matched with other participants of the same set into a subset, which we call a group. Each group will be composed of one supplier and two retailers. In a given round you will only interact with the participants in your group for that round.

Each round will be divided into two stages. In the first stage, the price setting stage, suppliers will decide the price that they will charge each retailer for their goods. In the second stage, the quantity setting stage, each retailer will decide the quantity to buy. In what follows we will explain to you the procedure in each round. Remember that 1) in each round you will only interact with the participants in your group and 2) in the beginning of each round you will be rematch to a new group within the same set.

Price Setting Stage. In this stage, only suppliers will have to make a decision. The supplier of each group will have to decide the price charged to each retailer of his/her group. Retailers will be called Retailer 1 and Retailer 2. These names will be randomly assigned each round. Suppliers can charge either a Low Price or a High Price. Therefore,

the suppliers have four possible decisions:

1. Charge a low price to both retailers.
2. Charge a low price to Retailer 1 and a high price to Retailer 2.
3. Charge a high price to Retailer 1 and a low price to Retailer 2.
4. Charge a high price to both retailers.

Quantity Setting Stage. In the quantity setting stage, only retailers will have to make a decision. Each retailer will have to decide the quantity to buy from the supplier of their group.

- If the supplier charged a low price, the retailer can buy either 0, 2 or 3 units.
- If the supplier charged a high price, the retailer can buy either 0 or 1 unit.

When deciding the quantity, retailers will not know the price offered to the other retailer of the group.

Payoffs for the Supplier. The supplier's payoff from sales to each of the retailers will depend on the price charged and on the quantity that the retailer decided to buy. Next table displays the payoff for each of the possible combinations of prices and quantities.

	Quantity Bought by the retailer			
	0	1	2	3
Low Price	0	–	20	40
High Price	0	28	–	–

In other words:

- If the price is Low, your payoff (as a supplier) is
 - 0, if the retailer buys 0 units
 - 20, if the retailer buys 2 units
 - 40, if the retailer buys 3 units
- If the price is High, your payoff (as a supplier) is
 - 0, if the retailer buys 0 units
 - 28, if the retailer buys 1 unit

The total payoff of a supplier is equal to the sum of payoffs from sales to each of the retailers.

Payoffs for the Retailers. The payoff of a retailer depends on the quantity bought by the retailer and the quantity bought by the other retailer of his group. Your payoff is summarized in the following table:

		Quantity Bought by the other retailer			
		0	1	2	3
Quantity Bought by you	0	28	28	28	28
	1	79	64	18	4
	2	175	83	55	19
	3	127	85	31	4

In other words:

- If you buy 0 unit and the other retailer buys 0, 1, 2 or 3 units, your payoff is 28.
- If you buy 1 unit, your payoff is:
 - 79, if the other retailer buys 0 units.
 - 64, if the other retailer buys 1 unit.
 - 18, if the other retailer buys 2 units.
 - 4, if the other retailer buys 3 units.
- If you buy 2 units, your payoff is:
 - 175, if the other retailer buys 0 units.
 - 83, if the other retailer buys 1 unit.
 - 55, if the other retailer buys 2 units.
 - 19, if the other retailer buys 3 units.
- If you buy 3 units, your payoff is:
 - 127, if the other retailer buys 0 units.
 - 85, if the other retailer buys 1 unit.
 - 31, if the other retailer buys 2 units.
 - 4, if the other retailer buys 3 units.

Remember that if you have been offered a high price you can only buy zero or one units, and that if you have been offered a low price, then you can buy either zero, two or three units. The same holds for the other retailer.

Information at the end of each Round. At the end of each round, all participants will receive the following information:

- (i) the prices offered to each retailer
- (ii) the quantities bought by each retailer
- (iii) your payoff

Final Earnings. At the end of the experiment, the computer will randomly select 4 rounds and you will earn the payoffs you obtained in these rounds. Each of the 50 rounds has the same chance of being selected. The total number of points accumulated in these 5 selected rounds will be transformed into euros by multiplying your earnings in talers by a conversion rate. For this experiment the conversion rate is 0.075, meaning that 100 talers

equal 7.5 Euros. Additionally; you will earn a show-up fee of 4.00€. Everyone will be paid in private and you are under no obligation to tell others how much you earned.

Control Questions. Before starting this part, you will have to fill in some control questions in the computer terminal. Click the button Ok after you have answered a question to move to the next question. In case you answer wrongly, a message will pop out with the right answer and you will have to answer it again. Once you and all the other participants have filled all the questions the experiment will start.

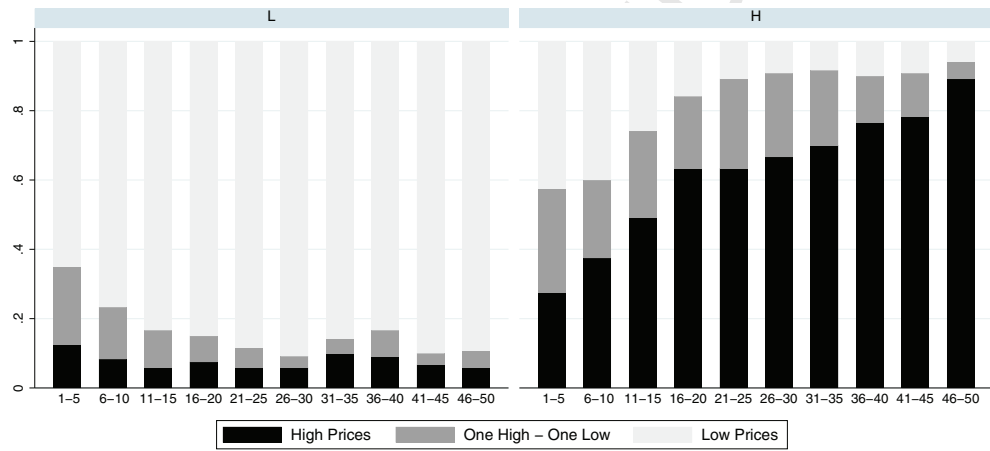


Figure1

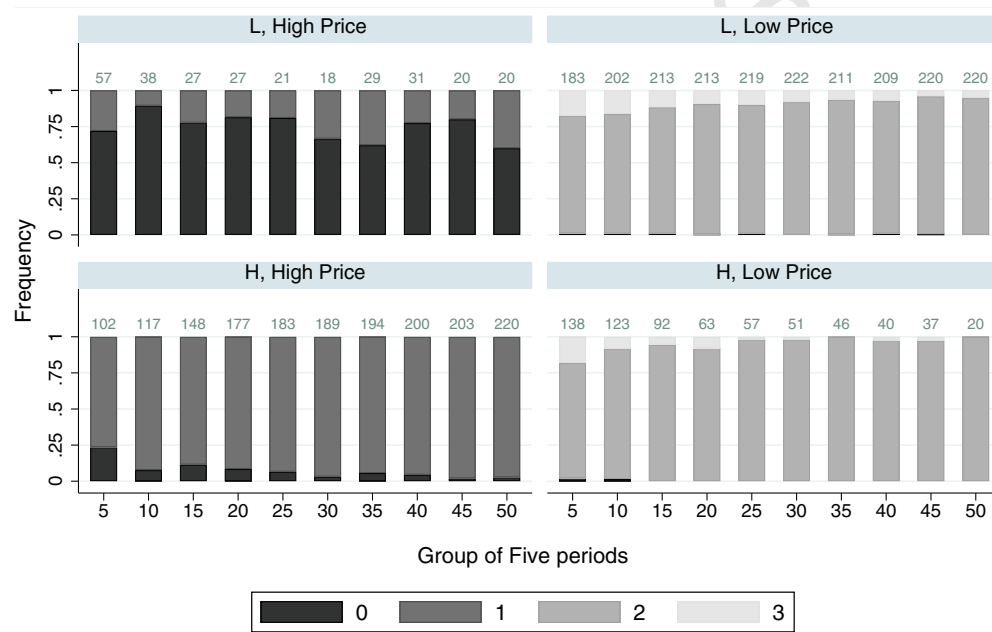


Figure2