



**Equilibrium Selection in
Sequential Games with
Imperfect Information**

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





Equilibrium Selection in Sequential Games with Imperfect Information

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

March 2014

Equilibrium Selection in Sequential Games with Imperfect Information*

Jon X. Eguia[†] Aniol Llorente-Saguer[‡] Rebecca Morton[§] Antonio Nicolò[¶]

March 24, 2014

Abstract

Games with imperfect information often feature multiple equilibria, which depend on beliefs off the equilibrium path. Standard selection criteria such as *passive beliefs*, *symmetric beliefs* 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 experiments to test the predictive power of the existing and our novel selection criteria in two applications: a game of vertical multi-lateral contracting, and a game of electoral competition. We find that our selection criterion outperforms the other selection criteria.

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

JEL Codes: C72, D86, H41, D72.

*We thank Antonio Cabrales, Christoph Engel, Ignacio Esponda, Kristoffel Grechenig, PJ Healy, Asen Ivanov, Matthew Jackson, Isabel Marcin, Marco Ottaviani, Nicolas Roux, Andrew Schotter, Chris Tyson, Piercarlo Zanchettin, seminar attendants at NYU-CESS, Carlos III, Venice and DICE, and participants at ASSET 2013 and the IV Workshop on Institutions, Individual Behavior and Economic Outcomes at Porto Conte for their helpful suggestions. We would also like to thank Dominic Land, Nicolas Meier and Maximilian Pohl for excellent assistance at running the experiments. We acknowledge financial support from NYU-NYC, NYU-AD and Università di Padova.

[†]Department of Economics, Univ. of Bristol. and Department of Politics, New York Univ. Email (corresponding author): jon.egua@bristol.ac.uk.

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

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

[¶]School of Social Sciences, Univ. of Manchester. Email: antonio.nicolo@manchester.co.uk.

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). These games typically feature multiple Perfect Bayesian and sequential equilibria. Equilibria depend on how, after observing a deviation, receivers update their beliefs about proposers' actions. Refinements that are useful for signaling games such as the intuitive criterion (Cho and Kreps (7)), D1 or universal divinity (Banks and Sobel (2)) have no bite in this context, because the lack of information is about the proposers' actions, not about their type (the game is one of imperfect information, not incomplete information; players' types are known).

Our main application of interest in this class of games is multi-lateral vertical contracting. 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 receives, it does not know what offers other firms have received (downstream firms operate under imperfect information).

The industrial organization literature on multi-lateral vertical contracting has dealt with this multiplicity of equilibria by imposing particular beliefs that are deemed more appealing off the equilibrium path, selecting equilibria that can be supported by these beliefs, and discarding all other equilibria. McAfee and Schwartz (22) propose three possible beliefs to consider: passive beliefs, symmetric beliefs, and wary beliefs.

Passive beliefs, sometimes called "passive conjectures" (Rey and Tirole (27)), are such that a downstream firm that receives an out of equilibrium offer does not update her beliefs on the offers received by all other players; rather, she believes that all the other unobserved actions remain as in equilibrium. The selection criterion based on selecting equilibria that can be supported by such passive beliefs is the one most frequently used in the literature, not only in games of multilateral vertical contracting (Hart and Tirole (16); O'Brien and Shaffer (23); Segal (30); Fontenay and Gans (10); Rey and Tirole (27); Inderst and Ottaviani (18)) but also in games of electoral competition in which two candidates make offers to voters

in multiple districts (Gavazza and Lizzeri (12)). The assumption of “passive beliefs” is also common in the consumer search literature: for instance consumers do not revise their beliefs about a firm’s quality when it charges an unexpected price (see for instance recent applications by Bar-Isaac, Caruana and Cuñat (4) or Buehler and Schuett (5)). However, “in many circumstances the ad hoc restriction to passive beliefs may not be compelling” (Segal and Whinston (31)). Indeed, while defending the assumption of passive beliefs in the particular game in which they use it, Rey and Tirole (27) 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 that receives an out of equilibrium offer believes that all other downstream players must have received the same out of equilibrium offer as well. If an upstream player in equilibrium was anticipated to offer x to all downstream players, a downstream player who receives an offer of y believes that the offers to all other downstream players is also y . This selection criterion is used, among others, by Pagnozzi and Piccolo (24). However, most of the literature seems to agree with McAfee and Schwartz’s (22) initial assessment: “[Symmetric] beliefs are not very compelling.”

A third suggestion is to consider equilibria supported by wary beliefs. These beliefs 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. This criterion has had scant following in the literature (Rey and Vergé (28), Avenel (1)).

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 very 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 not often appropriate. Martin, Normann and Snyder (21) 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.

A more cautious or modest approach is to accept that we cannot pin down the exact beliefs off the equilibrium path, beyond the restrictions given by standard refinements such as those implicit in a Perfect Bayesian or a sequential equilibrium (Kreps and Wilson (19)). Any assumption of specific equilibrium beliefs, even a plausible one, is ad hoc and it is difficult to justify it as superior to all others. In application of Bernoulli (3) and Laplace's (20) "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.¹

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 a larger set of different equilibrium beliefs. If we assume that each set of equilibrium beliefs is equally likely to emerge, the strategy profile most likely to emerge in equilibrium and the action profile most likely to be played along the equilibrium path are the ones that can be sustained by a largest variety of different equilibrium beliefs. This 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 still be played in a new equilibrium after the change of beliefs.

If a profile of actions a is played in equilibrium only under some very specific beliefs (very 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

¹As 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."

that lots of different equilibria result in playing a'), then we predict that in equilibrium, players will play according to a' , not a . For instance, in Section 3, we consider a vertical contracting game in which a producer can offer a product to each of two retailers at either a high or low price, and each retailer chooses how many units to buy. The producer can price-discriminate, and neither retailer knows the price offered to the other retailer. There are two equilibria in pure strategies in this game, one in which each retailer is offered a high price and buys one unit, and one in which each is offered a low price and buy two units. Both equilibria hold if a retailer who observes a deviation believes that the other retailer is offered a high price with a probability lower than a given threshold. This threshold depends on parameters of the games and differs across the two equilibria. For given specific parameters,² the equilibrium with high prices holds for any beliefs such that a retailer who observes a deviation (an unexpected offer of a low price) believes that the other retailer is offered a high price with probability in $[0, 0.92]$, while the equilibrium with low prices holds for any beliefs such that a retailer who observes a deviation believes that the other retailer is offered a high price with probability in $[0, 0.16]$. Since the first case encompasses a larger set of beliefs, we predict that the strategy profile played in equilibrium is the one in which the producer offers high prices and retailers buy one unit.

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

1. 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.
2. 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.
3. We select the action profile with the largest set of beliefs that support it as an equi-

²These parameters correspond to Treatment H in Section 3 (see Table 1).

librium action profile.

In games with finitely many equilibria, including any finite game, our criterion yields a non-empty prediction, a desirable property of any solution concept that selection by passive, symmetric or wary beliefs fail to meet.

Notice that equilibrium selection by the largest set of beliefs results in the selection of an equilibrium action profile and the strategy profiles in which it is played. The criterion, however, does not select the specific beliefs that accompany these strategy profiles in equilibrium. We remain agnostic and make no predictions about the specific equilibrium beliefs; selection by the largest set of beliefs predicts only the strategies and the actions played in equilibrium.³

Above and beyond the appealing theoretical properties of selection by the largest set of beliefs, we are interested in the predictive power of this criterion. Using laboratory experiments, we test our selection by the largest set of beliefs against selection by other criteria in two applications of interest: vertical multilateral contracting, and electoral competition over multiple districts. Our experiments consists in two very simple sequential games. Relative to the seminal work by Martin, Norman and Snyder (21), the simplicity of our games has the advantage that most groups of subjects play some equilibrium; whereas, not a single group of subjects play according to equilibrium in the complex games in the Martin, Norman and Snyder (21) experiments.⁴ We can therefore conduct a more meaningful analysis of equilibrium selection within the large collection of groups of subjects who play according to some equilibrium.

Overall we find that selection according to the largest set of beliefs has better predictive power than any of the other selection criteria. In the application to vertical contracting, in the aggregate 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

³We do not provide a theory of belief formation, nor do we assume that players regard all equilibrium beliefs as equally likely. Our interpretation is that players play one equilibrium, which they hold as certain and all others as entirely unlikely. But the analyst forecasting how a game will be played does not know how beliefs are formed, so a priori we treat all equilibrium beliefs as equally probable in order to make a prediction for the game.

⁴Their experiment is aimed at testing different theories of vertical foreclosure and the analysis of out of equilibrium beliefs and the relative performance of different selection criteria are not the goals of their paper.

only 54% play the equilibrium selected by passive or wary beliefs and only 46% play the equilibrium supported by symmetric beliefs. In the electoral competition application results are weaker: only 49% of groups play an equilibrium. Within these, selection according to the largest set of supporting beliefs (76% correct predictions) dominates selection by symmetric beliefs (4% correct predictions), but it does not predict significantly better or worse than selection by passive beliefs (78%) or by wary beliefs (67%). In each of the two applications we run two different treatments. We find strong and significant comparative statics across treatments in each of the two applications. 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.

For completeness, we have also considered a number of other refinements that, to our knowledge, have not been previously used in our applications of interest (vertical contracting and electoral competition), but are prominent in other game theoretic research. Harsanyi and Selten (14) 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.⁵ Risk dominance, extended by Harsanyi (15) and Selten (32) 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, Motta (6); Schmidt, Shupp, Walker and Ostrom (29)), it is not obvious how to apply the intuition underlying this concept to games with imperfect information in a compelling manner. Peski (25) introduces two variations of risk dominance that apply 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 (13) is applicable and makes a non-empty prediction. Alas, this prediction does not solve our multiplicity

⁵Payoff dominance only offers a prediction in one of the four treatments in our experiments, and when it does, it coincides with our prediction.

problem: all equilibria in the vertical contracting game we study satisfy forward induction.⁶ Myerson’s (21) Proper Equilibrium refinement is the only one of these concepts that is applicable and offers a sharper prediction in our games. However, its predictive power is underwhelming: its predictions coincide with those of selection by wary beliefs, and thus in both applications 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 (17) propose a new equilibrium concept to refine the set of Perfect Bayesian equilibria in a class of senders-receivers games. They name this concept “RI-equilibrium.” In the application to vertical contracting, the RI-equilibrium coincides with the equilibrium with wary beliefs and with the proper equilibrium (and more generally, the RI-equilibria are a superset of the proper equilibria), and therefore in our applications, the RI-equilibrium does not offer as good a prediction as our selection criterion.

In what follows we first define our selection criterion for a class of two-stage finite games, and then we provide the theory, empirical implications, and experimental tests in two applications of interest: vertical contracting, and electoral competition. An Appendix contains proofs to the theoretical propositions, and additional figures.

2 A New Selection Criterion

Consider a class of finite extensive form games \mathcal{G} with one or more upstream players (proposers), and two or more downstream players (receivers). Let P denote the set of proposers, R denote the set of receivers, $N = P \cup R$ denote the set of players, $n_P = |P|$ the number of proposers, and $n_R = |R|$ the number of receivers. Proposers move first, simultaneously: each proposer j chooses an action $a_j = (a_{j,1}, a_{j,2}, \dots, a_{j,n_R}) \in A_j \subset \mathbb{R}^{K \times n_R}$, where K is a strictly positive integer, $a_{j,i} \in \mathbb{R}^K$ is a vector for each $j \in P$ and $i \in R$, and A_j is a finite set of feasible actions. A proposer action a_j is a matrix with n_R vectors. We interpret each vector $a_{j,i} \in \mathbb{R}^K$ as an offer made to receiver i . Each receiver i observes each element of each matrix of offers with some probability in $[0, 1]$. These probabilities may differ across

⁶A proof of this statement is available from the authors.

elements of the matrix, they may be correlated across elements and across receivers, and they may be degenerate so that Nature plays no role and i observes some vectors but not others. We assume that there is at least one proposer j , one receiver i , and one element of a_j such that i observes this element of a_j with probability strictly less than one, so that the game is not one of perfect information. Our motivation is that receiver i is more likely to observe the offers made to her ($a_{j,i}$ by each proposer j) than the offers made to any other receiver, but we do not need to assume this as a restriction. At the second and last stage, all receivers take simultaneous actions in response to the information they have observed. We do not impose any restriction or assumption (beyond finiteness) on the action set of receivers.

For any finite extensive form game $\Gamma \in \mathcal{G}$, let s_P denote a proposers' pure strategy profile, s_R a receivers' pure strategy profile, and $s = (s_P, s_R)$ a pure strategy profile. Let s_i denote the strategy for agent i , and let s_{-i} denote a strategy profile for every other agent, so $s = (s_i, s_{-i})$. Let S_P , S_R and $S = S_P \times S_R$ be the set of all proposers' pure strategy profiles, the set of all receivers' strategy profiles, and the set of all strategy profiles. Let I be an information set in game Γ , 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 . Let $\mathcal{I}_P \equiv \{I \in \mathcal{I} : i(I) \in P\}$ be the collection of information sets in which the mover is a proposer. 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 action (or behavioral strategy) taken by $i(I)$ at I according to strategy profile s . Let a *proposers' action profile*, denoted a_P , be a list specifying the action taken at each information set in which the mover is a proposer. That is, a proposers' action profile is a list of all the actions taken by the set of proposers.

Given any proposers' action profile a_P , let $\mathcal{I}(a_P) \subset \mathcal{I}$ be the collection of all information sets contained in the branches of the game tree that follow after proposers play a_P . That is, each $I \in \mathcal{I}(a_P)$ is reached after proposers play a_P for some realization of Nature's moves, and no node outside $\mathcal{I}(a_P)$ is reached after proposers play a_P , regardless of Nature's moves.

Given a proposer's action profile a_P , let $\mathcal{O}(a_P) \subset \mathcal{I}$ be the collection of all information

sets that:

- i) cannot be reached if proposers play a_P , i.e. the intersection $\mathcal{O}(a_P) \cap \mathcal{I}(a_P)$ is empty,
- ii) are reached for some realization of Nature's moves if exactly one proposer j deviates from a_P , and
- iii) there are at least two actions a'_j and a''_j such that information set I can be reached if j plays a'_j or a''_j (hence receiver $i(I)$ must form beliefs about proposer j 's actions).

We interpret $\mathcal{O}(a_P)$ as the collection of information sets one step out of the path of a_P . Since the game Γ is finite, the size of $\mathcal{O}(a_P)$, denoted $|\mathcal{O}(a_P)|$, is also finite for any a_P . 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.

Since each proposer's set of pure strategies is the set of actions available to the proposer, the proposers' strategy profile and proposers' action profile coincide: given any candidate strategy profile s_P , the proposers' action profile is $a_P = s_P$. Given a proposers' action profile a_P , let a *receivers' action profile* be a list specifying the action to be taken at each information set $I \in \mathcal{I}(a_P)$, and let this receivers' action profile be denoted $a_R^{a_P}$. We then refer to the pair $a = (a_P, a_R^{a_P})$ as an *action profile*. Each equilibrium strategy profile s is associated with one equilibrium action profile a , which specifies the actions taken along the equilibrium paths. Notice, however, that if there is uncertainty, not all information sets along the equilibrium paths are reached along the path of play in every realization of the game, as Nature's resolution of uncertainty directs the game along one or another path of play. Let A be the set of all possible action profiles, and let $a \in A$ be an arbitrary action profile.

A pure **Perfect Bayesian Equilibrium** (PBE) of game Γ is a pair (s, μ) such that given μ and given s_{-i} , each player at each information set maximizes her expected utility playing according to s_i , and such that given s , beliefs μ are correct along the equilibrium path, and satisfy Bayes Rule and other consistency requirements off the equilibrium path (see, for instance, Fudenberg and Tirole (11) p. 331-332).

Given any Perfect Bayesian 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 one (or more) proposers are also pinned down by receiver's equilibrium strategies, as are the beliefs about the proposers that have not been observed to deviate from s^* . There is, however, some indeterminacy in the beliefs held by receivers about the actions taken by any proposer j who is observed to have deviated from s^* . For any information set $I \in \mathcal{O}(a_P^*)$, let j^I be the proposer who must deviate from a_P^* in order for I to be reached.⁷ Then A_{j^I} is the set of all feasible actions for j^I . Let $B_{j^I} \subset A_{j^I}$ be the subset of actions for proposer j^I such that information set I can be reached if every proposer in $P \setminus \{j^I\}$ plays according to a_P .

Let Δ_I be the set of all probability distributions over B_{j^I} and let $\omega_I \in \Delta_I$ be a probability distribution over B_{j^I} . Set Δ_I is the set of all beliefs that $i(I)$ might hold at I about what proposer j^I 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 $\Delta_{\mathcal{O}(a_P)} = \prod_{I \in \mathcal{O}(a_P)} \Delta_I$ be the set of all possible belief profiles over the collection of information sets $\mathcal{O}(a_P)$, and let $\omega = (\omega_1, \dots, \omega_{|\mathcal{O}(a_P)|}) \in \Delta_{\mathcal{O}(a_P)}$ be a belief profile, which specifies beliefs at each information set in $\mathcal{O}(a_P)$.

For each equilibrium strategy profile $s = (s_P, s_R)$ and for each information set $I \in \mathcal{O}(a_P)$, let $\Delta_I^s \subseteq \Delta_I$ be the set of beliefs of $i(I)$ such that a Perfect Bayesian Equilibrium with strategy profile s and with beliefs ω_I at each $I \in \mathcal{O}(a_P)$ exists if and only if $\omega_I \in \Delta_I^s$ for each $I \in \mathcal{O}(a_P)$. Then, $\prod_{I \in \mathcal{O}(a_P)} \Delta_I^s$ is the set of belief profiles that sustain a PBE with strategy profile s .

Let $\Delta_{\mathcal{O}(a_P)}^s = \prod_{I \in \mathcal{O}(a_P)} \Delta_I^s$. A PBE in which agents play s and hold belief profile ω exists if and only if $\omega \in \Delta_{\mathcal{O}(a_P)}^s$.

For any action profile $a = (a_P, a_R^{a_P})$, let $S^a = \{s \in S : s_P = a_P \text{ and } s_R(I) = a_R(I) \text{ for each } I \in \mathcal{I}(a_P)\}$ 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.

⁷If more than one proposer deviates, beliefs about each deviating proposer are determined independently. Therefore, a specification of beliefs in $\mathcal{O}(a_P^*)$ suffices to construct a complete belief system everywhere in \mathcal{I} .

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

Let L be the Lebesgue measure over $\Delta_{\mathcal{O}(a_P)}$, where $\Delta_{\mathcal{O}(a_P)}$ is the Cartesian product of $|\mathcal{O}(a_P)|$ simplexes and hence is itself a subset of a unit hypercube. Then, $L(\Delta_{\mathcal{O}(a_P)})$ represents the size of all the set of belief profiles over the collection of information sets $\mathcal{O}(a_P)$, and $L(\Delta_{\mathcal{O}(a_P)}^a)$ is the size of the subset of these that support action profile a in a PBE. The fraction $\frac{L(\Delta_{\mathcal{O}(a_P)}^a)}{L(\Delta_{\mathcal{O}(a_P)})}$ captures the relative size of the belief profiles that support action profile a in a Perfect Bayesian 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 beliefs that support an action profile a in equilibrium is $SSB(a) = \frac{L(\Delta_{\mathcal{O}(a_P)}^a)}{L(\Delta_{\mathcal{O}(a_P)})}$. An equilibrium action profile a^* has a largest set of supporting beliefs if $SSB(a^*) \geq SSB(a)$ for any $a \in A$.*

We refer to a^* as the equilibrium action profile with the largest set of beliefs.⁸ The motivation for this selection criterion is agnosticism 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 very particular beliefs. But the focality of these beliefs in exclusion of all others is often difficult to justify, as we discuss in the introduction. We take a more open-minded approach toward off-path beliefs. Following traditional pure game theory, we conjecture that agents may have any beliefs off the equilibrium path, and that we are unable to predict which of these beliefs agents will hold. If we have a uniform prior over these off equilibrium beliefs, any equilibrium may hold.

However, we highlight that if all equilibria are equally likely, some equilibrium strategy profiles or equilibrium action profile are more likely than others. An equilibrium is a pair (s, μ) , composed of a strategy profile and beliefs. If an equilibrium *action profile* a is played according to several different equilibrium *strategy profiles*, each of which holds for many

⁸Note that there can be more than one equilibrium action profiles with a largest set of beliefs. In the Appendix, however, we show that the equilibrium supported by a largest set of beliefs is most often unique. Hence, we refer to a^* as the equilibrium action profile with *the* largest set of beliefs.

different beliefs, action profile a holds in many equilibria. If action profile a' is played according to a few different strategy profiles, each of which holds for a few beliefs, a' holds in a few equilibria.

Consider the following algorithm: Randomly draw an action profile a from a uniform distribution over A . Then, randomly draw a profile of off-path beliefs ω from a uniform distribution over $\Delta_{\mathcal{O}(a_P)} = \prod_{I \in \mathcal{O}(a_P)} \Delta_I$. If a PBE with action profile a and beliefs ω exists, stop. If not, restart the algorithm drawing a new action profile. The action profile most likely to be the outcome of this algorithm is the one supported by the largest set of beliefs.

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 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 (27) for lack of existence of pure equilibria with passive beliefs).⁹ In contrast, given a non-empty set of pure equilibria in a finite game, our selection criterion always makes a non-empty selection.

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

More generally, in any game with a finite set of equilibria, selecting the equilibrium action profile supported by the largest set of beliefs guarantees delivery of a well-defined

⁹For 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 prefer an outcome in which one receiver accepts A and the other B to an outcome in which both accept A , and prefer both offers rejected to any other outcome; and let receivers prefer both to accept A more than to both accept B , and both to reject their offers more than any other combination. The only pure equilibrium is for both to accept A , and it is not sustained by symmetric beliefs.

prediction, while the alternative criteria of selection based on passive or symmetric beliefs do 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 equilibria 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 the Appendix.

Having established that our selection criterion better satisfies desirable theoretical properties (existence and uniqueness), in the next two sections we test how its predictive power compares to that of other alternatives in two applications of interest: vertical contracting, and electoral competition.

3 Application: Vertical Contracting

The main application to test our selection criterion is a theory of vertical contracting with imperfect information and externalities. The game form we present is adapted from Segal (30) and Rey and Tirole (27) (see as well Hart and Tirole (16); McAfee and Schwartz (22); or Rey and Vergé (28) among others).

3.1 Theory and Predictions

We present a game form of a simplified multi-lateral vertical contracting game that allows us to generate multiple equilibria and different predictions according to different selection

criteria for varying parameter values.

A proposer, interpreted as a producer or upstream firm, and labeled 0, makes independent offers to each of two receivers $\{1, 2\}$, interpreted as retailers or downstream firms. Let i denote an arbitrary retailer and let $-i$ denote the other retailer. The producer sells a good. An offer takes the form of a price, which is either high ($p^H \in \mathbb{R}_{++}$) or low ($p^L \in \mathbb{R}_{++}$). Let $p_1, p_2 \in \{p^H, p^L\}$ be the prices offered by the producer to retailers 1 and 2. The strategy set of the producer 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 producer incurs a transaction cost $c_0 \in \mathbb{R}_+$ for each executed trade, and each retailer incurs a transaction cost $c_r \in \mathbb{R}_+$ for each executed transaction (but they do not incur any cost if an offer is not accepted). 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 price of the good in the consumer market is a function $p : \{0, 1, 2, 3, 4, 5, 6\} \rightarrow \mathbb{R}_+$ of the aggregate supply $Q = q_1 + q_2$. We assume $p(Q) : \{0, 1, 2, 3, 4, 5, 6\} \rightarrow \mathbb{R}_+$ is strictly decreasing in Q .

Profits for the producer are $\sum_{i=1}^2 (p_i q_i - c_0 \mathbf{1}[q_i > 0])$, where $\mathbf{1}[P]$ is an indicator function such that $\mathbf{1}[P] = 1$ if P is true, and $\mathbf{1}[P] = 0$ if P is false. Retailer i profits are $(p(Q) - p_i)q_i - c_r \mathbf{1}[q_i > 0]$.

Because retailer profits depend on the number of units purchased by the other retailer, contracts between the producer and one retailer exert an externality to the other retailer. Because 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

producer: willingness to purchase depends on beliefs about the trades executed between the producer and the other retailer.

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

The first condition is that for the producer, selling one unit at a high price is better 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 .

In this game each equilibrium action profile is played in only one equilibrium strategy profile, and therefore each equilibrium action profile identifies an equilibrium strategy profile. It follows that $\Delta_{\mathcal{O}(a_P)}^s = \Delta_{\mathcal{O}(a_P)}^a$.

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

The first equilibrium action profile is $((p^L, p^L), 2, 2)$, that is, the producer 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. Producer profits are low; the producer would like to deviate offering a high price to both retailers. 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 producer is thus stuck in a low price equilibrium, 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 producer could be tempted to deviate to offer a low price to a retailer, who would then benefit from purchasing three units, which would increase the producer'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 got it as well, and thus that quantities in the market will be high, which limits their willingness to purchase to only two units, thwarting the producer's incentive to deviate in the first place. Off-path beliefs are critical to sustain one or the other equilibria. The most common assumption on off-path beliefs is that they are passive, i.e., retailer i does not update her beliefs on the offer received by retailer $-i$, when i receives a price offer p_i that is off equilibrium. Rather, retailer i believes that $-i$ still receives the equilibrium price offer.

Claim 3 *For any $\theta \in \Theta$, the unique action profile supported by **passive beliefs** in a pure Perfect Bayesian equilibrium of game Γ_θ is $((p^L, p^L), 2, 2)$.*

In the equilibrium action profile supported by passive beliefs, retailers buy two units at a low price in either game. **Wary beliefs** are such that if the producer offers an off equilibrium price \hat{p}_i to retailer i , retailer i believes that the offer to retailer $-i$ is an optimal one for the producer, given \hat{p}_i . Because contracts are unobservable, **wary beliefs** coincide with passive beliefs (McAfee and Schwartz (22)) and action profile L is also the unique pure equilibrium action profile supported by *wary beliefs*. We prove in the Appendix (Claim 11) that there is a unique action profile sustained in a **proper equilibrium** and

it coincides with the action profile supported in equilibrium by wary beliefs or by passive beliefs, and also by In and Wright's RI-equilibrium.

Alternatively, under symmetric beliefs, if retailer i receives an out of equilibrium offer p_i , then she believes that $p_{-i} = p_i$ with probability one (the producer must have deviated to make the same offer to all retailers).

Claim 4 *For any $\theta \in \Theta$, the unique action profile supported by **symmetric beliefs** in a pure Perfect Bayesian equilibrium of game Γ_θ is $((p^H, p^H), 1, 1)$.*

In the equilibrium action profile supported by symmetric beliefs, retailers purchase one unit at a high price.

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.

Out of equilibrium beliefs are a pair $(\omega_1(p^H|\tilde{p}_1), \omega_2(p^H|\tilde{p}_2)) \in [0, 1]^2$, where \tilde{p}_1 and \tilde{p}_2 are off-equilibrium offers observed by retailer 1 and retailer 2, and $\omega_i(p^H|\tilde{p}_i)$ is the probability assigned to $p_{-i} = p^H$ by agent i with beliefs ω_i who observes \tilde{p}_i .

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

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

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

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

$$\omega_i(p^H|\tilde{p}_i) \leq \frac{p^L + 3p(5) - 2p(4)}{2p(3) - 5p(4) + 3p(5)}.$$

Let

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

The hyperplane $\frac{p^H + c_r - p(3)}{p(1) - p(3)} = \frac{p^L + 3p(5) - 2p(4)}{2p(3) - 5p(4) + 3p(5)}$ divides the set of parameter vectors Θ into a subset Θ_L in which a larger set of beliefs sustains the equilibrium action profile L and a subset Θ_H in which a larger set of beliefs sustains the equilibrium action profile H . We obtain the following claim.

Claim 5 *For any $\theta \in \Theta$, the unique action profile supported by the **largest set of beliefs** in a pure Perfect Bayesian equilibrium of game Γ_θ is $((p^L, p^L), 2, 2)$ if $\theta \in \Theta_L$ and $((p^H, p^H), 1, 1)$ if $\theta \in \Theta_H$.*

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.

3.2 Experimental Design and Procedures

We use controlled laboratory experiments to evaluate the predictive power of the different selection criteria. We designed two treatments that only differ in their parameter constellations, one with parameter values $\theta_L \in \Theta_L$ and another with parameter values $\theta_H \in \Theta_H$. We label these Treatments L and H respectively.

The parameters related to producer payoffs are kept fixed across the two games: producer prices are $p^H = 36$, $p^L = 15$ and the producer'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 $k \in \{1, 2, 3, 4, 5, 6\}$.

Treatment L		Quantity bought by the other retailer				Treatment H		Quantity bought by the other retailer			
		<i>0</i>	<i>1</i>	<i>2</i>	<i>3</i>			<i>0</i>	<i>1</i>	<i>2</i>	<i>3</i>
Quantity bought	<i>0</i>	33	33	33	33	Quantity bought	<i>0</i>	28	28	28	28
	<i>1</i>	40	39	3	2		<i>1</i>	79	64	18	4
	<i>2</i>	116	44	42	28		<i>2</i>	175	83	55	19
	<i>3</i>	64	61	40	25		<i>3</i>	127	85	31	4

Table 1: Retailer’s payoffs.

		Quantity bought by the retailer			
		<i>0</i>	<i>1</i>	<i>2</i>	<i>3</i>
Price Charged by the supplier	High	0	28	–	–
	Low	0	–	20	40

Table 2: Supplier’s payoffs for each transaction.

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 producer and retailers’ expected payoffs by multiplying the producer’s payoffs by $4/3$. Tables 1 and 2 respectively summarize retailers’ payoffs and supplier’s payoffs. All these payoffs are expressed in talers, the experimental currency.¹⁰

¹⁰The 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.

All participants were given the role of either a supplier or a retailer, and kept that role throughout the experiment. Participants played the game for 50 rounds, being rematched after every round within matching groups of 12 subjects. After each round, subjects received full feedback about 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.

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. 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 along the same procedure. Subjects received detailed written instructions, which an instructor read aloud (see supplementary appendix). Subjects were asked to answer a questionnaire to check their full understanding of the experimental design before beginning the experiment. At the end of the experiment, subjects completed a short questionnaire which we describe in more detail below.

3.3 Experimental Results

We organize our discussion of the experimental results by focusing, in turn, on prices, quantities and beliefs.

3.3.1 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.

The figure makes two points quite clearly. First, there is stark contrast between the

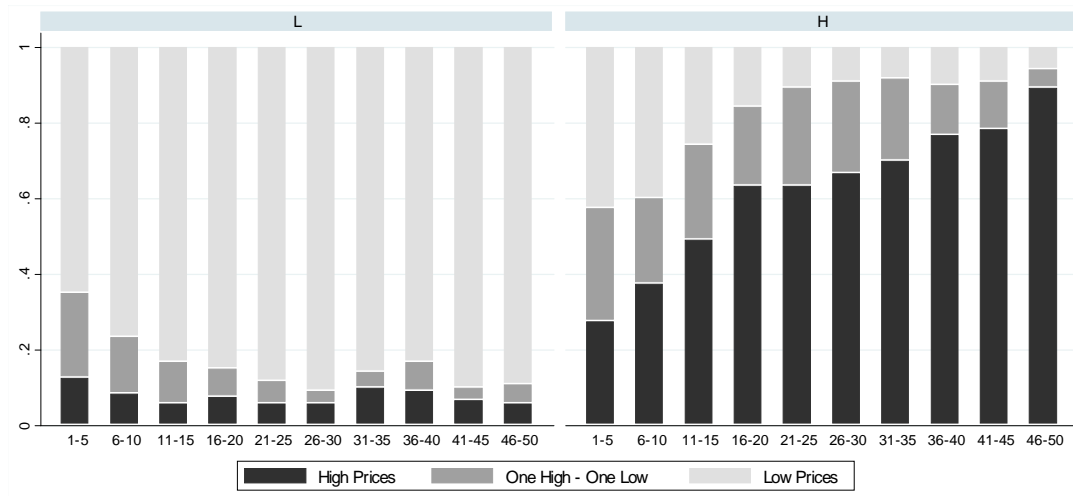


Figure 1: Aggregate prices in both treatments, aggregated by groups of five periods.

	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.

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, $z = 2.882$, $p = 0.0039$).¹¹ This difference across treatments is line with the prediction of the largest set of 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 formalizes these ideas by displaying the results of 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. Table 3 indicates the strong treatment effect found above and the convergence pattern observed in Figure 1.

3.3.2 Quantities

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

¹¹In 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.

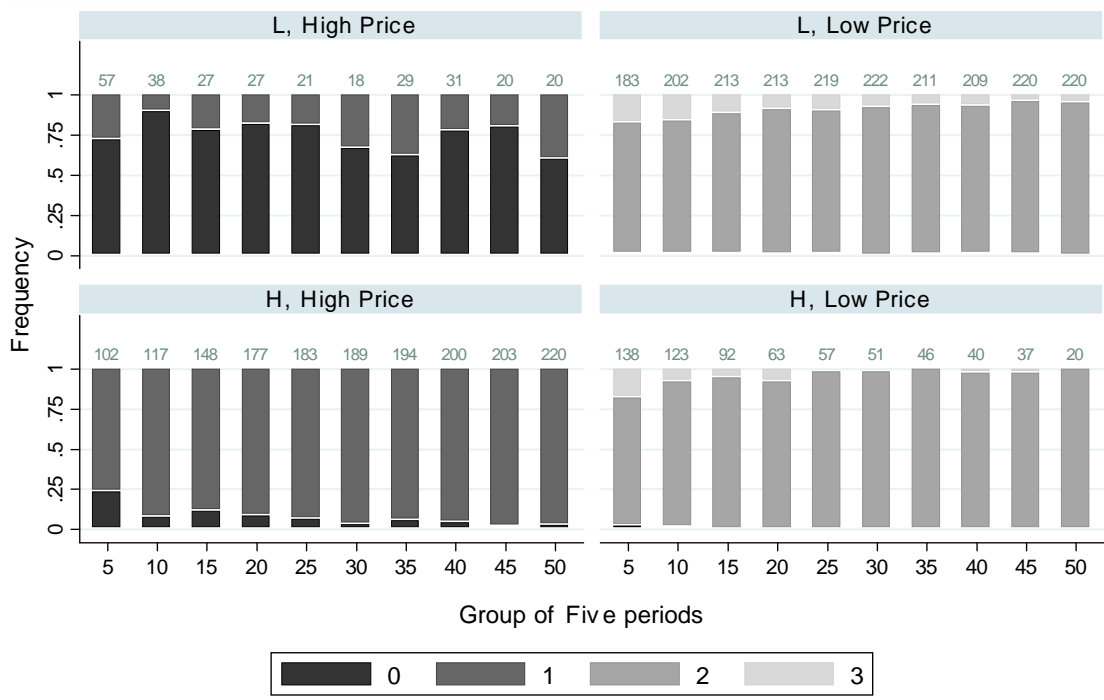


Figure 2: Quantities bought conditional on the price in each treatment.

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.

3.3.3 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 percentage of groups that played an

	All Data	L	H
Passive beliefs	54.10	99.57	10.43
Symmetric beliefs	45.90	00.43	89.56
Wary beliefs / Proper Eq.	54.10	99.57	10.43
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.

equilibrium action profile.¹²

Overall, we find that our selection criterion outperforms both the common prediction of passive/wary beliefs and proper or RI 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/proper equilibrium and symmetric beliefs make the right prediction in 54.10% and 45.90% of cases, respectively. 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.

3.3.4 Beliefs

Beliefs are a crucial element of equilibrium selection in the games presented. In order to assess whether game play was somehow related to participants' beliefs, we elicited beliefs of voters in a non-incentivized manner at the end of the experiment. Immediately after

¹²We restrict the measure to the second half of the experiment due to the convergence process observed in Section 3.3.1. Qualitatively similar results are obtained when considering the whole sample.

		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 ^{S}	30 ^{P}	20 ^{S}	7 ^{P}

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.

finishing the main part of the experiment, subjects completed a questionnaire. The two first questions of the questionnaire 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 substantial difference across 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 would 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.

4 Application: Electoral Competition

4.1 Theory and Predictions

The second application to test our selection criterion is a theory of electoral competition (adapted from Eguia and Nicolò (8)). We present a game of local public good provision. Two candidates $\{A, B\}$ play the role of upstream agents making offers, which in this case take the form of binding campaign promises. Three representative voters $\{1, 2, 3\}$, one from each of three districts, play the role of downstream agents with imperfect information who simultaneously vote for one of the two candidates, or abstain.

Each candidate $J \in \{A, B\}$ chooses whether or not to provide a local public good in each district. Let $S^J = \{0, 1\}^3$ be the strategy set of candidate J and let $s^J = (s_1^J, s_2^J, s_3^J)$ be a pure strategy, where $s_i^J = 0$ indicates that J does not offer the local public good to voter i , and $s_i^J = 1$ indicates that J does offer the public good to voter i . If candidate J uses a mixed strategy, let $p^J = (p_1^J, p_2^J, p_3^J) \in \{0, 1\}^3$ denote the action taken (the policy proposal chosen) by the candidate. If J uses a pure strategy s^J , then $p^J = s^J$.

Each voter i observes p_i^A and p_i^B , that is, whether or not each candidate proposes to provide the local public good in the voter's own district. With probability $1 - \pi$, this information is all that voters observe. With probability $\pi \in (0, \frac{1}{2})$, voters become fully informed and observe p^A and p^B , what candidates propose to all districts. Voters choose whether to vote for candidate A , B , or abstain.

Candidates obtain utility equal to one if they win, zero if they lose. The candidate who receives the most votes wins and implements her proposed policy. Ties are decided by a coin toss.

Voters obtain a benefit $\beta < 1$ if the local public good is provided in their own districts, and pay a cost $\frac{k}{3}$ if the public good is provided to k districts. Therefore, the utility for voter i if candidate J wins is $\beta p_i^J - \frac{p_1^J + p_2^J + p_3^J}{3}$. Notice that local public goods are inefficient: their total cost, normalized to one, surpasses their total benefit β . The socially efficient outcome that maximizes the sum of utilities is not to provide these inefficient local public goods to any district.

The timing is as follows. First, candidates simultaneously choose their policy proposals. Second, Nature determines whether voters become fully informed or not: with probability $\pi \in (0, \frac{1}{2})$ both candidates' proposals become public information, with probability $1 - \pi$, each voter $i \in \{1, 2, 3\}$ only observes the offers each candidate makes to his or her own district and remains unaware of what is offered to other districts. Third, the election is held and voters simultaneously choose one of the candidates or abstain. The candidate who receives the most votes, or a randomly selected candidate in case of a tie, wins the election and implements her policy proposal.

Candidates and voters are strategic, rational agents who evaluate lotteries according to standard expected utilities. We assume that voters choose the candidate whose proposal maximizes their expected payoffs (that is, voters are sequentially rational; see Cho and Kreps (7)), do not use weakly dominated strategies, and, if indifferent between the two candidates, abstain. The solution concept we use is a pure Perfect Bayesian Equilibrium in which no agent uses weakly dominated strategies.

If voters do not fully observe (p^A, p^B) , each voter i is unaware of the proposals made by candidates for districts other than i : voter i only observes either $p_i^J = 0$ or $p_i^J = 1$ for each candidate J ; that is, voter i observes $(p_i^A, p_i^B) \in \{0, 1\} \times \{0, 1\}$, so that i has four information sets in which to make a decision. Voters have imperfect information, and must assign beliefs in each of these information sets. Let $\Delta(\{0, 1\}^3 \times \{0, 1\}^3)$ be a probability distribution over the set of policy proposal pairs offered by the two candidates. Let $\omega_i : \{0, 1\} \times \{0, 1\} \rightarrow \Delta(\{0, 1\}^3 \times \{0, 1\}^3)$ be the set of beliefs of agent i as a function of what she observes. Equilibria depend on the beliefs that voters assign out of equilibrium path. Let $\omega = (\omega_1, \omega_2, \omega_3)$ be the belief profile of the three voters.

Multiple equilibria exist.

Claim 6 (Multiple Equilibria) *For any $\beta \in (\frac{2}{3}, 1)$ and any $k \in \{0, 2, 3\}$, there exists a pure Perfect Bayesian equilibrium (PBE) in which both candidates propose to provide the local public good to k districts. For any $\beta \in (\frac{1}{3}, \frac{2}{3})$ and any $k \in \{0, 1, 2, 3\}$, there exists a PBE in which both candidates propose to provide the local public good to k districts.*

In particular, both the efficient outcome with no local public good provision and the least

efficient outcome in which the public good is provided to every district can be supported in equilibrium by off-path beliefs such that a voter who observes a deviation believes that the deviating candidate offers the local public good to the other two districts.

Selection criteria based on alternative specific off-path beliefs yield sharper predictions. Let p_{-i}^J denote the offers by J to districts other than i .

Given a pure strategy equilibrium in which candidates use strategy pair (s^A, s^B) , a voter i has passive beliefs if she believes that the deviating candidate has only deviated to make a different proposal in i 's district, while all other districts receive their equilibrium offers. For each candidate J , if i observes a deviation to $\hat{p}_i^J = 1 - s_i^J$, i believes that candidate J has played \hat{p}^J such that $\hat{p}_{-i}^J = s_{-i}^J$.

Claim 7 *The unique action profile supported by **passive beliefs** in a pure Perfect Bayesian equilibrium is such that both candidates offer the local public good to all three districts and each voter abstains.*

Passive beliefs select the equilibria with maximum inefficiency.

Given a pure strategy equilibrium in which candidates use strategy pair (s^A, s^B) , a voter i has symmetric beliefs if she believes that the deviating candidate has deviated to make the same proposal to all districts. For each candidate J , if i observes a deviation to $\hat{p}_i^J = 1 - s_i^J$, i believes that candidate J has played \hat{p}^J such that $\hat{p}_{-i}^J = (\hat{p}_i^J, \hat{p}_i^J)$. Selection of equilibria supported by symmetric beliefs does not fully resolve the multiplicity problem in this application.

Claim 8 *If $\beta > \frac{2}{3}$, the unique action profile supported by **symmetric beliefs** in a pure Perfect Bayesian equilibrium is such that candidates do not offer the local public good to any district and each voter abstains. If $\beta \in (\frac{1}{3}, \frac{2}{3})$, there are two classes of action profiles supported by symmetric beliefs in a pure PBE: one in which candidates do not offer the local public good to any district, and another in which both candidates offer the local public good to only one (the same) district.*

Selection by symmetric beliefs is particularly problematic in the case of asymmetric

equilibria: it is difficult to justify why agents who receive different offers in equilibrium must believe that they all receive the same offers off equilibrium.

A voter i has **wary beliefs** if she believes that the deviating candidate has deviated to make a proposal that is optimal for the candidate given what voter i observes. If $\beta > \frac{2}{3}$, selection by passive beliefs, selection by wary beliefs, and the **proper** equilibrium refinement yield the same result: the selected equilibrium action profile is unique and is such that candidates propose to provide the local public good to all districts. If $\beta < \frac{2}{3}$, selection by wary beliefs and the **proper** equilibrium refinement coincide once again, but they do not offer a sharp prediction. They both select the union of all the equilibrium action profiles that can be supported in equilibrium either by passive beliefs or by symmetric beliefs: an action profile in which candidates do not provide the local public good to any district, another in which candidates provide the public good to all districts, and a class of action profiles in which both candidates offer the local public good to only one (the same) district.

We next identify the action profile with the largest set of supporting beliefs.

Comparing the size of the set of beliefs that support each equilibrium action profile is made simpler by noting that a voter with imperfect information who observes a deviation only cares about the expected number of projects in other districts. This expectation, without details of the distribution of these projects across other districts, is a sufficient statistic to form a best response given beliefs. For any observed deviation by a candidate, voter i votes for the deviating candidate if she expects (given her beliefs and the observed deviation) that the number of other projects is sufficiently low. Therefore, any equilibrium strategy profile holds if beliefs are such that following any deviation, the expected number of projects in other districts calculated by any voter observing the deviation is sufficiently high. For example, assume $\beta = 5/6$. Let $E[k|p_i^J] \equiv E[\sum_{j=1}^3 p_j^J - p_i^J | p_i^J, \omega_i]$ denote the expected number of projects offered by J in districts other than i calculated by voter i with beliefs ω_i who observes an offer $p_i^J \in \{0, 1\}$ by candidate J to district i . Consider the action profile in which candidates propose $((0, 0, 0), (0, 0, 0))$ and each voter abstains. In order for this action profile to be sustained in equilibrium, a voter who observes a deviation to $p_i^J = 1$ must not vote for J . The expected utility for voter i who observes a deviation to p_i^J is then

$\beta - \frac{1+E[k|p_i^J]}{3}$ and the expected utility of voting for $-J$ is zero. Voter i votes for $-J$ or abstains (so the deviation fails and the equilibrium holds) if and only if $\beta - \frac{1+E[k|p_i^J]}{3} \leq 0$ or equivalently $E[k|p_i^J] \geq 3\beta - 1 = \frac{3}{2}$. Therefore, the action profile is supported in equilibrium if $E[k|p_i^J] \in [\frac{3}{2}, 2]$ for each voter i and candidate J . Consider now an action profile in which candidates propose $((1, 1, 1), (1, 1, 1))$ and each voter abstains. In order for this action profile to be sustained in equilibrium, a voter who observes a deviation to $p_i^J = 0$, must vote for $-J$ or abstain, which requires $-\frac{E[k|p_i^J]}{3} \leq \beta - 1 = \frac{-1}{6}$ or equivalently, $E[k|p_i^J] \in [\frac{1}{2}, 2]$. Thus, the size of the set of beliefs supporting $((1, 1, 1), (1, 1, 1))$ in equilibrium is larger than the size of the set of beliefs supporting $((0, 0, 0), (0, 0, 0))$.¹³ Therefore, the selection criterion based on the size of beliefs selects $((1, 1, 1), (1, 1, 1))$ ahead of $((0, 0, 0), (0, 0, 0))$ if $\beta = \frac{5}{6}$. In fact, this criterion selects $((1, 1, 1), (1, 1, 1))$ among all equilibria for any $\beta > 2/3$.

Claim 9 *If $\beta \in (\frac{2}{3}, 1)$, the unique action profile supported by the **largest set of beliefs** in a Perfect Bayesian equilibrium (PBE) is such that candidates offer the local public good to all three districts (and each voter abstains). If $\beta \in (\frac{1}{3}, \frac{2}{3})$, the unique action profile supported by the largest set of beliefs in a pure PBE is such that candidates do not offer the local public good to any district (and each voter abstains).*

Selection by passive beliefs fails to recognize that the two cases with β larger or smaller than $\frac{2}{3}$ are fundamentally different: if β is close to one, there is much to gain ($\beta - \frac{1}{3}$) from supporting a deviation away from the efficient policy (no provision) to one that provides the local public good to one district: few beliefs sustain the action profile that leads to an efficient outcome. In contrast, if β is low, there is little to gain ($\beta - \frac{1}{3}$) from supporting a deviation away from the efficient policy to one that offers to provide the local public good to one district, and the action profile is sustained for almost any beliefs. Our selection criterion selects the equilibrium action profile in which candidates do not provide the public good to any district when its provision is highly inefficient, but discards this action profile when public good provision is almost efficient. Selection by passive beliefs treats both situations as identical even though the incentives are quite different. The symmetric beliefs criterion,

¹³As shown in the Appendix, the size of beliefs supporting $((0, 0, 0), (0, 0, 0))$ is $\frac{1}{16}$, while the size of beliefs supporting $((1, 1, 1), (1, 1, 1))$ is $\frac{15}{16}$.

Treatment <i>I</i>		Total Investment in other districts			Treatment <i>NI</i>		Total Investment in other districts		
		<i>0</i>	<i>1</i>	<i>2</i>			<i>0</i>	<i>1</i>	<i>2</i>
Investment in Own District	Yes	90	70	50	Investment in Own District	Yes	60	40	20
	No	56	36	16		No	56	36	16

Table 6: Voter’s payoffs.

implausibly, predicts more public good provision when β is low and public good provision is more inefficient. Selection by wary or symmetric beliefs, or properness, fails to yield a unique prediction.

4.2 Experimental Design and Procedures

We designed two treatments to test the predictive power of the different selection principles in the voting game presented in the previous section. The two treatments only differed on the level of inefficiency of the public good. In Treatment *I* the public good was slightly inefficient ($\beta = 0.9$) while in Treatment *NI* the public good was highly inefficient ($\beta = 0.4$). Our selection criterion predicts that candidates propose to Invest in every district in Treatment *I*, and to Not Invest in any district in Treatment *NI*. The remaining parameters were the same across treatments. Candidates earned 10 experimental talers if they lost an election and 94 if they won. Voters were informed about investment promises in other districts with probability $\pi = 0.25$. The initial endowment of voters was 56 and the price of the investment in each district was 60.¹⁴ Table 6 summarizes voters’ payoffs in both treatments.

All participants were given the role of either a voter or a candidate, and kept that role throughout the experiment. Participants played the game for 40 rounds, being rematched after every round within matching groups of 10 subjects. After each round, subjects received full feedback about actions of all subjects in their subgroup and their payoff for that round. To determine payment, the computer randomly selected four periods for the final payment.

The experiment was conducted at the BonnEconLab of the University of Bonn in Oc-

¹⁴The total amount earned was transformed into euros through the conversion rate of 0.05. In total, subjects earned an average of €12.56, including a show-up fee of €3.

tober 2012. We ran a total of 6 sessions with 20 subjects each. No subject participated in more than one session. We therefore have 6 independent matching groups for each 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 along the same procedure as in the previous experiment. Each experimental session lasted approximately one hour.

4.3 Experimental Results

4.3.1 Politicians' Proposals and Aggregate Investment

Figure 3 plots the aggregate proposals made by politicians in both treatments, aggregated by groups of five periods. The proposals are similar in the first periods: average investment proposed by politicians in the first five periods is not significantly different across treatments (Mann-Whitney test, $z = 0.816$, $p = 0.414$). However, as subjects gain experience playing the game, substantial differences arise. While the average investment proposal in Treatment *I* amounts to 2.62, it is 1.80 in Treatment *NI*. This difference is statistically significant (Mann-Whitney test, $z = 2.242$, $p = 0.025$). This strong and significant comparative static result is predicted by our selection criterion, as well as by wary beliefs and proper equilibrium. Instead, passive beliefs predicts no difference between treatments and symmetric beliefs predicts a difference in the opposite direction.

The clear and opposite trend in the different treatments is evident in Table 7. In Table 7 we estimate a similar multilevel mixed effects logistic regression as Table 3 above with the amount of investment of politicians as the dependent variable.

In Treatment *I*, we observe a significant increase in the total amount of proposed investment through time, clearly converging to the full investment equilibrium. Furthermore, behavior across different independent groups is relatively homogeneous in Treatment *I*.¹⁵ Investment proposals are substantially different in Treatment *NI*. First, we observe a significant decrease in the total amount of investment proposals over time. However, we do

¹⁵Figure 5 in Appendix B displays investment proposals made by politicians disaggregated for each independent group in Treatment *I*.

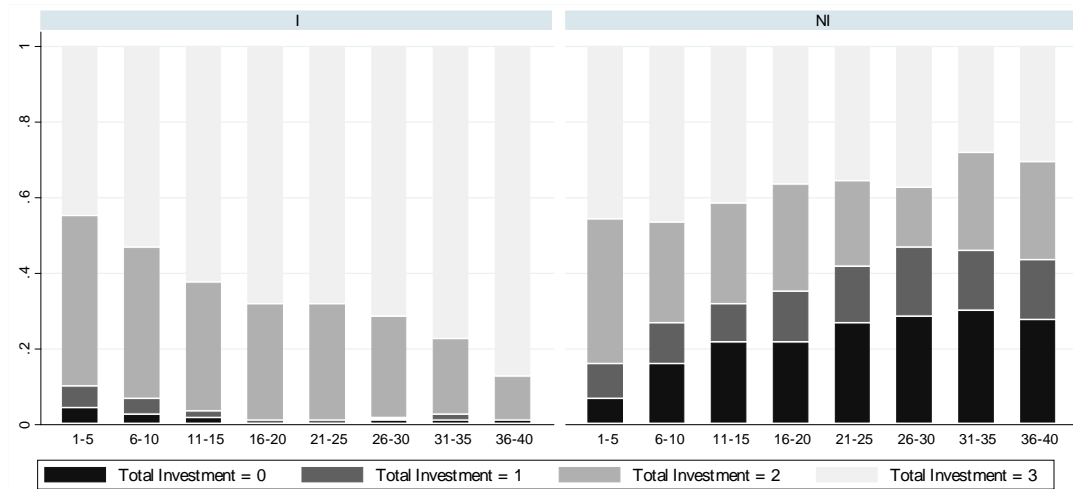


Figure 3: Investment proposals by politicians in each treatment, aggregated in groups of five periods.

	Coef.	Std. Err.	Odds Ratio	z	Pr > $ z $
Treatment I	-0.19	1.07	0.82	-0.18	0.86
Period	0.07	0.01	1.07	9.18	0.05
Treatment I * Period	-0.12	0.01	0.89	-11.97	0.00
Cut 1	-3.90	0.76		-5.14	0.00
Cut 2	-2.67	0.75		-3.54	0.00
Cut 3	-0.04	0.75		-0.06	0.96
Group variance	2.59	1.42			
Subject variance	3.36	0.96			
Number of Observations			1920		
Number of Groups			12		
Number of Subjects			48		
Log Likelihood			-1458.36		
Wald χ^2			147.90		
Hausman test χ^2 *			0.00		

Table 7: Three-level (overall, group, and subject) mixed effects ordered logistic regression of investment offered by politicians as a function of treatment, period, and treatment and period interacted. Treatment I is a dummy variable that takes value one if the treatment is I and zero otherwise. * The Hausman test is done on a linear mirror version of the model.

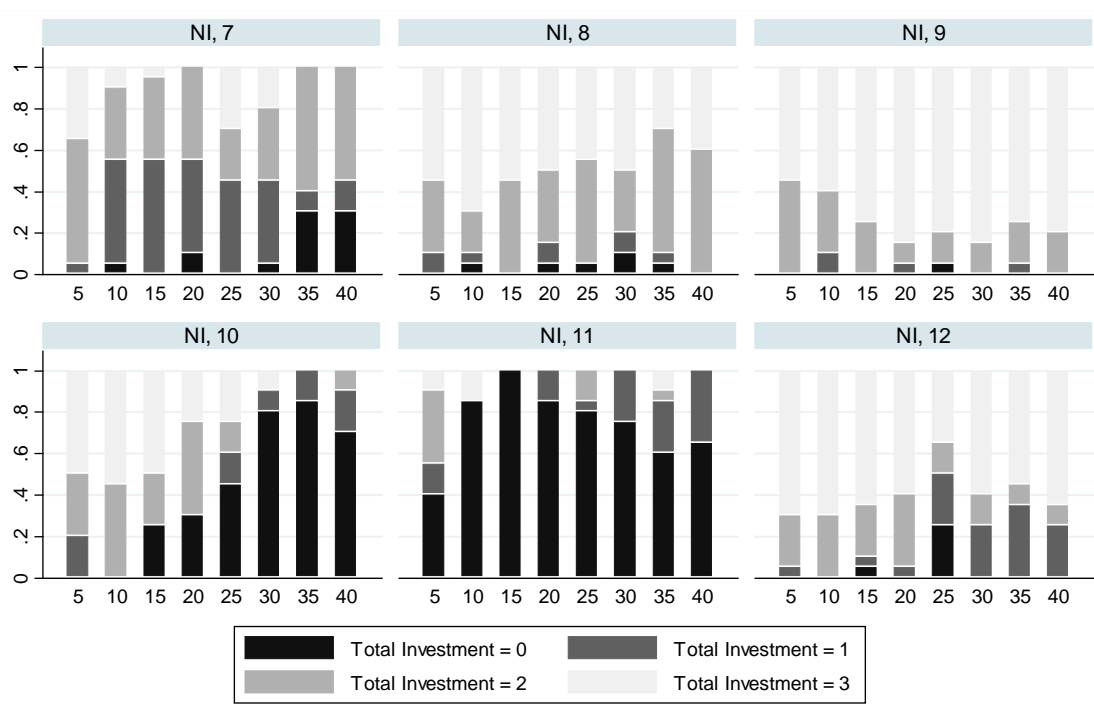


Figure 4: Investment proposals by politicians in each group in Treatment NI, aggregated in groups of five periods.

not observe full convergence to the zero investment equilibrium. Second, as evidenced in Figure 4, we observe heterogeneity across independent groups. Figure 4 displays politician's proposals disaggregated by each independent group.

The analysis above describes proposed policies made by politicians. We now turn to the policies of the winning candidates which were implemented and whether they differ from the proposals. We find the comparison varies by treatment. In Treatment *I*, we find that implemented policies are significantly higher than politicians' proposals (Wilcoxon test, $z = 2.201$, $p = 0.0277$). As we describe in detail in the next subsection, in Treatment *I* voters choose such that they select the politicians who offer a higher total investment. In Treatment *NI* instead, we find no significant difference between the proposed and implemented policies (Wilcoxon test, $z = 0.524$, $p = 0.6002$).

		Conditional Vote						
		Treatment <i>I</i>			Treatment <i>NI</i>			
		a	b	\emptyset	a	b	\emptyset	
FULL	Best	a	91	3	6	91	5	4
INFORMATION	Alternative:	b	3	96	1	4	92	4
		equal	38	16	45	26	26	48
PARTIAL	Highest :	a	93	2	5	60	29	11
INFORMATION	Investment	b	1	96	3	27	65	8
		equal	41	29	30	36	26	38

Table 8: Conditional Vote in the different treatments.

4.3.2 Voters' Behavior

Table 8 plots subjects' voting behavior across different treatments and across different information conditions. In line with the theoretical predictions, subjects mostly vote for the alternative that yields the highest payoff whenever they have full information about the investments in each district (93.60% of the cases in Treatment *I*, and in 91.22% in Treatment *NI*). The difference is not significant (Mann-Whitney test, $z = 0.401$, $p = 0.6884$).

We observe however, a significant difference between treatments in voting behavior whenever subjects are uninformed about investments in other districts. Although in Treatment *I* the percentage of subjects voting for the candidate who offers them a higher investment is 94.63%, the percentage is much lower, 62.94% in Treatment *NI*. This significant difference in behavior (Mann-Whitney test, $z = 2.741$, $p = 0.0061$) stems from the different natures of the equilibri to which groups converge. As seen in the previous section, all groups in Treatment *I* converge to the inefficient equilibrium. In such an equilibrium (supported by passive beliefs, wary beliefs and proper equilibrium), voters should indeed vote for the candidate offering a higher investment in their district. In contrast, in Treatment *NI*, some groups converge to the efficient equilibrium (supported by symmetric beliefs as well as by wary beliefs, proper equilibrium and by largest beliefs). In this alternative equilibrium, voters should optimally vote for the candidate offering no investment. For groups in Treatment *NI*, we observe a significant correlation between voting for the candidate offering higher investment and the implemented level of investment (Spearman's $\rho = 0.8286$, $p = 0.0416$).

	All Data	I	NI
Passive beliefs	77.68	94.00	48.19
Symmetric beliefs	3.97	0	11.14
Wary beliefs / Proper Eq.	67.12	94.00	18.55
Largest set of beliefs	75.97	94.00	43.37
Random	10.14	12.50	5.88
% Equilibrium play	48.54	62.50	34.58

Table 9: Measure of fit of the different selection criteria. Row "Random" displays a measure of fit of a randomly picked equilibrium. The data is restricted to the second half of the experiment.

This difference in behavior explains, in turn, the asymmetric relation between proposed policies and implemented policies across treatments.

4.3.3 Goodness of Fit

In this section we apply the same measure of fit introduced in Section 3.3.3 to the new setting. As a reminder, the measure of fit that we use is the percentage of observations in which an entire group behaved as prescribed by the respective equilibrium. Note, however, that in the games presented in this section, selection by symmetric beliefs, wary beliefs, and the proper equilibrium refinement do not deliver unique predictions in Treatment *NI*. In order to correct for the bigger set of predicted equilibria, we divide the percentage of observations in which a whole group behaved as prescribed by one of the predicted equilibria by the total number of predicted equilibrium profiles. Table 9 displays the results for the second half of the experiment.¹⁶

The results differ considerably from the results observed in the contracting experiment. First, equilibrium predictions have less predictive power in the electoral competition setting than in the contracting one. Overall, the fraction of groups which play according to an equilibrium is 48.54%, compared to 73.92% observed in the contracting experiment. One explanation for the difference in results might be the slower convergence in the voting game, particularly in Treatment *NI*. Additionally, the contracting game was played 10

¹⁶We restrict the measure to the second half of the experiment due to the convergence process observed in Section 4.3.1.

extra periods.¹⁷ The second difference between the voting and contracting games relates to the qualitative ranking of the selection criteria. As in the contracting setting, our selection criterion outperforms symmetric beliefs in terms of predictive power (Wilcoxon test, $z = 2.412$, $p = 0.0159$). But unlike the first setting, the predictive power of our selection criterion (75.97%) does not differ significantly from the predictive power of passive beliefs (77.68%), or wary beliefs or proper equilibrium (67.12%), (Wilcoxon test, $z = 0.042$, $p = 0.9661$ in the first case and $z = 0.127$, $p = 0.8989$ in the second).

The picture does not change much when we disaggregate by treatment. In Treatment *I*, the prediction by the largest set of beliefs coincides with the prediction using passive beliefs, wary beliefs and proper equilibrium. Most groups converge to this prediction of full investment and, therefore, all these selection criteria vastly outperform the prediction of symmetric beliefs. In Treatment *NI*, subjects' behavior is more diverse, equilibrium play is rarer, and accurate predictions more difficult. The predictive power of selection by the largest set of supporting beliefs is not significantly different from that of passive beliefs or wary beliefs, but it remains significantly better than selection by symmetric beliefs. Results in Section 4.3.1 help to explain these findings. On the one hand, most groups in Treatment *NI* start with relatively high levels of investment (which is consistent with the prediction of passive beliefs). On the other hand, all but one group show a significant decrease of total investment over time, but most groups did not fully converge to a stable situation within the 40 periods played. Playing for additional periods might result in convergence to the zero investment equilibrium, which is the equilibrium supported by the largest set of beliefs. The combination of both forces leads to such mixed results.

4.3.4 Voters' Beliefs

As in the first experiment, we elicited beliefs of voters (in a non-incentivized manner). In particular, we asked the following questions: "Suppose that you play an additional period as a voter, and that you only receive information about investment in your district. If a

¹⁷Actually, the voting experiment was conducted first. In light of the slow convergence, we extended the number of periods in the contracting game in order to increase the likelihood of observing, if at all, convergence to equilibrium.

		Total Investment in other districts if there is no Investment in own district					
		Treatment <i>I</i>			Treatment <i>NI</i>		
		0	1	2	0	1	2
Total Investment in	0	0	1	0	6	0	0
other districts if there is	1	1	1	7	5	7	3
Investment in own district	2	4 ^{<i>S</i>}	0	22 ^{<i>P</i>}	1 ^{<i>S</i>}	5	9 ^{<i>P</i>}

Table 10: Joint distribution of voter’s beliefs (in absolute number of observations).

politician offers *investment* / *no investment* in your district, in how many other districts do you think he will offer investment?” To which they could answer “In none of the others”, “In one other district” or “In all the other districts”. Table 10 shows the joint distributions of beliefs.

As in the contracting game, we observe a large difference across treatments. In Treatment *I*, 61.11% of the subjects’ beliefs are in line with the prediction of passive beliefs. In Treatment *NI*, although still the modal response, the percentage of subjects whose beliefs are in line with passive beliefs is only 25%. Surprisingly though, the percentage of subjects whose beliefs are in line with symmetric beliefs is also lower in Treatment *NI*. Still, the significant difference in beliefs across games points in the same direction: off-path beliefs depend on the particular game.

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. To the extent that we cannot predict with confidence which specific 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 these out of equilibrium beliefs, any equilibrium may hold,

but not all equilibrium strategy profiles and action profiles are equally likely to emerge in the equilibrium actually played: those that require very specific beliefs are less likely to be played than action profiles that hold in equilibrium for a large set of beliefs. The probability that a particular strategy profile is played in equilibrium increases in the size of the set of beliefs that support this strategy profile as an equilibrium profile. The equilibrium action profile that is more likely to be played is the one which is supported by the largest set of beliefs. This is the equilibrium action profile we select.

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. For each information set I , we calculate the fraction of the Lebesgue measure of the set of beliefs at information set I that support the strategy profile in equilibrium, over the Lebesgue measure of the set of all possible beliefs at information set I . We then calculate the Cartesian product of these fractions over all information sets that follow from an individual deviation from the equilibrium strategy profile by a first mover in the sequential game under consideration. This Cartesian product defines the size of supporting beliefs of the strategy profile s .

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 this paper we have used the size of the set of beliefs only to identify the action profile(s) with the largest set of supporting beliefs (without seeking to construct an order of action profiles) because we wished to compare our criterion to criteria such as selection by passive,

symmetric or wary beliefs, which offer only a binary partition of the set of equilibrium action profiles (selected / not selected), not a full order. Future experiments can test the order prediction, and measure how the predictive success of selection by the largest set of beliefs depends on the relative size of the largest set of supporting beliefs over the size of the smaller sets of beliefs that support other equilibrium action profiles.

The only previous experiment on equilibrium selection in vertical contracting games is by Martin, Normann and Snyder (21).¹⁸ Their finding that no specific restriction on the set of beliefs can fit well with the data is consistent with our theory. Beyond this negative finding, it is difficult to conduct an equilibrium analysis with their data, because in their experiment, no group of players play according to any equilibrium. Since all their subject groups, without exception, are out of equilibrium, comparisons of equilibrium selection predictions are problematic. Martin, Normann and Snyder (21) suggest that subjects have heterogeneous beliefs and that their behavior can be explained by a (non-equilibrium) model in which a fraction of subjects hold passive beliefs and the remaining subjects hold symmetric beliefs. We agree that agents may have different beliefs (see our Table 5 and Table 10) and this explanation can ex-post rationalize their observed experimental results. However, a theory based on heterogeneous beliefs has poor predictive power as long as the source of this heterogeneity remains unexplained and the fraction of agents who have each class of beliefs is ex-ante unknown. We propose a selection criterion that offers a clear and in most cases unique prediction, determined ex-ante by the characteristics and payoff matrix of each specific game. Our selection criterion not only rationalizes the observed data ex-post, but more importantly, it also predicts ex-ante the results we expect to obtain.

Our theory can be extended to larger classes of finite sequential games, including games with more than two stages or with a more complex information structure. To make it applicable to continuous games, we must overcome two challenges. First, if the set of strategies available to proposers is uncountable, beliefs at out of equilibrium information sets are now a probability measure over an uncountable set. Calculating and comparing sizes of such probability measures is impractical. We resolve this difficulty by considering only

¹⁸To our knowledge, there is no previous experimental work on equilibrium selection in games of electoral competition.

degenerate beliefs that assign all probability mass to a pure strategy. Calculating the size of degenerate beliefs that support an equilibrium at a given information set is sufficiently straightforward. The second challenge is that with an uncountable set of strategies, there are uncountably many out of equilibrium information sets. The Cartesian product over them is not defined, so the Lebesgue measure of beliefs over this Cartesian product is also not defined. Therefore, we must use a different measure of beliefs across information sets. We suggest to define the size of the set of supporting beliefs as the infimum of the size that supports an action profile at a particular information set, where the minimization is over all the information sets in $\mathcal{O}(a_P)$. In this manner, we obtain a size of supporting beliefs for games in which the strategy set of proposers is uncountable.¹⁹

The motivation to use our selection criterion is that we cannot predict players' out of equilibrium beliefs, and it is therefore more appropriate to consider that any beliefs might emerge, than to make ad hoc assumptions about them. However, if previous experience in a particular game, or survey data, provide good reasons to anticipate that players are more likely to hold some beliefs than others, we can consider a more general family of selection criteria based on weighted sizes of beliefs. For each information set I with k information nodes, let $\mu^I = (\mu_1^I, \mu_2^I, \dots, \mu_k^I)$ be a vector of weights ($\sum_{i=1}^k \mu_i^I = 1$ and $\mu_i^I \in [0, 1]$ for each $i \in \{1, 2, \dots, k\}$), which is determined by the prior (held by a theorist modeling the game) over the beliefs that players may hold at that information set. For each strategy profile s , we can then calculate the weighted size of the set of beliefs that support s as an equilibrium strategy profile, where beliefs that are deemed a priori more likely carry a greater weight. This approach generates a family of selection criteria, characterized by the weights μ . Selection by passive, symmetric or wary beliefs are special cases of this family, in which the prior is degenerate and assigns all weight to passive, symmetric or wary beliefs.

We interpret our selection criterion as the benchmark member of this family with a uniform prior, which we see as the default to be used in any game in which we lack solid grounds to form good conjectures on out of equilibrium beliefs. Our experiments show that

¹⁹This modified definition of the size of supporting beliefs can also be adopted as an alternative to simplify the calculations in games where the number of strategies available to the proposers is finite but very large. This alternative leads to the same predictions as the ones presented in both applications.

vertical contracting is one application in which using passive, symmetric or wary beliefs is unwarranted: the predictive power of equilibrium selection by 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. While using a weighted measure could in theory (if we knew the correct prior and weights) lead to better predictions, we have shown that in our vertical contracting games, selecting the equilibrium strategy profile that can be supported by the unweighted largest set of beliefs performs remarkably well.

References

- [1] Eric Avenel. 2012. “Upstream capacity constraint and the preservation of monopoly power in private bilateral contracting.” *Journal of Industrial Economics* 60(4): 578–598.
- [2] Jeffrey Banks and Joel Sobel. 1987. “Equilibrium selection in signalling games.” *Econometrica* 55: 647–662.
- [3] Jakob Bernoulli. 1713. *Ars Conjectandi*. Thurneysen Brothers Press, Basel.
- [4] Heski Bar-Isaac, Guillermo Caruana and Vicente Cuñat, 2012. “Search, Design and Market Structure.” *The American Economic Review*, 102 (2): 1140–1160.
- [5] Benno Buehler and Florian Schuett 2012. “Certification and minimum quality standards when some consumers are uninformed.” TILEC Discussion Paper DP 2012-040.
- [6] Antonio Cabrales, Walter Garcia-Fontes and Massimo Motta. 2000. “Risk dominance selects the leader: an experimental analysis.” *International Journal of Industrial Organization* 18: 137–162.
- [7] In-Koo Cho and David M. Kreps. 1987. “Signaling games and stable equilibria.” *Quarterly Journal of Economics* 102(2): 179–222.
- [8] Jon X. Eguia and Antonio Nicolò. 2012. “On the efficiency of partial information in

- elections.” SSRN Working paper 1931040.
- [9] Urs Fischbacher. 2007. “z-Tree - Zurich toolbox for ready-made economic experiments.” *Experimental Economics*, 10: 171–178.
- [10] Catherine C. de Fontenay and Joshua S. Gans. 2005. “Vertical integration in the presence of upstream competition.” *RAND Journal of Economics* 36(3): 544–572.
- [11] Drew Fudenberg and Jean Tirole. 1991. *Game Theory*. MIT Press, Cambridge.
- [12] Alessandro Gavazza and Alessandro Lizzeri. 2009. “Transparency and economic policy.” *Review of Economic Studies* 76(3): 1023–1048.
- [13] Srihari Govindan and Robert Wilson. 2009. “On forward induction.” *Econometrica* 77(1): 1–28.
- [14] John C. Harsanyi and Reinhard Selten. 1988. *A General Theory of Equilibrium Selection in Games*. MIT Press.
- [15] John C. Harsanyi. 1995. “A new theory of equilibrium selection for games with complete information.” *Games and Economic Behavior* 8: 91–122.
- [16] Oliver Hart and Jean Tirole. 1990. “Vertical integration and market foreclosure.” *Brookings Papers on Economic Activity*, Special Issue, 205–276.
- [17] YOUNGHWAN IN and Julian Wright. 2012. “Signaling private choices.” *KAIST Business School wp 2012-010*.
- [18] Roman Inderst and Marco Ottaviani. 2012. “Competition through commissions and kickbacks.” *American Economic Review* 102(2): 780–809.
- [19] David M. Kreps and Robert Wilson. 1982. “Sequential equilibria.” *Econometrica* 50: 863–894.
- [20] Pierre S. de Laplace. 1820. *Theorie analytique des probabilités*. Courcier, Paris.
- [21] Stephen Martin, Hans-Theo Normann and Christopher M. Snyder. 2001. “Vertical foreclosure in experimental markets.” *RAND Journal of Economics* 32(3): 466–496.

- [22] R. Preston McAfee and Marius Schwartz. 1994. “Opportunism in multilateral vertical contracting: nondiscrimination, exclusivity, and uniformity.” *American Economic Review* 84: 210–230.
- [23] Daniel P. O’Brien and Greg Shaffer. 1992. “Vertical control with bilateral contracts.” *RAND Journal of Economics* 23(3): 299–308.
- [24] Marco Pagnocci and Salvatore Piccolo. 2011. “Vertical separation with private contracts.” *Economic Journal* 122: 173–207.
- [25] Marcin Peski. 2010. “Generalized risk-dominance and asymmetric dynamics.” *Journal of Economic Theory* 145: 216–248.
- [26] Roger B. Myerson. 1978. “Refinements of the Nash equilibrium concept.” *International Journal of Game Theory* 15:133–154.
- [27] Patrick Rey and Jean Tirole. 2007. “A primer on foreclosure” in *Handbook of Industrial Organization* (eds. M. Armstrong and R. Porter) vol. 3: 2145–2220.
- [28] Patrick Rey and Thibaud Vergé. 2004. “Bilateral control with vertical contracts.” *RAND Journal of Economics* 35(4): 728–746.
- [29] David Schmidt, 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.
- [30] Ilya Segal. 1999. “Contracting with externalities.” *Quarterly Journal of Economics* 114: 337–388.
- [31] Ilya Segal and Michael D. Whinston. 2003. “Robust predictions for bilateral contracting with externalities.” *Econometrica* 71(3): 757–791.
- [32] Reinhard Selten. 1995. “An axiomatic theory of a risk dominance measure for bipolar games with linear incentives.” *Games and Economic Behavior* 8: 213–263.

Appendices

Appendix A: Proofs

Proof of Proposition 1

Proof. Given a finite game, there are only finite 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. ■

In Section 2 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 $n = n_P + n_R$ players (n_P proposers and a 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 payoff of each proposer 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 $\varepsilon \geq 0$, $\mathcal{P}_\varepsilon(U) = \{U' \in \mathbb{R}^{|A| \times n} : U_{ak} = U'_{ak} \text{ for any } k \in \{1, 2, \dots, n_P\}, \text{ and } |U_{ak} - U'_{ak}| \leq \varepsilon \text{ for any } k \in \{n_P + 1, n_P + 2, \dots, n_P + n_R\}, \text{ for any } a \in A\}$. 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 10 *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 , $\tilde{U}_{a''k} = \tilde{U}_{a''k}^\lambda$ 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 . ■

Proof of Claim 2

Proof. Consider an arbitrary game Γ_θ ; we first prove that there are two pure strategy profiles that are supported in a Perfect Bayesian Equilibrium (PBE) 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 producer deviates proposing a high price to some retailer, then she decreases her payoff because the retailer buys zero unit and the deviation is not observed by the other retailer. When the price offered is high, to buy zero unit is a best response for a retailer who assigns probability one that the producer is proposing a low price to the other retailer.

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 producer 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 producer is proposing a low price to the other retailer, too.

Consider now any other strategy profile. First, notice that if the producer 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 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 the retailer in equilibrium 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 producer 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 producer to retailer j . Therefore the deviation is profitable for the producer. It follows that there are not asymmetric equilibria such that the producer offers a high price to a retailer and a low price to the other one. If the producer 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 producer offers a low price to both retailers, if retailer i buys three units, retailer j 's best response is to buy zero unit. However, to buy three units for retailer i is not the best response when retailer j buys zero unit. Hence, there are no equilibria such that the producer offers a low price to both retailers and some retailer buys a quantity different than two. ■

Proof of Claim 3 and Claim 4

These claims follow immediately from the proof of Claim 2 and the definitions of passive beliefs, and symmetric beliefs.

Claim 11 *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 Perfect Bayesian Equilibrium (PBE) 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 PBE 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 producer is such that the producer is still playing the strategy $p_{-i} = p^H$ and therefore in any proper equilibrium, the retail 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 are such that when i observes $p_i = p^L$ should buy three units and not two units as prescribed by the equilibrium, making the deviation profitable for the producer. Therefore this equilibrium strategy profile cannot be supported as a proper equilibrium.

Consider the equilibrium $(p^L, p^L), (0, 2), (0, 2)$. A PBE 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 producer is such that the producer is still playing the strategy $p_{-i} = p^L$, then this strategy profile is supported by beliefs that satisfy the ε -proper equilibrium refinement. ■

Claim 5 is proved directly in the text preceding the claim.

Proof of Claim 6

Proof. The result follows as a corollary from Proposition 8 in Eguia and Nicolò (8). ■

Proof of Claim 7

Proof. We first show that the action profile such that candidates promise $s^A = s^B = (1, 1, 1)$, and voters abstain, is supported in equilibrium by passive beliefs, and then we prove that it is unique. Consider the candidate action profile $s^A = s^B = (1, 1, 1)$ and suppose voters have passive beliefs. If candidate J deviates to \hat{s}^J with $\hat{s}_i^J = 0$, i votes for the other candidate. Therefore for any possible deviation, candidate J loses the election if the electorate does not become fully informed. Since $\pi < \frac{1}{2}$, no deviation is profitable. Given Claim 6 we have to show that action profiles such that either $k = 0$, or $k = 2$, or $k = 1$ (an action profile with $k = 1$ is part of an equilibrium only if $\beta < \frac{2}{3}$), cannot be supported by passive beliefs. Consider first any candidate equilibrium strategy profile such that either $k = 0$, or $k = 1$. Consider a voter i with $s_i^J = 0$ for both $J = A, B$. If candidate J deviates to \hat{s}^J with $\hat{s}_i^J = 1$ (and $\hat{s}_l^J = s_l^J$ for both $l \neq i$), voter i votes for candidate J and J wins the election if the electorate does not become fully informed. Since $\pi < \frac{1}{2}$, then the deviation is profitable. Consider now a candidate action profile with $k = 2$ and for all $i = 1, 2, 3$ there exists J such that $s_i^J = 1$ (there is no equilibrium with $k = 2$ and both candidates promising the public goods to the same voters). Consider a voter i such that $s_i^J = 1$ and $s_i^{-J} = 0$. If candidate J deviates to $\hat{s}_i^J = 0$ (and $\hat{s}_l^J = s_l^J$ for both $l \neq i$), voter i with passive beliefs still votes for candidate J and therefore candidate J ties the election if the electorate does not become fully informed and wins the election if the electorate becomes fully informed. Hence the deviation is profitable. ■

Proof of Claim 8

Proof. First, we prove that the action profile such that $s^A = s^B = (0, 0, 0)$ and voters abstain is supported in equilibrium if voters have symmetric beliefs. In equilibrium, candidates tie the election. If candidate J deviates to \hat{s}^J with $\hat{s}_i^J = 1$ voter i with symmetric beliefs votes for the other candidate. For any possible deviation J loses the election if the electorate does not become fully informed, and therefore no deviation is profitable.

Second, we prove that the candidate action profile such that $k = 1$ and both candidates offer the public good to the same voter i is supported in equilibrium if and only if $\beta \in (\frac{1}{3}, \frac{2}{3})$. Suppose that $\beta \in (\frac{1}{3}, \frac{2}{3})$. If candidate J deviates and proposes $\hat{s}_i^J = 0$ voter i votes for the other candidate. If candidate J deviates and proposes $\hat{s}_l^J = 1$ to some voter $l \neq i$, then voter l votes for the other candidate if the electorate does not become fully informed. Since $\pi < \frac{1}{2}$, it follows that no deviation

is profitable. Suppose that $\beta > \frac{2}{3}$. Let l be a voter such that $s_l^J = 0$. If candidate J deviates to $\hat{s}_l^J = 1$, then voter l vote for J , irrespective of being informed or not. Therefore, it is profitable for candidate J to deviate to \hat{s}^J , with $\hat{s}_i^J = 0$ and $\hat{s}_l^J = 1$ for both $l \neq i$.

We prove now that there are not any other action profiles supported in equilibrium by symmetric beliefs. Suppose that $\beta \in (\frac{1}{3}, \frac{2}{3})$, $k = 1$ and each candidate proposes the public good to a different district. Consider voter i such that $s_i^J = 1$. If candidate J deviates to $\hat{s}^J = (0, 0, 0)$, then voter i still votes for J irrespective of being informed or not. Hence J ties the election when the electorate is not fully informed and wins the election when the electorate is fully informed, and the deviation is profitable. Consider a candidate action profile with $k = 3$ for both candidates. If candidate J deviates to $\hat{s}^J = (0, 0, 0)$ all voters vote for candidate J both in case they are fully informed and when they are not fully informed and, therefore, the deviation is profitable. Consider, finally, a candidate action profile such that $k = 2$ for both candidates, and for all $i = 1, 2, 3$, there exists J such that $s_i^J = 1$ (the action profile with $k = 2$ and both candidates offering the public good to the same districts cannot be supported in equilibrium by any beliefs). Consider a voter i such that $s_i^J = 1$ and $s_i^{-J} = 0$. If candidate J deviates to $\hat{s}_i^J = 0$, (and $\hat{s}_l^J = s_l^J$ for both $l \neq i$), then voter i still votes for candidate J irrespective of being informed or not, and therefore candidate J ties the election if the electorate does not become fully informed and wins the election if the electorate becomes fully informed. Hence the deviation is profitable. ■

Proof of Claim 9

Proof. We calculate the size of the set of belief profiles that support each equilibrium action profile.

Part I: $\beta \in (\frac{2}{3}, 1)$. There are three classes of equilibrium action profiles.

- i) Equilibrium action profiles such that both candidates promise $(0, 0, 0)$ and all voters abstain.
- ii) Equilibrium action profiles such that both candidates promise $(1, 1, 1)$ and all voters abstain.
- iii) Equilibrium action profiles such that candidate action profile is an element of the set $\{((0, 1, 1), (1, 0, 1)), ((0, 1, 1), (1, 1, 0)), ((1, 0, 1), (0, 1, 1)), ((1, 0, 1), (1, 1, 0)), ((1, 1, 0), (0, 1, 1)), ((1, 1, 0), (1, 0, 1))\}$ and the voters' action profile is such that each voter i votes for J if $s_i^J > s_i^{-J}$ and abstains if $s_i^J = s_i^{-J}$.

I.i. Equilibria with candidate strategy profile $((0, 0, 0), (0, 0, 0))$.

Suppose voter 1 observes off equilibrium $\hat{s}_1^J = 1$. There are four strategies by candidate J consistent with what is observed: $(1, 0, 0)$, $(1, 0, 1)$, $(1, 1, 0)$, $(1, 1, 1)$. Let beliefs assigned to these strategies be denoted $\omega^1, \omega^2, \omega^3$ and ω^4 . The Lebesgue measure of the set of probability distributions over these four strategies (the 3 dimensional simplex) is $\int_0^1 \frac{x^2}{2} dx = \frac{1}{6}$. Let $\omega_1 = 1 - (\omega^2 + \omega^3 + \omega^4)$. The payoff voting for candidate $-J$ is zero.

The expected payoff voting for candidate J is

$$\beta - \frac{1 + \omega^2 + \omega^3 + 2\omega^4}{3}.$$

An equilibrium with candidate action profile $((0, 0, 0), (0, 0, 0))$ holds if

$$\beta - \frac{1 + \omega^2 + \omega^3 + 2\omega^4}{3} \leq 0, \text{ or}$$

$$\omega^2 + \omega^3 + 2\omega^4 \geq 3\beta - 1.$$

The absolute size of the set of beliefs such that $\omega^2 + \omega^3 + 2\omega^4 \geq 3\beta - 1$, relabeling ω^2 as x , is

$$\int_0^{3(1-\beta)} \frac{\left(1 - \frac{3\beta-1}{2} - \frac{x}{2}\right)(3(1-\beta) - x)}{2} dx = \frac{9}{4} (1-\beta)^3.$$

The relative size with respect to the whole simplex is then $\frac{27}{2} (1-\beta)^3$.

Each voter has only one information set out of equilibrium beliefs reached by single-player deviation. Hence, the absolute size of the set of belief profiles $\omega = (\omega_1, \omega_2, \omega_3)$ supporting the equilibrium is

$$\left[\frac{3^3}{2} (1-\beta)^3\right]^3 = \frac{(3(1-\beta))^9}{8}.$$

I.ii. Equilibria with candidate strategy profile $((1, 1, 1), (1, 1, 1))$.

Suppose voter 1 observes off equilibrium $p_1^J = 0$. There are four strategies consistent with what is observed: $(0, 0, 0), (0, 0, 1), (0, 1, 0), (0, 1, 1)$. Let beliefs assigned to these strategies be denoted $\omega^1, \omega^2, \omega^3$ and ω^4 and let $\omega^1 = 1 - \omega^2 - \omega^3 - \omega^4$.

The payoff if candidate $-J$ wins is $\beta - 1$.

The expected payoff voting if candidate J wins is

$$-\frac{\omega^2 + \omega^3 + 2\omega^4}{3}.$$

The equilibrium holds if

$$\begin{aligned} -\frac{\omega^2 + \omega^3 + 2\omega^4}{3} &\leq \beta - 1 \\ \omega^2 + \omega^3 + 2\omega^4 &\geq 3(1-\beta). \end{aligned}$$

The absolute size of the set of beliefs such that $\omega^2 + \omega^3 + 2\omega^4 \geq 3(1-\beta)$, relabeling ω^2 as x is

$$\frac{1}{6} - \int_0^{3(1-\beta)} \frac{\left(\frac{3(1-\beta)}{2} - \frac{x}{2}\right)(3(1-\beta) - x)}{2} dx = \frac{1}{6} - \frac{9}{4} (1-\beta)^3.$$

The relative size is six times larger, $1 - \frac{27}{2} (1-\beta)^3$.

Since there are three voters, the size of the set of beliefs profiles $\omega = (\omega_1, \omega_2, \omega_3)$ supporting the equilibrium is $\left[1 - \frac{27}{2} (1-\beta)^3\right]^3$. Easy computation shows that for all $\beta \in (\frac{2}{3}, 1)$,

$$\begin{aligned} \left[1 - \frac{27}{2} (1-\beta)^3\right]^3 &> \left[\frac{27}{2} (1-\beta)^3\right]^3 \\ \frac{1}{27} &> (1-\beta)^3 \end{aligned}$$

and therefore the equilibrium $((1, 1, 1), (1, 1, 1))$ has a larger set of beliefs than $((0, 0, 0), (0, 0, 0))$.

I.iii. Equilibria with any permutation of candidate action profile $((1, 1, 0), (0, 1, 1))$.

This equilibrium requires that voter 1 believe with probability 1 that $p^A = (0, 1, 1)$ given an

observed deviation $p_1^A = 0$. Otherwise, voter 1 continues to vote for A after a deviation to $(0, 1, 0)$ and A wins if information is revealed. Thus, the size of beliefs that support this equilibrium is zero.

Therefore, if $\beta \in (\frac{2}{3}, 1)$, the equilibrium action profile supported by the largest set of beliefs is such that both candidates promise $(1, 1, 1)$ and voters abstain.

Part II: $\beta \in (\frac{1}{3}, \frac{2}{3})$. There are five classes of equilibria: i) $((0, 0, 0), (0, 0, 0))$, ii) $((1, 1, 1), (1, 1, 1))$, iii) permutations of $((1, 1, 0), (0, 1, 1))$, iv) permutations of $((1, 0, 0), (1, 0, 0))$ and v) permutations of $((1, 0, 0), (0, 1, 0))$.

II.i. Equilibria with candidate strategy profile $((0, 0, 0), (0, 0, 0))$.

Suppose voter 1 observes off equilibrium $\hat{s}_1^J = 1$. There are four strategies consistent with what is observed: $(1, 0, 0), (1, 0, 1), (1, 1, 0), (1, 1, 1)$. Let beliefs assigned to these strategies be denoted $\omega^1, \omega^2, \omega^3$ and ω^4 .

The payoff voting for candidate $-J$ is zero. The expected payoff voting for candidate J is $\beta - \frac{1 + \omega^2 + \omega^3 + 2\omega^4}{3}$.

The equilibrium holds if

$$\begin{aligned} \beta - \frac{1 + \omega^2 + \omega^3 + 2\omega^4}{3} &\leq 0 \\ \omega^2 + \omega^3 + 2\omega^4 &\geq 3\beta - 1. \end{aligned} \tag{1}$$

The absolute size of the set of beliefs such that $\omega^2 + \omega^3 + 2\omega^4 \geq 3\beta - 1$, relabeling ω^2 as x is

$$\frac{1}{6} - \int_0^{3\beta-1} \frac{\left(\frac{3\beta-1}{2} - \frac{x}{2}\right)(3\beta-1-x)}{2} dx = \frac{1}{6} - \frac{1}{12} (3\beta-1)^3.$$

So the relative size of beliefs is $1 - \frac{1}{2} (3\beta - 1)^3$.

Since there are three voters, the relative size of the set of beliefs profiles $\omega = (\omega_1, \omega_2, \omega_3)$ supporting the equilibrium is $\left[1 - \frac{1}{2} (3\beta - 1)^3\right]^3$.

II.ii. Equilibria with candidate strategy profile $((1, 1, 1), (1, 1, 1))$.

Suppose voter 1 observes off equilibrium $p_1^J = 0$. There are four strategies consistent with what is observed: $(0, 0, 0), (0, 0, 1), (0, 1, 0), (0, 1, 1)$. Let beliefs assigned to these strategies be denoted $\omega^1, \omega^2, \omega^3$ and ω^4 .

The payoff if candidate $-J$ wins is $\beta - 1$. The expected payoff voting if candidate J wins is

$$-\frac{\omega^2 + \omega^3 + 2\omega^4}{3}.$$

The equilibrium holds if

$$\begin{aligned} -\frac{\omega^2 + \omega^3 + 2\omega^4}{3} &\leq \beta - 1 \\ \omega^2 + \omega^3 + 2\omega^4 &\geq 3(1 - \beta). \end{aligned}$$

The absolute size of the set of beliefs such that $\omega^2 + \omega^3 + 2\omega^4 \geq 3(1 - \beta)$, relabeling ω^2 as x is

$$\int_0^{3\beta-1} \frac{\left(1 - \frac{3(1-\beta)}{2} - \frac{x}{2}\right) (3\beta - 1 - x)}{2} dx = \frac{1}{12} (3\beta - 1)^3.$$

So the (relative) size of beliefs is $\frac{1}{2} (3\beta - 1)^3$.

Since there are three agents, the size of the set of beliefs profiles $\omega = (\omega_1, \omega_2, \omega_3)$ supporting the equilibrium is $\frac{1}{8} (3\beta - 1)^9$. For all $\beta \in (\frac{1}{3}, \frac{2}{3})$, we have that $\left[1 - \frac{1}{2} (3\beta - 1)\right]^3 > \frac{1}{8} (3\beta - 1)^9$ and, therefore, the equilibrium action profile such that candidates promise $((0, 0, 0), (0, 0, 0))$ and voters abstain, has a larger set of beliefs than the equilibrium action profile such that candidate promise $((1, 1, 1), (1, 1, 1))$ and voters abstain.

II.iii. Equilibria with any permutation of candidate strategy profile $((1, 1, 0), (0, 1, 1))$.

This equilibrium requires that voter 1 believe with probability 1 that $p^A = (0, 1, 1)$ given an observed deviation $p_1^A = 0$. Otherwise, voter 1 continues to vote for A after a deviation to $(0, 1, 0)$ and A wins if information is revealed. Thus, the size of beliefs that support this equilibrium is zero.

II.iv. Equilibria with any permutation of candidate strategy profile $((1, 0, 0), (1, 0, 0))$.

Deviations in district 1 are never successful, for any beliefs. A deviation to $p_2^A = 1$ is consistent with beliefs that assign probability one to the set of strategies $\{(0, 1, 0), (0, 1, 1), (1, 1, 0), (1, 1, 1)\}$. Let the probabilities for each of these four strategies be respectively labeled $\omega^1, \omega^2, \omega^3, \omega^4$.

The expected payoff for voter 2 voting for B is $-\frac{1}{3}$. The expected payoff voting for A is $\beta - \frac{1 + \omega^1 + \omega^2 + 2\omega^3}{3}$. Hence the equilibrium holds if

$$\begin{aligned} -\frac{1}{3} &\geq \beta - \frac{1 + \omega^1 + \omega^2 + 2\omega^3}{3} \\ \omega^1 + \omega^2 + 2\omega^3 &\geq 3\beta. \end{aligned}$$

The absolute size of the set of beliefs such that $\omega^1 + \omega^2 + 2\omega^3 \geq 3\beta$ is (defining $\omega^1 = x$)

$$\int_0^{2-3\beta} \frac{\left(1 - \frac{3\beta}{2} - \frac{x}{2}\right) (2 - 3\beta - x)}{2} dx = \frac{1}{12} (2 - 3\beta)^3.$$

So the (relative) size of beliefs is $\frac{1}{2} (2 - 3\beta)^3$.

Since there are two agents receiving zero in equilibrium, the size of the set of beliefs profiles $\omega = (\omega_1, \omega_2, \omega_3)$ supporting the equilibrium is $\frac{1}{4} (2 - 3\beta)^6$.

II.v. Equilibrium with any permutation of candidate strategy profile $((1, 0, 0), (0, 1, 0))$.

Without loss of generality, consider a deviation by candidate A . First consider a deviation to $p_1^A = 0$. In order for the equilibrium to hold, it must be that voter 1 does not vote for A . Let $\omega^1, \omega^2, \omega^3$ and ω^4 be the beliefs that voter 1 assigns to, respectively, $(0, 0, 0), (0, 0, 1), (0, 1, 0)$ and $(0, 1, 1)$ subject to observing $p_1^A = 0$. The expected utility for 1 if B wins the election is $-\frac{1}{3}$. The

expected utility if A wins is $\frac{-\omega^2 - \omega^3 - 2\omega^4}{3}$. So in order for the equilibrium to hold, it must be

$$\begin{aligned} -\omega^2 - \omega^3 - 2\omega^4 &\leq -1 \\ \omega^2 + \omega^3 + 2\omega^4 &\geq 1. \end{aligned}$$

The relative size of beliefs that satisfy this inequality is $\frac{1}{2}$.

A deviation to $p_2^A = 1$ is only successful with degenerate beliefs by voter 2 who believes with probability one that the deviation is to $(0, 1, 0)$. So the equilibrium holds for a relative size of beliefs of 1.

A deviation to $p_3^A = 1$ fails and the equilibrium holds if voter 3 does not vote for A after the deviation. Let $\alpha^1, \alpha^2, \alpha^3, \alpha^4$ be the the beliefs that voter 3 assigns to, respectively, $(0, 0, 1), (0, 1, 1), (1, 0, 1)$ and $(1, 1, 1)$ after observing $p_3^A = 1$. The expected payoff for voter 3 is B wins is $-\frac{1}{3}$. The expected payoff if A wins is $\beta - \frac{1 + \alpha^1 + \alpha^2 + 2\alpha^3}{3}$. So in order for the equilibrium to hold, it must be

$$-\frac{1}{3} \geq \beta - \frac{1 + \alpha^1 + \alpha^2 + 2\alpha^3}{3}.$$

The relative size of beliefs that satisfy this inequality, as noted in case II.4., is $\frac{1}{2}(2 - 3\beta)^3$.

So the size of the set of belief profiles that sustain the equilibrium is $\frac{1}{2} \frac{1}{2} (2 - 3\beta)^3 = \frac{1}{4} (2 - 3\beta)^3$

Notice that $\frac{1}{4}(2 - 3\beta)^3 > \frac{1}{4}(2 - 3\beta)^6$ for any $\beta \in (\frac{1}{3}, \frac{2}{3})$. Then compare $\frac{1}{4}(2 - 3\beta)^3$ to $1 - \frac{1}{2}(3\beta - 1)^3$:

$$\begin{aligned} 1 - \frac{1}{2}(3\beta - 1)^3 &> \frac{1}{4}(2 - 3\beta)^3 \\ 4 &> (2 - 3\beta)^3 + 2(3\beta - 1)^3 > 1 + 2. \end{aligned}$$

So the equilibrium supported by the largest set of beliefs is $((0, 0, 0), (0, 0, 0))$ for any $\beta \in (\frac{1}{3}, \frac{2}{3})$. ■

Appendix B

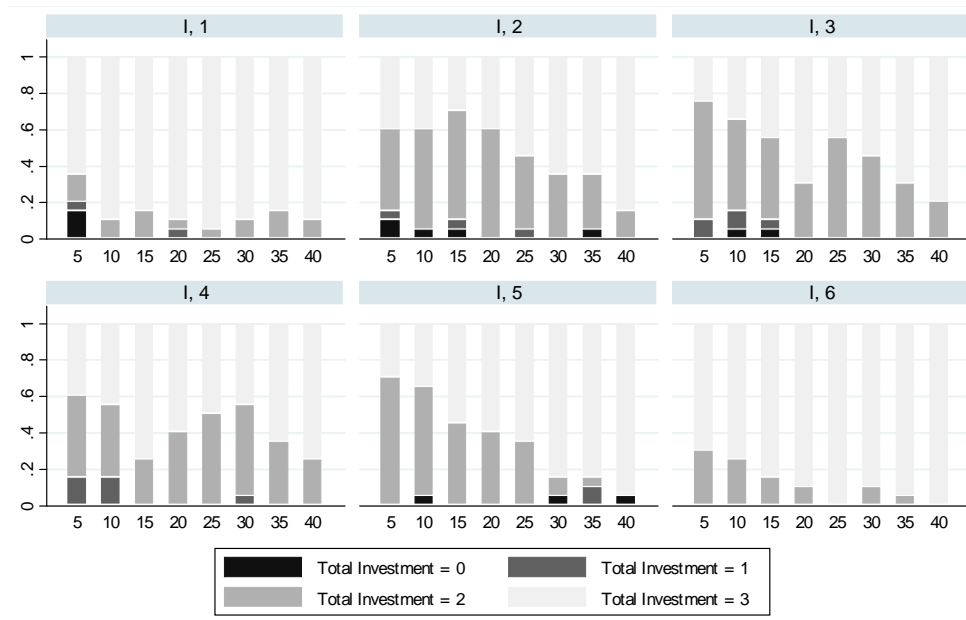


Figure 5: Investment proposals by politicians in each group in Treatment I, aggregated in groups of five periods.

Appendix C. Not For Publication

Equilibria with Wary Beliefs in the Electoral Game

Claim 12 *If $\beta \in (\frac{2}{3}, 1)$, the unique equilibrium that can be supported by wary beliefs is such that candidates promise $((1, 1, 1), (1, 1, 1))$. If $\beta \in (\frac{1}{3}, \frac{2}{3})$ and $\pi \leq \frac{1}{3}$, the set of candidate strategy profiles that can be supported in an equilibrium with wary beliefs is $\{((0, 0, 0), (0, 0, 0)), ((1, 1, 1), (1, 1, 1)), ((1, 0, 0), (1, 0, 0)), ((0, 1, 0), (0, 1, 0)), ((0, 0, 1), (0, 0, 1))\}$; while if $\beta \in (\frac{1}{3}, \frac{2}{3})$ and $\pi \in (\frac{1}{3}, \frac{1}{2})$, the set is $\{((0, 0, 0), (0, 0, 0)), ((1, 1, 1), (1, 1, 1))\}$.*

Proof Part I: Suppose $\beta \in (\frac{2}{3}, 1)$.

There are three classes of equilibria with candidate strategy profiles, respectively: $((0, 0, 0), (0, 0, 0))$, $((1, 1, 1), (1, 1, 1))$, and permutations of $((1, 1, 0), (0, 1, 1))$.

1. Equilibria with candidate strategy profile $((0, 0, 0), (0, 0, 0))$.

Candidate strategy profile $s^A = s^B = (0, 0, 0)$ is supported by beliefs such that every voter i who observes 1 believes that the public good is promised to at least 3β districts. Suppose district i observes candidate J proposing $\hat{s}_i^J = 1$. The optimal strategy for candidate J subject to $\hat{s}_i^J = 1$ is to propose the local public good to two districts because J wins the election if information is revealed. Therefore wary beliefs do not support this strategy profile.

2. Equilibria with candidate strategy profile $((1, 1, 0), (0, 1, 1))$ or permutations of the order of districts.

Consider a voter i such that $s_i^J = 1$ and $s_i^{-J} = 0$. Suppose candidate J deviates and proposes $\hat{s}_1^J = 0$. The optimal strategy for candidate J conditional on \hat{s}_1^J is to propose to provide the public good to a lower number of districts to win the election when information is fully revealed and therefore wary beliefs do not support this strategy profile as an equilibrium.

3. Equilibria with candidate strategy profile $((1, 1, 1), (1, 1, 1))$.

Suppose voter i observes J proposing $\hat{s}_i^J = 0$. For all possible deviation J loses the election when information is not revealed. If candidate J proposes $\hat{s}_l^J = 1$ for both $l \neq i$, then J wins the election when information is revealed. Voter i 's beliefs that assign probability one that the deviating candidate is proposing $k = 2$ (and therefore voter i votes for candidate $-J$) satisfy the notion of wary beliefs and therefore the equilibrium strategy profile $((1, 1, 1), (1, 1, 1))$ is supported by wary beliefs.

Part II: Suppose $\beta \in (\frac{1}{3}, \frac{2}{3})$.

1. Equilibria with candidate strategy profile $((0, 0, 0), (0, 0, 0))$.

candidate strategy profile $s^A = s^B = (0, 0, 0)$ is supported by beliefs such that every voter who observes 1 believes that the public good is promised with some positive probability to another district. Suppose voter i observes $\hat{s}_i^J = 1$. All strategies \hat{s}^J such that $\hat{s}_i^J = 1$ provide the same payoff to candidate J , because with any of them candidate J loses the election whether information is revealed or not. Therefore voter i 's beliefs that assign probability one to the deviation $\hat{s}^J = (1, 1, 1)$ satisfy the notion of wary beliefs, and sustain the equilibrium.

2. Equilibria with candidate strategy profile $((1, 1, 1), (1, 1, 1))$.

Suppose voter i observes $\hat{s}_i^J = 0$. An optimal deviation for candidate J is to propose $\hat{s}_l^J = 1$ for both $l \neq i$; in this case candidate J wins the election when information is revealed. Voter i 's beliefs that assign probability one that the deviating candidate is proposing $k = 2$ satisfy the notion of wary beliefs, and thus sustain this equilibrium.

3. Equilibria with candidate strategy profile $((1, 0, 0), (0, 0, 1))$ or permutations.

Suppose voter i with $s_i^J = 1$ observes a deviation $\hat{s}_i^J = 0$. The optimal strategy for candidate J is to propose $\hat{s}^J = (0, 0, 0)$ because candidate J wins the election when information is fully revealed. Therefore voter i should still vote for candidate J when observing a deviation: wary beliefs do not support this strategy profile as an equilibrium.

4. Equilibria with candidate strategy profile $((1, 0, 0), (1, 0, 0))$ or permutations.

Consider an equilibrium such that both $J = A, B$ promise $(1, 0, 0)$, and given that information is not fully revealed, voter 1 votes for J if $s_1^{-J} = 1$ and $s_i^J = 0$, and abstains otherwise, and voters $l = 2, 3$ always abstain. Notice that for voter $l = 2, 3$ to abstain when $\hat{s}_l^J = 1$ and $s_l^J = 0$ is a best response only if l believes that the number k of districts to which the deviating candidate J promises to provide the local public good is such that $-\frac{1}{3} = \beta - \frac{k}{3}$, or $k = 3\beta + 1$. Suppose candidate A deviates to $\hat{s}_2^A = 1$. The unique deviation $(\hat{s}_1^A, 1, \hat{s}_3^A)$ such that candidate A wins the election if information is fully revealed is $\hat{s}^A = (0, 1, 1)$; candidate A ties the election if $\hat{s}_1^A + \hat{s}_2^A + \hat{s}_3^A = 1$ and loses for any other deviation. If information is not revealed, candidate A loses the election if $\hat{s}_1^A = 0$ and ties the election in all other cases. It follows that if A deviates playing $(0, 1, 1)$, she loses the election if information is not revealed and wins if the information is revealed; if A deviates playing $(1, 1, 1)$, $(1, 1, 0)$, or $(1, 0, 1)$, A ties the election if information is not revealed and loses if information is revealed; the deviation strategies according to which $\hat{s}_1^A = 1$ provide the highest payoff to candidate A if and only if $(1 - \pi)\frac{1}{2} \geq \pi$, that is $\pi \leq \frac{1}{3}$. Therefore if $\pi \leq \frac{1}{3}$ wary beliefs support this strategy profile.

5. Equilibria with candidate strategy profile $((1, 1, 0), (1, 0, 1))$ or permutations.

Consider a voter i such that $s_i^J = 1$ and $s_i^{-J} = 0$. Suppose candidate J deviates and proposes $\hat{s}_1^J = 0$. The optimal strategy for candidate J is to propose the public good to a lower number of districts to win the election if information is fully revealed and therefore wary beliefs do not support this strategy profile as an equilibrium.

Proper Equilibria in the Electoral Game

Claim 13 *Assume $\beta \in (\frac{2}{3}, 1)$. The unique proper equilibrium is such that both candidates propose three projects.*

Proof: For any $\pi \in (0, \frac{1}{2})$, there are three classes of equilibria. We consider each in turn.

1. Equilibria with candidate strategy profile $((0, 0, 0), (0, 0, 0))$.

This strategy profile is supported as an equilibrium only if a voter $i \in N$ who observes $\hat{s}_i^J = 1$ believes that candidate J promises $k \geq 3\beta$. If candidate J deviates to $\hat{s}_i^J = 1$, J wins the election when information is revealed if and only if J promises $k = 2$ public goods. For any other deviations J loses the election when information is not revealed and when information is revealed. In any ε -proper equilibrium, voter i assigns probability which gives a higher weight (of the order of $\frac{1}{\varepsilon}$ for $\pi > 0$ infinitesimally small) to strategies according to which the deviating candidate J promise two public goods than to any other strategy. It follows that the expected number of public goods promised by the deviating candidate J must be lower than 3β for all $\beta > \frac{2}{3}$ and therefore in any ε -proper equilibrium the best response of voter i when observing $\hat{s}_i^J = 1$ is voting for candidate J . Hence the strategy profile $s^A = s^B = (0, 0, 0)$ cannot be part of a proper equilibrium.

2. Equilibria with candidate strategy profile $((1, 0, 1), (1, 1, 0))$ or permutations.

Assume without loss of generality that $s^J = (1, 0, 1)$ and $s^{-J} = (1, 1, 0)$. This strategy profile is part of an equilibrium only if, when observing $s_c^J = 0$, voter c does not vote for candidate J , otherwise candidate J can profitably deviate by promising $s^J = (1, 0, 0)$. Consider a deviation by candidate J such that $\hat{s}_c^J = 0$. If candidate J deviates and promises $(0, 0, 0)$, then J loses the election when information is revealed; if candidate J promises $(1, 1, 0)$, then candidate J ties the election when information is revealed, and, finally, if candidate J promises $(1, 0, 0)$ or $(0, 1, 0)$, J wins the election when information is revealed. Hence voter c should assign a probability that gives a higher weight to strategies according to which the deviating candidate J promises one project than two projects. It follows that in any ε -proper equilibrium the best response of voter i when observing $s_i^J = 1$ is voting for candidate J .

3. Equilibria with candidate strategy profile $((1, 1, 1), (1, 1, 1))$.

Suppose candidate J deviates and promises $\hat{s}_i^J = 0$, candidate J wins the election when information is revealed, either promising zero projects or two projects in the other districts, while J loses the election promising only one project. Beliefs such voter i assigns sufficiently high probability that the deviating candidate is promising $k = 2$, satisfy the proper refinement and therefore this strategy profile is supported by beliefs that satisfy the ε -proper requirements.

Claim 14 *Assume $\beta \in (\frac{1}{3}, \frac{2}{3})$ and $\pi \leq \frac{1}{3}$; then the set of proper equilibria is such that candidate strategy profile are respectively $\{((0, 0, 0), (0, 0, 0)), ((1, 1, 1), (1, 1, 1)), ((1, 0, 0), (1, 0, 0)), ((0, 1, 0), (0, 1, 0)), ((0, 0, 1), (0, 0, 1))\}$. Assume $\beta \in (\frac{1}{3}, \frac{2}{3})$ and $\pi \in (\frac{1}{3}, \frac{1}{2})$; then the set of proper equilibria is such that candidate strategy profile are $\{((0, 0, 0), (0, 0, 0)), ((1, 1, 1), (1, 1, 1))\}$.*

Proof For any $\pi \in (0, \frac{1}{2})$, there are five classes of equilibria. We consider each in turn.

1. Equilibria with candidate strategy profile $((0, 0, 0), (0, 0, 0))$.

This strategy profile is supported as an equilibrium only if no voter $i \in N$ votes for a candidate $j = A, B$ when $\hat{s}_i^J = 1$. Therefore, each voter i has to believe that the expected total number of projects k promised by candidate J is such that $\beta - \frac{k}{3} \leq 0$, that is $k \geq 3\beta$. When information is revealed no deviation is profitable, because the deviating candidate always loses the election, and therefore any belief satisfies the proper refinement.

2. Equilibria with candidate strategy profile $((1, 1, 0), (1, 0, 1))$ or permutations.

Suppose without loss of generality that $s^A = (1, 1, 0)$ and $s^B = (1, 0, 1)$. Suppose candidate A deviates to $\hat{s}_c^A = 1$. In order to be an equilibrium, voter c should still vote for candidate B , and therefore according to his beliefs the number of projects promised by the deviating candidate A must be larger than two $k > 2$. However, when information is revealed, J wins the election if $k = 1$, ties the election if $k = 2$, and loses if $k = 3$. Therefore voter c should assign higher probability to a deviation with $k = 2$ and consequently vote for A . It follows that this strategy profile cannot be supported by beliefs that satisfy the ε -proper requirement.

3. Equilibria with candidate strategy profile $((1, 1, 1), (1, 1, 1))$.

This strategy profile is supported by beliefs such that if voter i observes a deviation $\hat{s}_i^J = 0$, he votes for candidate $-J$. If a deviating candidate deviates to \hat{s}^J such that $\hat{s}_i^J = 0$ and $\hat{s}_l^J = 1$ for both $l \neq i$, J wins the election when information is fully revealed. Therefore to assign a high probability that the candidate is still promising the project to the other two districts satisfies the proper requirement.

4. Equilibria with candidate strategy profile $((1, 0, 0), (1, 0, 0))$ or permutations.

Consider an equilibrium such that both $J = A, B$ promise $(1, 0, 0)$, and when information is not fully revealed, voter 1 votes for J if $s_1^{-J} = 1$ and $s_i^J = 0$, and abstains otherwise, voters $l = 2, 3$ always abstain. Notice that for voter $l = 2, 3$ to abstain when $\hat{s}_l^J = 1$ and $s_l^J = 0$ is a best response only if in case l believes that the number k of public goods promised by the deviating candidate J is such that $-\frac{1}{3} = \beta - \frac{k}{3}$, or $k = 3\beta + 1$. Suppose candidate J deviates to $\hat{s}_2^J = 1$. When information is fully revealed, the unique deviation such that candidate J wins the election is $\hat{s}^J = (0, 1, 1)$. If \hat{s}^J is such that $k = 1$ candidate J ties the election and for all other possible deviation candidate J loses the election. When information is not revealed, candidate J loses the election if $\hat{s}_a^J = 0$ and ties the election in all other cases. It follows that if J deviates by playing $(0, 1, 1)$, he loses the election when information is not revealed and wins when the information is revealed; if J deviates by playing $(1, 1, 1)$, or $(1, 1, 0)$, J ties the election when information is not revealed and loses when the information is revealed; the deviation strategies according to which $\hat{s}_a^J = 1$ provide the highest payoff to candidate J if and only if $(1 - \pi)\frac{1}{2} \geq \pi$ that is $\pi \leq \frac{1}{3}$. Therefore if $\pi \leq \frac{1}{3}$ beliefs of voters b and c that induce these voters to abstain in case of deviation, satisfy the proper refinement.

5. Equilibria with candidate strategy profile $((1, 0, 0), (0, 0, 1))$ or permutations.

Suppose voter i with $s_i^J = 1$ and $s_i^{-J} = 0$ observes a deviation $\hat{s}_i^J = 0$. The optimal strategy for candidate J is to propose $\hat{s}^J = (0, 0, 0)$ because candidate J wins the election when information is fully revealed. Therefore voter i should assign larger probability to this deviation and vote for the deviating candidate J .