Equilibrium Selection in Imperfect Information Games with Externalities

Jon X. Eguia (NYU)* Aniol Llorente-Saguer (Max Planck)†
Rebecca Morton (NYU)‡ Antonio Nicolò (U. Padova)§

April 9, 2013

Abstract

Games with externalities and imperfect information often feature multiple equilibria, which depend on beliefs off the equilibrium path. Standard selection criteria, such as *passive beliefs* equilibrium and *symmetric beliefs* equilibrium rest on ad hoc restrictions on beliefs. We propose a new selection criterion that imposes no restrictions on beliefs: we select the equilibrium that is supported 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: Games with externalities, equilibrium selection, passive beliefs, symmetric beliefs, vertical contracting.

JEL Codes: C72, D86, H41, D72.

–INCOMPLETE: WORK IN PROGRESS. Please Do Not Circulate–

*eguia@nyu.edu (corresponding author) 19 West 4th St. 2nd fl. Dept. of Politics, NYU. New York, NY 10012, US.

†aniol.llorente.saguer@gmail.com, Max Planck Institute, Kurt-Schumacher-Str 10, 53113 Bonn, Germany.

‡rebecca.morton@nyu.edu 19 West 4th St. 2nd fl. Dept. of Politics, NYU. New York, NY 10012, US.

§Department of Economics "M. Fanno", Università degli Studi di Padova, Via del Santo 33, 35123 Padova (Italy).

1 Introduction

In this paper, we propose a solution to the problem of multiplicity of equilibria in imperfect information games with externalities. An important application of these games is multilateral 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 ignores what offers have other firms received (downstream firms operate under imperfect information).

These games typically feature multiple Perfect Bayes and sequential equilibria. Equilibria depend on the beliefs that a downstream firm holds when it receives an unexpected, out of equilibrium offer. Refinements that are useful for signaling games such as the intuitive criterion (Cho and Kreps (2)) have no bite in this context, because the lack of information is about the upstream firms' actions, not about their type (the game is one of imperfect information, not incomplete information; players' types are known).

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 (10) propose three possible beliefs to consider: passive beliefs, symmetric beliefs, and wary beliefs.

Passive beliefs, sometimes called "passive conjectures" (Rey and Tirole (13)), 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 beliefs that all the other unobserved actions remain as in equilibrium. 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 are indeed x, as expected in equilibrium. The selection criterion based on selecting equilibria that can be supported by such passive beliefs is the most frequently used in the literature, not only in games of multilateral vertical contracting (Hart and Tirole (7); O'Brien and Shaffer (11); Segal (15); Fontenay and Gans (5); Rey and Tirole (13)) but also in games of electoral competition in which two candidates make offers to voters in each one of multiple districts (Gavazza and Lizzeri (6)). However, "*in many circumstances the ad hoc restriction to passive beliefs may not be compelling*" (Segal and Whinston (16)). Indeed, while defending the assumption of passive beliefs in the particular game in which they use it, Rey and Tirole (13) 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 beliefs 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 (12). However, most of the literature seems to agree with McAfee and Schwartz's (10) 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 Verge (14), Avenel (1)).

The problem common to all these criteria is the lack of a convincing argument to insist that only one particular set of beliefs is admissible out of the equilibrium path. The very existence of several alternative choices of specific beliefs that have received consideration in the literature as the most natural, or plausible, or appropriate underscores that none of these beliefs are an obvious choice that allows the theorist to assume them without fear and to eliminate all equilibria that rest on different beliefs. We agree that passive beliefs, or wary beliefs, perhaps even symmetric beliefs, may be plausible in a given particular application. The assumption we question, and reject, is that all other beliefs are inappropriate, that all equilibria but those sustained by passive (or symmetric, or wary) beliefs must be discarded.

A more cautious, or modest, approach, is to accept that we cannot pin down the exact beliefs out of 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 (8)). Any assumption of specific beliefs, even a plausible one, is ad hoc and it is difficult to justify it as superior to all others.

We solve the multiplicity problem by suggesting a new selection criterion that does not make ad hoc assumptions on specific beliefs. In fact, the motivation for our solution is that we cannot discern which beliefs are most plausible off equilibrium. Given a solution concept (i.e. Perfect Bayesian or sequential equilibrium) that leads to a multiplicity of equilibria problem, all admissible off-path beliefs that are consistent with the requirements of the given equilibrium concept are equally admissible to us. We make no questionable selection about which beliefs seem more or less natural or intuitive to us.

Rather, we take the conservative view of treating all beliefs as equally plausible, and then we select the equilibrium that can be supported by a largest set of different beliefs. That is, if equilibrium E_1 can only be supported by some very specific beliefs such that after any given deviation, agents must assign probability one or close to one to a particular strategy profile, while on the other hand equilibrium E_2 can be supported regardless of agents' off-path beliefs, or for a wide array of possible beliefs, then we select equilibrium E_2 ahead of equilibrium E_1.

This equilibrium selection criterion has several theoretical features that make it more appealing than selection resorting to passive, symmetric or wary beliefs. First, in finite games, it yields a non-empty prediction, a desirable feature of any solution concept that selection by passive, symmetric or wary beliefs fail to meet. Second, it almost always produces a unique prediction (cases in which it does not are knife-edge, infinitesimal perturbations to payoffs yield a unique prediction). Third, it avoids imposing dubious assumptions on beliefs.

Above and beyond its appealing axiomatic properties, we are interested in its predictive power. We test it against the previous selection criteria in laboratory experiments that match two applications of interest: vertical multilateral contracting, and electoral competition over multiple districts.

Overall we find that our selection criterion has higher predictive power than the other selection criteria. In the first application our selection criterion clearly outperforms the alternative criteria: the largest set of supporting beliefs correctly predicts the behavior of 74% of observations in the experiment; while selection by passive beliefs or wary beliefs correctly predict the behavior of 43% of observations and selection by symmetric beliefs correctly predicts the behavior of 36% of observations in the experiment. In our second application results are somehow a bit weaker given predictive power of the largest set of supporting beliefs still dominates the one by symmetric beliefs, but doesn't differ significantly from passive beliefs. Still, as predicted by the novel equilibrium selection criterion (and not by the other selection criteria) we find strong and significant comparative statics across treatments.

In what follows we first describe our selection criterion in a general class of games, and then we provide the theory, empirical implications, and experimental tests in two applications of interest: vertical contracting, and electoral competition.

2 A New Selection Criterion

Consider a class of games $\mathcal{G}$ with one or more upstream players (proposers), and two or more downstream players (receivers). Proposers make a set of offers to receivers. Each receiver privately receives some information about the set of offers, but may receive only imperfect information about the full set of offers, and all receivers take simultaneous actions in response to the information they individually receive. Each receiver's action generates externalities to other receivers.

For any game $\Gamma \in \mathcal{G}$, let s be a strategy profile. Let S be the set of all strategy profiles. Let ω be a set of beliefs held by receivers (beliefs are probability measures over the set of proposer strategies). Let Ω be the set of all possible beliefs. Let (s, ω) be a Perfect Bayesian equilibrium of game Γ. For any strategy profile s, let W_s be the set of beliefs that supports s in equilibrium, so that $\omega \in W_s$ if and only if (s, ω) constitutes a pure Perfect Bayesian equilibrium of the game.

Let μ be a finite measure over Ω. Then $\mu(W_s)$ denotes the measure or size of the set of beliefs W_s.

Definition 1 *A strategy profile s can be supported by a largest set of beliefs if $\mu(W_s) \geq \mu(W_{s'})$ for any $s' \in S$.*

This selection criterion is in fact a family of selection criteria, as it leaves open the definition of the measure μ. The appropriate measure varies application by application: a benchmark is the standard Lebesgue measure over the simplex on pure strategies.

The motivation for this selection criterion is agnosticism about beliefs off the equilibrium path. Equilibrium 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 will agents hold. If we have a uniform prior over these off equilibrium beliefs, any equilibrium may hold, but not all equilibria are equally likely: equilibria that require very specific beliefs are less likely to be played that equilibria that hold for a vast set of beliefs. Following this logic, the equilibrium most likely to hold is the one supported by the largest set of off-path beliefs.

Consider the following algorithm: a pair (s, ω) of strategies and beliefs is randomly drawn from a uniform distribution over $S \times \Omega$. If this pair constitutes an equilibrium, this is the solution. If not, a new pair is drawn. The equilibrium most likely to be the outcome of this algorithm is the one supported by a largest set of beliefs, measured by the Lebesgue measure.

The purpose of any equilibrium refinement or selection criterion is to solve the problem that multiple equilibria exist 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 (13) for lack of existence of pure equilibria with passive beliefs)[1]. 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 finite game $\Gamma \in \mathcal{G}$ with a non-empty set of pure Perfect Bayesian equilibrium, there exists at least one strategy profile that can be supported by a largest set of beliefs.*

Proof. Given a finite game, there are only finite many strategy profiles, and thus, only finitely many of them are equilibrium strategy profiles. Let k be the number of different strategy profiles that can be sustained in equilibrium. For an arbitrary strategy profile s that can be supported in equilibrium, the value $\mu(W_s)$ is well defined. There are at most k different values, and thus, there exists a maximum among them. ■

A bit more generally, in any game with a finite set of equilibria, selecting the equilibrium supported by a largest set of beliefs guarantees to deliver a well-defined prediction, while the alternative criteria of selection based on passive or symmetric beliefs do not. This prediction avoids making ad hoc assumptions on beliefs off the equilibrium path. In the next two sections, we compare its predictive power in two applications in which alternative selection criteria succeed, at least, in making a prediction.

3 Application: Vertical Contracting

The first application to test our selection criterion is to a theory of vertical contracting with imperfect information and externalities. The game form we present is adapted from Segal

[1] For 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 better than to both accept B, and both to reject their offers other than any other combination. The only pure equilibrium is for both to accept B, and it is not sustained by symmetric beliefs.

(15) and Rey and Tirole (13) (see as well Hart and Tirole (7); McAfree and Schwartz (10); or Rey and Vergé (14) among others).

3.1 Theory and Predictions

A proposer, interpreted as a producer or upstream firm, makes independent offers to each of two receivers, interpreted as retailers or downstream firms. The producer sells a good. An offer takes the form of a price, which is either high (H) or low (L). Let $p_1, p_2 \in \{H, L\}$ be the prices offered by the producer to retailers 1 and 2. The strategy set of the producer is $\{HH, HL, LH, LL\}$, where the first letter 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 seek to describe the simplest multi-lateral contracting game in which the problem of multiplicity of equilibria arises, we constrain the strategy set of retailers, excluding any non-essential strategy. If the price is H, retailers can buy zero or one units; and if the price is L, zero, two or three units. Hence the strategy set for each retailer $i \in \{1, 2\}$ is $S_i = \{0, 1\} \times \{0, 2, 3\}$, where each strategy $(x, y) \in S_i$ corresponds to how much to purchase following an H offer (first coordinate) and following an L offer (second coordinate). Let q_i be the quantity purchased by retailer i.[2]

Retailers sell their units of the good in the consumer market, at a price that is determined by their aggregate supply, which coincides with the total number of units they purchased from the producer. 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.

To test the predictive power of alternative selection criteria, we assume specific parame-

[2] We have checked that the set of equilibria and the equilibria supported by passive beliefs, symmetric beliefs, and by a largest set of beliefs, are invariant if we expand the game to let retailers purchase one unit at a low price. Since adding this superfluous strategy does not change the strategic environment or the results, we use the simpler game in the exposition and in the experiments.

ter values so as to obtain exact predictions that we can test. We consider two variations of the same contracting game form, with different payoffs. We label these contracting games Game $C1$ and Game $C2$. In both games we assume that producer prices are $H = 36$, $L = 15$ and the producer's transaction cost for any executed trade is 15. Payoffs to the producer for a trade with each of the two retailers are then given by the following table, where the row indicates the price, and the column the quantity:

	0	1	2	3
H	0	21	–	–
L	0	–	15	30

Total payoff to the producer are the sum of the payoff for the trade with each of the two retailers.

The two games vary in the retailers' payoffs. In Game $C1$, the retailer's transaction cost for any positive purchase is 29 and the vector of consumer market prices is $(72, 71, 35, 34, 27, 22)$, where coordinate k denotes the price if k units are sold in the market, for any number of units $k \in \{1, 2, 3, 4, 5, 6\}$. In Game $C2$, the retailer's transaction cost for any positive purchase is 33, and the vector of consumer market prices is $(120, 105, 59, 45, 27, 18)$. These parameters generate the following payoff matrix for retailer i for Game C1 (left) and Game C2 (right), where the row indicates (p_i, q_i), the column indicates the number of units purchased by the other retailer, and payoffs are normalized so the payoff for no transaction is zero.

Game C1

	0	1	2	3
$H, 1$	7	6	-30	-31
$L, 2$	83	11	9	-5
$L, 3$	31	28	7	-8

Game C2

	0	1	2	3
$H, 1$	51	36	-10	-24
$L, 2$	147	55	27	-3
$L, 3$	99	57	3	-24

The timing and information is as follows: first, the producer chooses prices $(p_1, p_2) \in$

$\{(H,H),(H,L),(L,H),(L,L)\}$. Retailer $i \in \{1,2\}$ only observes the price p_i offered to her. Second, retailers simultaneously choose a quantity to purchase, and payoffs accrue.

An equilibrium strategy profile is an element of $\{H,L\}^2 \times (\{0,1\} \times \{0,2,3\})^2$ indicating the prices offered by the producer, and the number of units that each retailer buys given a High price and a Low price offer. An equilibrium must specify a strategy profile, and beliefs off the equilibrium path. Let $\omega_i(H|H)$ be the belief (probability) that the producer offered price H to the other retailer, given that $p_i = H$, and similarly let $\omega_i(H|L)$ be the belief (probability) that the producer offered price H to the other retailer, given that $p_i = L$.

We assume that agents maximize expected utility when choosing lotteries, and that they are sequentially rational (Kreps and Wilson (8)), that is, they maximize their expected utility calculated according to their beliefs at any information node, including those out of equilibrium. The solution concept we use is Perfect Bayesian equilibrium in which agents do not use weakly dominated strategies. Multiple equilibria can emerge, depending on beliefs.

The set of equilibrium strategy profiles in the two games coincides.

Claim 2 *The set of equilibrium strategy profiles is the same for Game $C1$ and Game $C2$, and is $\{((L,L),(0,2),(0,2)),((H,H),(1,2),(1,2))\}$.*

That is, in equilibrium, either both retailers buy one unit at a high price, or both retailers buy one unit at a low price. A threat to the equilibrium with high prices is that the producer and one retailer could both become better off by colluding to deviate to a trade of three units at a low price. The equilibrium is sustained by beliefs such that a retailer offered a low price purchases two, not three units, which is optimal under a belief that the producer made the same deviating offer to the other retailer, who then purchased two units as well. The equilibrium with low prices is sustained by beliefs such that if the producer deviates to offer the high price to a retailer, the retailer does not purchase any unit, which is an optimal action given beliefs that the offer to the other retailer continues to be a low price so the other retailer continues to buy two units.

Claim 3 *The unique equilibrium pure strategy profile supported by* ***passive beliefs*** *is $((L,L),(0,2),(0,2))$ in either Game $C1$ or Game $C2$.*

Proof. The proof follows immediately from the proof of claim 2 and the definition of passive beliefs refinement. ■

In the equilibrium with passive beliefs, retailers buy two units at a low price in either game.[3]

Claim 4 *The unique equilibrium pure strategy profile supported by **symmetric beliefs** is* $((H,H),(1,2),(1,2))$ *in either Game* $C1$ *or Game* $C2$*.*

Proof. The proof follows immediately from the proof of claim 2 and the definition of symmetric beliefs refinement. ■

In the equilibrium with symmetric beliefs, retailers purchase one unit at a high price in either game.

Whether an equilibrium holds or not hinges on the belief that the other retailer was offered a high price, following a deviation. If the probability assigned following a deviation to a high price for the other retailer is low enough, the equilibrium holds. Otherwise, it fails.

Given any pure strategy profile s, beliefs along the equilibrium path are correct. Out of equilibrium beliefs are a pair $(\omega_1(H|\tilde{p}_1), \omega_2(H|\tilde{p}_2)) \in [0,1]^2$, where $\tilde{p}_1$ and $\tilde{p}_2$ are off-equilibrium offers observed by retailer 1 and retailer 2, and $W_s = \{\omega \in [0,1]^2 : (s;\omega)$ is an equilibrium$\}$, that is, W_s is the set of beliefs that support s in equilibrium.

Let μ_L be the Lebesgue measure over the unit square $[0,1]^2$. In this application, we let the measure on set of beliefs be this ordinary Lebesgue measure.

Definition 2 *The measure* μ *on the set of beliefs* W_s *that supports an equilibrium with strategy profile* s *is* $\mu(W_s) = \mu_L(W_s)$*.*

Notice that for any s, there exist $w_1^s, w_2^s \in [0,1]$ such that $W_s = [0, w_1^s] \times [0, w_2^s]$. That is, the set of beliefs that support an equilibrium with strategy profile s is the set of beliefs

[3] Because contracts are unobservable, wary beliefs coincide with passive beliefs (McAfee and Schwartz (10)) and Equilibrium $((L,L),(0,2),(0,2))$ is also the unique pure equilibrium supported by *wary beliefs* (wary beliefs are such that if the producer offers an off equilibrium price $\hat{p}_i$ to retailer i, retailer i believers that the offer to the other retailer is an optimal one for the producer, given $\hat{p}_i$).

such as each retailer's belief (belief that the other retailer has received an H offer out of equilibrium) is sufficiently low. So

$$\mu_L(W_s) = w_1^s w_2^s.$$

For instance, in Game $C1$, the strategy profile $((H,H),(1,2),(1,2))$ is supported in equilibrium only by beliefs that make it a best response to purchase two units given $p_i = L$. The expected utility of purchasing two units is

$$11\omega_i(H|L) + 9(1 - \omega_i(H|L)), \tag{1}$$

whereas, the expected utility of purchasing three units is

$$28\omega_i(H|L) + 7(1 - \omega_i(H|L)), \tag{2}$$

hence purchasing two units is a best response if and only if (1) is no less than (2), or, equivalently $\omega_i(H|L) \leq \frac{2}{19}$, so $w_i^s = \frac{2}{19}$ for $i \in \{1,2\}$ and the size of the set of beliefs that supports an equilibrium with strategy profile $((H,H),(1,2),(1,2))$ is $\left(\frac{2}{19}\right)^2$. The strategy profile $((L,L),(0,2),(0,2))$ is supported in equilibrium by beliefs such that purchasing zero units is a best response following $p_i = H$, which requires

$$0 \geq 7\omega_i(H|L) - 30(1 - \omega_i(H|L)),$$

that is, $\omega_i(H|L) \leq \frac{30}{37}$, and the size of of the set of beliefs that supports an equilibrium with strategy profile $((L,L),(0,2),(0,2))$ is $\left(\frac{30}{37}\right)^2$. Since this size is larger, this is the equilibrium supported by a largest set of beliefs in Game $C1$, and is the equilibrium we select.

Selection by passive beliefs (or by symmetric beliefs) yields the same prediction in Game $C1$ and Game $C2$. However, the two games are different. While the set of equilibria are the same in the two games, the cardinality of the payoffs, and thus the quantitative incentive to play or another strategy, is very different in the two games. Our selection criterion recognizes this difference in the payoff structure of the two games, providing a different

Treatment C1

Quantity bought by the other retailer

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

Treatment C2

Quantity bought by the other retailer

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

Table 1: Retailer's payoffs.

prediction in either game.

Claim 5 *The equilibrium strategy profile supported by a* ***largest set of beliefs*** *is* $((L, L), (0, 2), (0, 2))$ *in Game* $C1$ *and* $((H, H), (1, 2), (1, 2))$ *in Game* $C2$*.*

In short, we predict that retailers buy two units at a Low price in Game $C1$, and one unit at a High price in Game $C2$.

3.2 Experimental Design and Procedures

We use controlled laboratory experiments to evaluate the predictive power of the different selection criteria. We ran two treatments ($C1$ and $C2$) corresponding to the two games ($C1$ and $C2$) presented in section 3.1. 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.

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 payoff for that round.

		Quantity bought by the other retailer			
		0	1	2	3
Price Charged to the Retailer	High	0	28	–	–
	Low	0	–	20	40

Table 2: Supplier's payoffs for each of the retailers.

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). Before starting, subjects were asked to answer a questionnaire to check their full understanding of the experimental design. In the end, subjects had to fill a short questionnaire for statistical purposes.

To determine payment, the computer randomly selected five periods for the final payment. The total amount earned in this periods was transformed into euros through the conversion rate of 0.03 in Treatment $C1$ and 0.045 in Treatment $C2$. In total, subjects earned an average of €12.87, including a show-up fee of €4. Each experimental session lasted approximately one hour.

3.3 Experimental Results

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

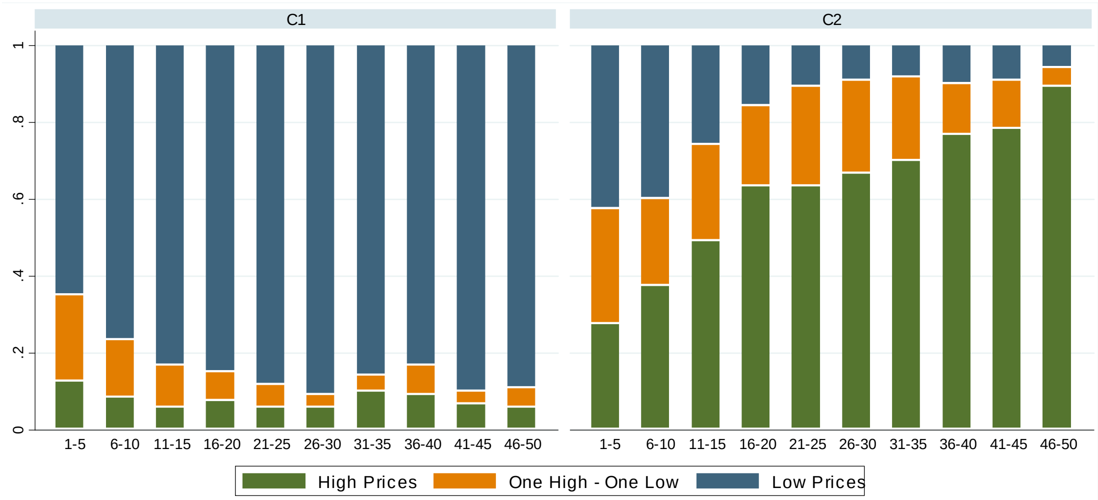

Figure 1: Aggregate prices in both treatments, aggregated by groups of five periods.

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 colors: the green part represents the frequency of suppliers which set both prices high, the yellow part represents the frequency of suppliers which set a high price for a retailer and a low price for the other, and the blue part represents the frequency of suppliers who charge low price to both retailers.

The figure makes two points quite clearly. First, there is stark contrast between the pricing strategies used in both treatments: while in Treatment $C1$ prices are high only in 12% of the cases, in Treatment $C2$ they are high in 72.21% of the cases. This difference is clearly significant (Mann-Whitney test, $z = 2.882$, $p = 0.0039$)[4]. 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 $C1$ is as low as 8.33%, while in Treatment $C2$ is as high as 91.67%.

Table 3 formalizes these ideas by displaying the results of a GLS random effects regression of the amount of high prices set by suppliers as a function of the number of period,

[4] In 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.

		Period		Constant	
Treatment	Group	Coef.	St.Err.	Coef.	St.Err.
C1	1	$-.0015$	.0022	.2339	.1425
	2	$-.0124^{***}$	.0022	$.4812^{***}$	.0817
	3	$-.0077^{***}$	.0019	$.3730^{***}$	.0909
	4	$.0081^{**}$	.0035	.4918	.3319
	5	$-.0062^{***}$	.0017	$.2585^{***}$	.0503
	6	$-.0063^{***}$	.0017	$.2681^{***}$	.0802
C2	7	$.0197^{***}$	.0037	$.6508^{***}$	.1226
	8	$.0332^{***}$	.0029	$.7261^{***}$	.1434
	9	$.0119^{***}$	.0023	1.4857^{***}	.1110
	10	$.0113^{***}$	.0037	$.7857^{***}$	.2203
	11	$.0255^{***}$	.0029	$.8973^{***}$	.1720
	12	$.0188^{***}$	.0029	1.0388^{***}	.1053

Table 3: GLS Random effects regression on the amount of high prices set by suppliers on the period and a constant independently for each group. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$

independently for each matching group. Table 3 indicates that this learning/convergence pattern observed in Figure 1 is observed in almost every independent group: all groups but one display a negative trend in Treatment $C1$ (four of them significantly) while all groups in Treatment $C2$ show a significantly positive one. Besides, the estimated starting levels of prices (i.e. the constant) in C2 are higher than those ones in C1.

3.3.2 Quantities

Figure 2 displays the aggregate quantities bough conditional on the price in each treatment in groups of five periods. The graphs in the top (bottom) belong to Treatment $C1$ ($C2$). The graphs on the left (right) correspond to situations in which retailers were offered a high (low) price. Each bar is divided into different colors which represent the amount bought: yellow if the retailer bought zero units, green for one, blue for two and red for three.

Figure 2 shows two clear patterns. First, retailer's 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 $C1$ and 91.90% in Treatment $C2$. The small difference across treatments is not significant (Mann-Whitney test, $z = 0.961$,

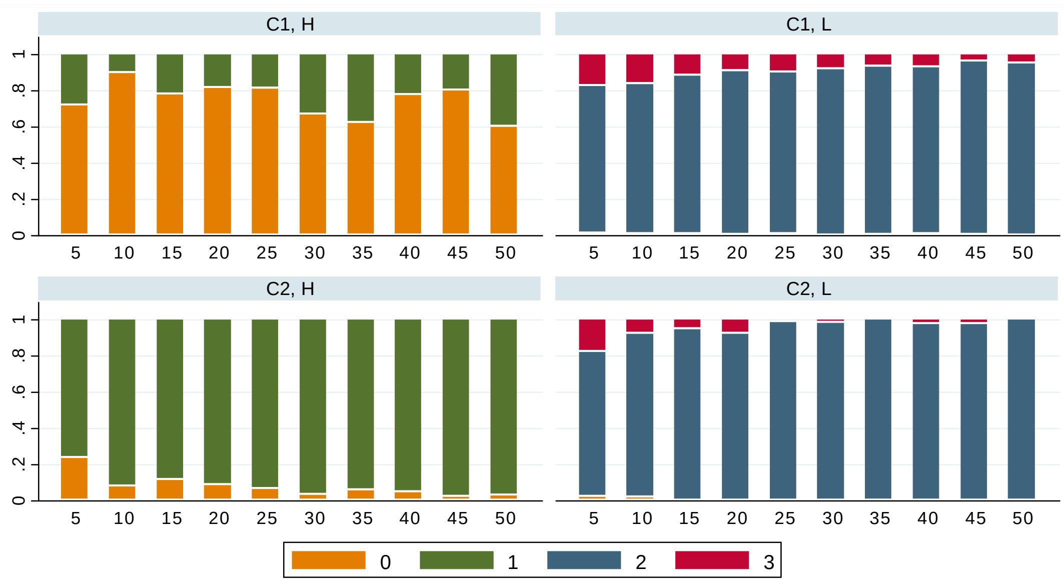

Figure 2: Quantities bought conditional on the price in each tratment.

$p = 0.3367$). Second, there is a substantial difference in retailer's behavior across treatments when receiving a high price. In treatment $C1$, retailers demand one unit 24.65% of the time. This percentage however, goes up to 93.48% in treatment $C2$. This percentage is of course highly significant (Mann-Whitney test, $z = 2.882$, $p = 0.0039$).

3.3.3 Goodness of Fit

The goal of this experimental part is to evaluate the predictive power of the different selection criteria described in the previous section. We try to answer this question in this subsection by comparing a measure of fit. The measure of fit that we use is the percentage of observations which a whole group behaved as preserved by the different equilibria. Because this measure is quite restrictive, we also include the measure separately for only suppliers or only retailers. Note that these additional measures allows us to consider out-of-equilibrium behavior. Table 4 displays the results for the second half of the experiment.[5]

Overall, we find that our selection criterion outperforms both passive beliefs and sym-

[5] We restrict the measure to the second half of the experiment due to the convergence process observed in section 3.3.1. SHOULD WE ADD THE TABLE FOR THE WHOLE SAMPLE IN THE APPENDIX?

		All Data	C1	C2
Whole Group	Passive beliefs	**42.33**	76.33	8.33
	Symmetric beliefs	**35.92**	0.33	71.50
	Largest set of beliefs	**73.92**	76.33	71.50
Only Suppliers	Passive beliefs	48.17	87.83	8.50
	Symmetric beliefs	41.83	7.50	76.17
	Largest set of beliefs	82.00	87.83	76.17
Only Retailers	Passive beliefs	55.00	91.08	18.92
	Symmetric beliefs	92.00	87.25	96.75
	Largest set of beliefs	93.92	91.08	96.75

Table 4: Percentage of events in which the whole group, the suppliers or the retailers behave according to the different selection criteria in the second half of the experiment.

metric beliefs. While our selection criterion makes the right prediction in 73.92% of cases, passive beliefs and symmetric beliefs make the right prediction in 42.33% and 35.92% 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, once can see that the predictive power of our selection criterion equalizes the best of the other criteria. Recall that in $C1$, the prediction of our criterion coincides with the prediction of the equilibrium under passive beliefs, while in $C2$, 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 it matches in each case the best performer of the other two.

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 the beliefs that participants had, we elicited beliefs of voters in a non-incentivized manner at the end of the experiment. Right after finishing the main part of the experiment, subjects had to fill a questionnaire. The two first questions of the questionnaire related to their beliefs. In particular, we asked

		Other Retailer's price if offered a high price			
		Treatment $C1$		Treatment $C2$	
		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.

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, the prediction by symmetric beliefs (in equilibrium) would be to expect the other player to be offered a low price when offered a low price himself and a high price when offered a high price himself. With passive beliefs, retailers should would expect a low price regardless of the price offered himself.

Note that in Treatment $C1$, most retailers beliefs are in line with the beliefs predicted by passive beliefs while in Treatment $C2$, instead, most retailer's beliefs are in line with the prediction of symmetric beliefs. This feature reinforces the results shown in 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 (3)). 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 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 $p^J = (s_1^J, s_2^J, s_3^J) \in \{0,1\}^3$ 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 is all that voters observe. With probability $\pi \in \left(0, \frac{1}{2}\right)$, voters become fully informed and observe P^A and p^B, what candidates propose to all districts. Voters choose which candidate to vote for, or they may abstain.

Candidates obtain utility one if they win, zero if they lose. The candidate who wins 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 district, 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 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 $\pi \in \left(0, \frac{1}{2}\right)$ both candidates' proposals become public information, with probability $1-\pi$, each voter $i \in \{1,2,3\}$ only observes what each candidate offers to the voter's district, and remains unaware of what candidates offer to other districts. Third, the election is held and voters simultaneously vote for one of the candidates or abstain. The candidate who receives 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 vote for the candidate whose proposal maximize their expected payoff (that is, voters are sequentially rational; see Cho and Kreps (2)), 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 remains unaware about what each candidate proposes in 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 $\Delta\left(\{0,1\}^3 \times \{0,1\}^3\right)$ be a probability distribution over the set of policy proposal pairs offered by the two candidates. Let $\omega_i : \{0,1\} \times \{0,1\} \longrightarrow \Delta\left(\{0,1\}^3 \times \{0,1\}^3\right)$ a 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.

There exists multiple equilibria.

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

Proof. The result follows as a corollary from Proposition 8 in Eguia and Nicolò (3). ■

In particular, both the efficient outcome with no local public good provision and the least efficient outcome in which all districts 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 a sharper prediction. 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 pure equilibrium supported by* ***passive beliefs*** *is such that both candidates offer the local public good to all three districts.*

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 pure equilibrium supported by* ***symmetric beliefs*** *is such that candidates do not to offer the local public good to any district. If $\beta \in \left(\frac{1}{3}, \frac{2}{3}\right)$, there are two pure equilibria supported by symmetric beliefs: 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.[6]

We propose as a selection criterion, without making ad hoc assumptions on off equilibrium path beliefs, the equilibrium strategy profile that can be supported by a largest set

[6] As a third alternative, 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. An equilibrium in which candidates propose to provide the local public good to all districts can be supported by wary beliefs; if $\beta < \frac{2}{3}$, an equilibrium in which candidates do not provide the local public good to any district can also be sustained.

of beliefs. In this application, the ordinary Lebesgue measure over beliefs is inappropriate. A belief ω_i is a probability distribution over four events: neither of the other two districts receives an offer for a public good project; only one of them receives it; only the other one receives it; or they both receive it. Any set of beliefs that assigns probability zero to any of these four events has Lebesgue measure zero.

The voter is not concerned about which other districts receive a public good offer, only about *how many* other districts receive a public offer. The expectation over this quantity given by beliefs is all that matters about these beliefs. For each voter i and candidate J, let $\theta_i^J : W_s \longrightarrow [0, 2]$ be a function such that for any $\omega \in W_s$, $\theta_i^J(\omega)$ is the expected number of offers by candidate J of local public good projects to districts other than i by candidate J, calculated by voter i who observes an off-equilibrium offer p_i^J, does not observe p^J fully, and holds beliefs ω_i. Then $\theta_i^J(W_s) \subseteq [0, 2]$ is the set of out of equilibrium path expectations that i may form give a deviation by J, and given beliefs ω_i that sustain s in equilibrium. To this set we can apply the ordinary Lebesgue measure μ_L, and obtain a size of sets of expectations such that i does not support a deviation by J. Since it takes only one (candidate, voter) pair to deviate to break an equilibrium, we find that the appropriate measure of beliefs that sustain an equilibrium is the minimum over all voters and candidates of $\mu_L\ (\theta_i^J(W_s))$.

Definition 3 *The measure μ on the set of beliefs W_s that can support strategy profile s in equilibrium is* $\mu(W_s) = \min_{i\in\{1,2,3\}} \left\{ \min_{J\in\{A,B\}} \left\{\mu_L(\theta_i^J(W_s))\right\} \right\}$.[7]

It can be checked that $\theta_i^J(W_s)$ is a closed interval with upper boundary at 2, that is, if an equilibrium holds for any given beliefs ω, it continues to hold if a voter develops more pessimistic beliefs (higher expected number of unwanted projects in other districts θ_i^J) following a deviation. Let $w_s^{J,i}$ denote the lower boundary of this interval, so that

[7]Nevertheless, we also calculate the Lebesgue measure over the Cartesian product of $\theta_i^J(W_s)$ over all voters and both candidates. Let $\mu_2(W_s) = \prod_{i=1}^{3} \prod_{J\in\{A,B\}} \mu_L(\theta_i^J(W_s))$.

We find that for the parameter values we test in the experimental section ($\beta = 0.4$ and $\beta = 0.9$), the equilibrium that is supported by a largest set of beliefs using measure μ_2 coincides with the equilibrium that is supported by a largest set of beliefs using measure μ.

$\theta_i^J(W_s) = [w_s^{i,J}, 2]$. Then

$$\mu(W_s) = \min_{i \in \{1,2,3\}} \left\{ \min_{J \in \{A,B\}} \left\{ 2 - w_s^{i,J} \right\} \right\}.$$

The intuition is that in order for equilibrium s to hold, following any deviation, each voter must have a sufficiently pessimistic expectation about the deviation, where "pessimistic" in this context means that a high quantity of unwanted projects in other districts are offered. The equilibrium is most fragile to the one deviation by candidate J to voter i that requires voter i to have the highest expectation about this quantity of unwanted projects in order for the equilibrium to hold. The set of expectations for which the equilibrium holds against this most threatening deviation is our measure of size of beliefs that support the equilibrium.

Calculating this size is simple. For example, assume $\beta = 5/6$. Let s be an equilibrium strategy profile in which candidates propose $((0,0,0),(0,0,0))$. Given s, a voter who observes a deviation to $p_i^J = 1$ does not vote for J. The expected utility for voter i who observes a deviation $p_i^J = 1$ from voting for J given beliefs ω_i is then $\beta - \frac{1+\theta_i^J(\omega)}{3}$ 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 $\beta - \frac{1+\theta_i^J(\omega)}{3} \leq 0$ or equivalently $\theta_i^J(\omega) \geq 3\beta - 1 = \frac{3}{2}$. Therefore, $w_s^{i,J} = \frac{3}{2}$, $\theta_i^J(W_s) = [\frac{3}{2}, 2]$, and, because the equilibrium is symmetric, $\mu_L(\theta_i^J(W_s)) = \frac{1}{2}$ for each voter i and candidate J, so $\mu(W_s) = \frac{1}{2}$. Whereas, given an equilibrium strategy profile s' in which the candidate strategy pair $((1,1,1),(1,1,1))$ and a deviation to $p_i^J = 0$, voter i votes for $-J$ or abstains if and only if $-\frac{\theta_i^J(\omega)}{3} \leq \beta - 1 = \frac{-1}{6}$ or equivalently, $\theta_i^J(\omega) \geq \frac{1}{2}$, so $w_s^{i,J} = \frac{1}{2}$ and the size of beliefs supporting $((1,1,1),(1,1,1))$ in equilibrium is $\mu(W_s) = 2 - \frac{1}{2} = \frac{3}{2}$. Therefore, the equilibrium 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 $\beta \in \left(\frac{2}{3}, 1\right)$, the equilibrium strategy profile supported by a **largest set of beliefs** is such that candidates offer the local public good to all three districts. If $\beta \in \left(\frac{1}{3}, \frac{2}{3}\right)$, the equilibrium strategy profile supported by a largest set of beliefs is such that candidates do* *not to offer the local public good to any district.*

Selection by passive or symmetric beliefs fails to recognize that the two cases with beta 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 that offers to provide the local public good to a district. The voter must have very pessimistic beliefs about the number of offers in other districts to reject the offer: few beliefs sustain the equilibrium. In contrast, if β is low, little above $\frac{1}{3}$, there is little to gain $(\beta - \frac{1}{3})$ from supporting a deviation away from the efficient policy that offers to provide the local public good to a district, and the equilibrium is sustained for almost any beliefs. Our selection criterion selects the equilibrium in which candidates do not offer public goods when these goods are very inefficient, but discards this equilibrium when these goods are almost efficient. Selection by passive or symmetric beliefs treats both situations as similar,[8] when the incentives are so different.

4.2 Experimental Design

We designed further experiments to test the prediction power of the different selection principles in the voting game presented in the previous section. We ran two treatments which only differed on the level of inefficiency of the public good. In Treatment $V1$ the public good was slightly inefficient ($\beta = 0.9$) while in treatment $V2$ the public good was highly inefficient ($\beta = 0.4$). The rest of parameters was the same across treatments. Candidates would earn 10 when losing the elections and 94 when winning elections. 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. 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

[8] The symmetric beliefs criterion admits a second equilibrium if $\beta \in (\frac{1}{3}, \frac{2}{3})$, thus failing to resolve the multiplicity problem.

Treatment V1		Total Investment in other districts		
		0	*1*	*2*
Investment in Own District	Yes	90	70	50
	No	56	36	16

Treatment V2		Total Investment in other districts		
		0	*1*	*2*
Investment in Own District	Yes	60	40	20
	No	56	36	16

Table 6: Voter's payoffs.

full feedback about actions of all subjects in their subgroup and their payoff for that round.

Experiments were conducted at the BonnEconLab of the University of Bonn in October 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. To determine payment, the computer randomly selected four periods for the final payment. The total amount earned in this periods 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. 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 politician in both treatments, aggregated by groups of five periods. As one can see, politicians 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 learn about the game, substantial differences arise. While the average investment proposal in treatment $V1$ amounts to 2.62, it is 1.80 in treatment $V2$. This difference is statistically significant (Mann-Whitney test, $z = 2.242$, $p = 0.025$).[9]

[9] *[To Do: Emphasize this. Passive predicts no difference between treatments. Symmetric predicts greater public good provision in Treatment V2 than in V1. We predict greater public good provision in treatment V1*

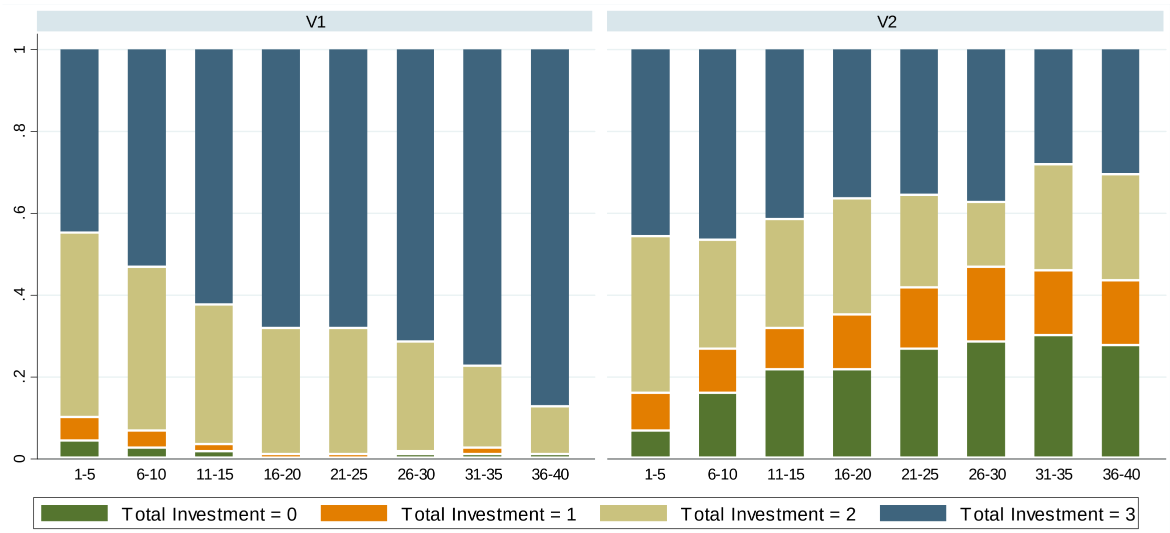

Figure 3: Investment proposals by politicians in each treatment, aggregated in groups of five periods.

Treatment	Ind Group	Period Coef.	Period St.Err.	Constant Coef.	Constant St.Err.
$V1$	1	.0101***	.0031	2.6048***	.1678
	2	.0179***	.0041	2.0760***	.1868
	3	.0182***	.0037	2.1029***	.1644
	4	.0092***	.0034	2.3356***	.1863
	5	.0170***	.0036	2.2394***	.1665
	6	.0077***	.0020	2.7240***	.0819
$V2$	7	−.0168***	.0047	1.9375***	.2775
	8	−.0076*	.0042	2.5250***	.2661
	9	.0067**	.0032	2.5683***	.1720
	10	−.0722***	.0057	2.8106***	.2992
	11	−.0095*	.0052	.6125***	.1857
	12	−.0134***	.0049	2.6442***	.3105

Table 7: GLS Random effects regression on the amount of high prices set by suppliers on the period and a constant independently for each group. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$

The clear and opposite trend in the different treatments is well captured in Table 7. Table 7 plots the results of a GLS random effects regression of the total amount of investment set by politicians as a function of the number of period, independently for each matching group.

In treatment $V1$, each and every group experiences a significant increase in the total amount of proposed investment thought time, clearly converging to the full investment equilibrium. It is also worth pointing out that behavior across different independent groups is relatively homogeneous in Treatment V1.[10] Investment proposals are substantially different in treatment $V2$. First, all but one groups experience a significant decrease in the total amount of investment proposals over time. However, we do not observe full convergence to the zero investment equilibrium. Second, we observe quite some heterogeneity across independent groups. Once can see this in the variance of converging speeds and estimated starting investment levels, but it is perhaps more evident in Figure 4. Figure 4 displays politician's proposals disaggregated by each independent group.

So far we have been describing the proposed policies made by politicians. What about the actually implemented ones? Do they differ in any significative manner from the proposals? The effect, again, varies across treatments. In treatment $V1$, we find that finally implemented policies are significantly higher that politicians proposals (Wilcoxon test, $z = 2.201$, $p = 0.0277$). As we will see in detail in next section, voters vote in such a way that they select the politicians who offer higher total investment. In treatment $V2$ instead, we don't find any significant difference between the proposed and implemented policies (Wilcoxon test, $z = 0.524$, $p = 0.6002$).

4.3.2 Voter's 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

than in V2. There is significantly greater public good provision in V1 than in V2.]

[10] Figure XXX in the Appendix displays investment proposals made by politicians disaggregated for each independent group in treatment V1.

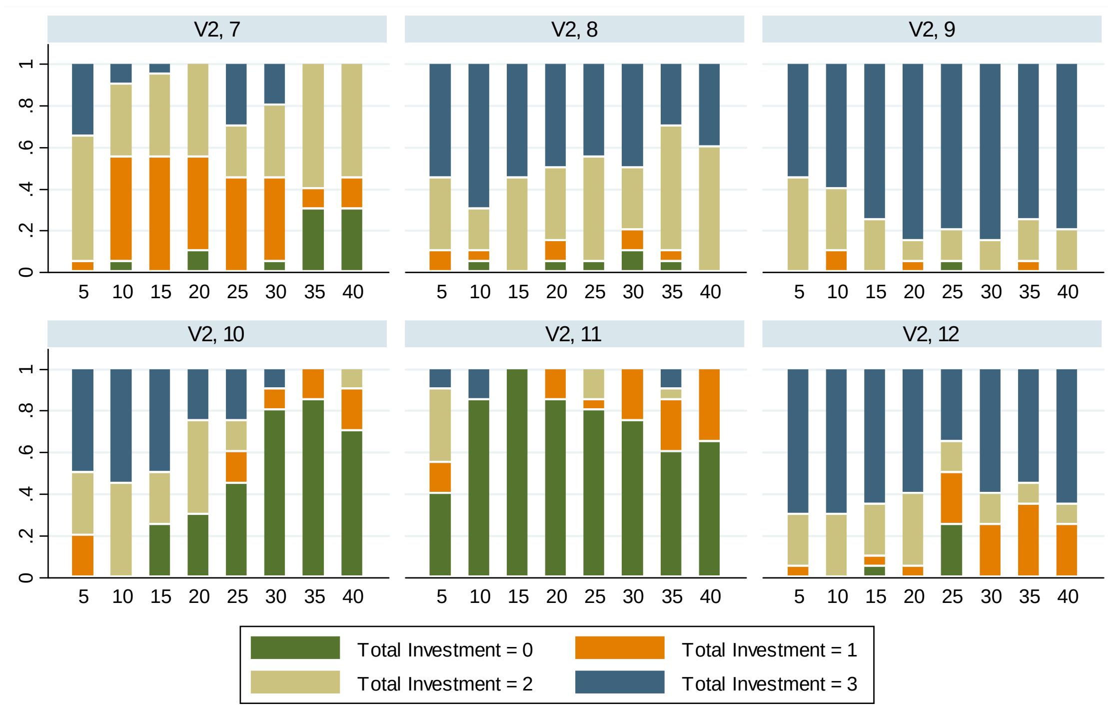

Figure 4: Investment proposals by politicians in each group in treatment V2, aggregated in groups of five periods.

			Conditional Vote					
			Treatment $V1$			Treatment $V2$		
			a	b	∅	a	b	∅
FULL INFORMATION	Best Alternative:	a	91	3	6	91	5	4
		b	3	96	1	4	92	4
		equal	38	16	45	26	26	48
PARTIAL INFORMATION	Highest : Investment	a	93	2	5	60	29	11
		b	1	96	3	27	65	8
		equal	41	29	30	36	26	38

Table 8: Conditional Vote in the different treatments.

investment in each district. This happens in 93.60% of the cases in treatment $V1$, and in 91.22% in treatment $V2$. This 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 investment in other districts. While in treatment $V1$ the percentage of subjects voting for the candidate that offer them a higher investment stays 94.63%, this percentage goes down to 62.94% in treatment $V2$. This significant difference in behavior (Mann-Whitney test, $z = 2.741$, $p = 0.0061$) relates to the different nature of the equilibrium that groups converge to. As seen in previous section, all groups in treatment $V1$ converge to the inefficient equilibrium. In such equilibrium (supported by passive beliefs), voter's should indeed vote for the candidate offering a higher investment in their district. In contrast, in treatment $V2$, some groups converge to the efficient equilibrium (supported by symmetric beliefs). In this alternative equilibrium, voters should optimally vote for the candidate offering no investment. For groups in treatment $V2$, we actually observed a significant correlation between voting the candidate offering higher investment and the implemented level of investment (Spearman's $\rho = 0.8286$, $p = 0.0416$). This difference in behavior explains, in turn, the asymmetric relation between policies proposed and implemented across treatments.

		All Data	V1	V2
Whole Group	Passive beliefs	**37.70**	58.75	16.67
	Symmetric beliefs	**7.50**	0	15.00
	Most supported	**36.88**	58.75	15.00
Only Politicians	Passive beliefs	54.68	76.25	33.13
	Symmetric beliefs	14.38	0.63	28.13
	Most supported	52.19	76.25	28.13
Only Voters	Passive beliefs	94.10	98.47	89.72
	Symmetric beliefs	83.82	86.53	81.11
	Most supported	89.79	98.47	81.11

Table 9: Percentage of events in which the whole group, the politicians or the voters behave according to the different selection criteria in the second half of the experiment.

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 a whole group behaved as prescribed by the respective equilibrium.[11] Similarly to what we did in section 3.3.3, we also include the measure separately for only politicians or only voters. Table 9 displays the results for the second half of the experiment.[12]

The results differ considerably from the results observed in contracting experiment. First, equilibrium predictions of all selection criteria have less predicting power in the voter setting than in the contracting one. Part of this might be due to the fact that convergence seems to be slower in this alternative setting, particularly in treatment $V2$. Additionally, the contracting game was played 10 extra periods.[13] The second difference 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, $z =$

[11] For the symmetric beliefs prediction, we take the symmetric equilibrium supported by symmetric beliefs. Results including all equilibria supported by symmetric beliefs are similar, because asymmetric equilibria supported by symmetric beliefs are rarely played.

[12] We restrict the measure to the second half of the experiment due to the convergence process observed in section 4.3.1.

[13] Actually, the set of experiments presented in this part were ran before the contracting experiments. In of the first results here, we decided to extend the number of periods in order to increase the likelihood of observing, if at all, convergence to equilibrium.

2.412, $p = 0.0159$). But unlike the first setting, the predictive power of our selection criterion (36.88%) doesn't differ significantly from the predictive power of passive beliefs (37.70%), (Wilcoxon test, $z = 0.443$, $p = 0.6579$).

Disaggregating by the different games sheds some light on these differences across settings. In treatment $V1$, our criterion coincides with the prediction of the equilibrium under passive beliefs. In this case, both selection criteria outperform the prediction of symmetric beliefs vastly. In treatment $V2$, our criterion coincides with the prediction of the equilibrium under symmetric beliefs. It is in this second setting where 1) the performance of any of the selection criteria differ substantially from any other and 2) all selection criteria perform poorly. Results in section 4.3.1 might explain why this is the case. On the one hand, most groups in treatment $V2$ start with relatively high levels of investment (which is consistent with the prediction of passive beliefs). On the other, all but one groups show a significant decrease of total investment over time, but most groups didn't fully converge to a stable situation within the 40 period played. Playing for additional periods might make them converge to the zero investment equilibrium (which is consistent with symmetric beliefs). The combination of both forces is probably the reason for such mixed results.

4.3.4 Voter's Beliefs

As we did with 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 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 we saw with the contracting game, we observe that there is a big difference across treatments. In treatment $V1$, 61.11% of the subjects' beliefs are in line with the prediction of passive beliefs. In treatment $V2$, still being the mode, the percentage of subjects whose beliefs are in line goes down to 25%. Surprisingly though, the percentage of subjects whose

		Total Investment in other districts if there is no Investment in own district					
		Treatment $V1$			Treatment $V2$		
		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^S	0	22^P	1^S	5	9^P

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

beliefs are in line with symmetric beliefs is also lower in treatment $V2$. Still, the significant difference in beliefs across games points in the same direction: off-path beliefs depend on the particular game.

5 Discussion

The one and only previous experiment on equilibrium selection in vertical contracting games with externalities and imperfect information that we are aware of (Martin, Normann and Snyder (9)) finds that "the downstream subjects' responses were neither uniformly consistent with symmetric beliefs nor uniformly consistent with passive beliefs." This is consistent with our results, and with our prediction: we show that instead, downstream subjects' responses are uniformly consistent with *our* selection criterion: the equilibrium played is the equilibrium that can be supported by a largest set of beliefs.

6 Appendix

Proof of Claim 2

Proof. We first prove that these strategy profiles are part of an equilibrium, and then we show that no other strategy profile is part of an equilibrium.

Case 1 $((L, L), (0, 2), (0, 2))$: In both games 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 any retailer, she decreases her payoff because the retailer buys zero unit and the deviation is not observed by the other retailer. To buy zero unit when the price offered is high is a best response for a retailer who assigns probability one that the producer is proposing a low price to the other retailer.

Case 2 $((H, H), (1, 2), (1, 2))$: In both games for each retailer to buy two units is the unique best response when the price is low 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 therefore the produces gets a payoff of 15 lower than the equilibrium payoff of 21. 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 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 do not exist asymmetric equilibria such that the producer offers a high price to a retailer and a low price to the other. 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 5

Proof. We already showed that the equilibrium supported by the largest set of beliefs in game C1. Consider game C2: the strategy profile $((H, H), (1, 2), (1, 2))$ is supported in equilibrium only by beliefs that make it a best response to purchase two units given $p_i = L$. The expected utility of purchasing two units is

$$55\omega_i(H|L) + 27(1 - \omega_i(H|L)), \tag{3}$$

whereas, the expected utility of purchasing three units is

$$57\omega_i(H|L) + 3(1 - \omega_i(H|L)), \tag{4}$$

hence purchasing two units is a best response if and only if (3) is no less than (4), or, equivalently $\omega_i(H|L) \leq \frac{12}{13}$, so $w_i^s = \frac{2}{19}$ for $i \in \{1, 2\}$ and the size of the set of beliefs that supports an equilibrium with strategy profile $((H, H), (1, 2), (1, 2))$ is $\left(\frac{12}{13}\right)^2$. The strategy profile $((L, L), (0, 2), (0, 2))$ is supported in equilibrium by beliefs such that purchasing zero units is a best response following $p_i = H$, which requires

$$0 \geq 51\omega_i(H|L) - 10(1 - \omega_i(H|L)),$$

that is, $\omega_i(H|L) \leq \frac{10}{61}$, and the size of of the set of beliefs that supports an equilibrium with strategy profile $((L, L), (0, 2), (0, 2))$ is $\left(\frac{10}{61}\right)^2$. Since this size is smaller, the equilibrium supported by a largest set of beliefs in Game C2 is the former one with strategy profile $((H, H), (1, 2), (1, 2))$. ■

Proof of Claim 7

Proof. If $s^A = s^B = (1, 1, 1)$ in equilibrium, candidates tie the election. With passive beliefs, if candidate J deviates to $\hat{s}^J$ with $s_i^J = 0$, i votes for the other candidate and J loses the election if the electorate does not become fully informed, so the deviation is not profitable. To prove uniqueness the following two remarks are useful. First, in any pure strategy equilibrium, both candidates win with equal probability. Suppose the contrary, and consider any candidates' strategy pair (s^A, s^B) such that candidate J wins with probability less than $\frac{1}{2}$. Since in equilibrium voters hold correct beliefs, the probability that J wins conditional on full information being revealed, or not revealed, is less than $\frac{1}{2}$ in each case. By assumption, voters abstain whenever indifferent, thus the probability of victory is in the set $\{0, \frac{1}{2}, 1\}$: if it is less than $\frac{1}{2}$, then it is zero. Deviating to $s^J = s^{-J}$, candidate J ties the election if full information is revealed, so the probability of winning is at least $\frac{\pi}{2}$. Since for any candidate type, to win with positive probability gives a positive payoff, there is no equilibrium where a candidate wins with zero probability. Second, in any pure strategy equilibrium both candidates propose to carry out the same number of projects. By the previous argument, in equilibrium both candidates win with equal probability. In fact, if candidates propose to implement a different number of projects, then no voter is indifferent between the two candidates if $\beta \neq \frac{2}{3}$ and hence no voter abstains, so that either A or B win for sure, which is a contradiction of the previous argument. If $\beta = \frac{2}{3}$ suppose, without loss of generality, that candidate A promises a public good to voters 1 and 2 and candidate B to voter 3. Then voters $i = 1, 2$ such that $(s_i^A, s_i^B) = (1, 0)$ are indifferent and abstain, while voters 3 prefers B, and therefore A loses for sure and therefore this strategy profile cannot be part of an equilibrium.

The previous remarks (see proof of Proposition 8 in (3) for a comprehensive argument) imply that there are four class of equilibria depending on the number $k = 0, 1, 2, 3$ of public goods promised by both candidates. It is easy to show that i) if $\beta > \frac{2}{3}$ the are not equilibria in which $k = 1$ because a candidate J wins by deviating to $\hat{s}^J = (1, 1, 1)$ and ii) if $k = 2$ then there exists one voter i such that $s_i^J = 1$ and $s_i^{-J} = 0$ (if not, a candidate who deviates and promises the public good to this voter wins with probability $1 - \pi > \frac{1}{2}$). Consider first the strategy profile $s^J = (0, 0, 0)$ for both $J = A, B$. If voter i is endowed with passive beliefs, voter i votes for candidate J if candidate J deviates to $\hat{s}_i^J = 1$ because she assigns probability one that candidate J promises only one public good,and since $\pi \in (0, \frac{1}{2})$ this deviation is profitable. Consider a strategy profile such that both candidates promise two public goods. As argued above there exist voter $i$such that $s_i^J = 1$ and $s_i^{-J} = 0$. Candidate J has a profitable deviation by promising $\hat{s}^J$ such that $\hat{s}_i^J = 0$ and $\hat{s}_k^J = s_k^J$ for $k \neq i$. If voter i is endowed with passive beliefs, voter i still votes for candidate J when candidate J deviates to $\hat{s}_i^J = 0$ because she assigns probability one that candidate J promises one public good to another district and candidate $-J$ promises a public good to both the other districts. It follows that candidate J ties the election with probability π and wins with probability $1-\pi$. Consider finally the case in which both candidates promise $k = 1$ in equilibrium (and therefore $\beta < \frac{2}{3}$). There exists one voter i such that $s_i^J = 0$ for both $J = A, B$. Candidate J has a profitable deviation by promising $\hat{s}^J$ such that $\hat{s}_i^J = 1$ and $\hat{s}_k^J = s_k^J$ for both $k \neq i$. ■

Proof of Claim 8

Proof. If $s^A = s^B = (0,0,0)$ in equilibrium, candidates tie the election. If candidate J deviates to $\hat{s}^J$ with $s_i^J = 1$, i votes for the other candidate and J loses the election if the electorate does not become fully informed, so the deviation is not profitable. If $s_i^A = s_i^B = 1$ and $s_k^A = s_k^B = 0$ for $k \neq i$, in equilibrium, candidates tie the election. A deviation such that $\hat{s}_i^J = 0$ is not profitable because voter i votes for candidate $-J$. If candidate J deviates to $\hat{s}_k^J = 1$, voter k votes for the other candidate and J loses the election if the electorate does not become fully informed, so the deviation is not profitable. Consider any equilibrium in which both candidates promise two public goods: there exists one voter i such that $s_i^J = 1$ and $s_i^{-J} = 0$. Candidate J has a profitable deviation by promising $\hat{s}^J$ such that $\hat{s}_i^J = 0$ and $\hat{s}_k^J = s_k^J$ for $k \neq i$. Voter i still votes for candidate J when candidate J deviates to $\hat{s}_i^J = 0$ because she assigns probability one that candidate J promises zero public good and candidate $-J$ promises a public good to both the other districts. It follows that candidate J ties the election with probability p and wins with probability $1-p$. Consider the strategy profiles such that $s^A = s^B = (1,1,1)$. Candidate J has a profitable deviation by promising $\hat{s}^J = (0,0,0)$. Each voter i votes for candidate J irrespective of the realization of π. Finally consider any equilibrium with $k = 1$ ($\beta < \frac{2}{3}$) and there is not any voter i such that $s_i^A = s_i^B = 1$. Consider any voter i such that $s_i^J = 1$. If candidate J deviates to $\hat{s}^J = (0,0,0)$ voter i such that $s_i^J = 1$still votes for candidate J under symmetric beliefs and therefore the deviation is profitable because candidate J ties the election with probability $1-\pi$ but wins the election with probability π. ■

Proof of Claim 9

Proof. We first introduce some useful notation: For each candidate $J \in \{A, B\}$, the set of pure strategies S^J consists of the following eight 3-dimensional vectors. $s_1 = (0,0,0)$; $s_2 = (1,0,0)$; $s_3 = (0,1,0)$; $s_4 = (0,0,1)$; $s_5 = (1,1,0)$; $s_6 = (1,0,1)$; $s_7 = (0,1,1)$ and $s_8 = (1,1,1)$. it is also useful to classify strategy pairs in classes of strategic equivalence, as follows: $S_1 = \{(s_1,s_1)\}$; $S_2 = \{(s_1,s_2),(s_1,s_3),(s_1,s_4),(s_2,s_1),(s_3,s_1),(s_4,s_1)\}$;

$S_3 = \{(s_1,s_5),(s_1,s_6),(s_1,s_7),(s_5,s_1),(s_6,s_1),(s_7,s_1)\}$;

$S_4 = \{(s_1,s_8),(s_8,s_1)\}$; $S_5 = \{(s_2,s_2),(s_3,s_3),(s_4,s_4)\}$;

$S_6 = \{(s_2,s_3),(s_2,s_4),(s_3,s_4),(s_3,s_2),(s_4,s_2),(s_4,s_3)\}$;

$S_7 = \{(s_2,s_5),(s_2,s_6),(s_3,s_5),(s_3,s_7),(s_4,s_6),(s_4,s_7),(s_5,s_2),(s_6,s_2),(s_5,s_3),(s_7,s_3),(s_6,s_4),(s_7,s_4)\}$;

$S_8 = \{(s_2,s_7),(s_3,s_6),(s_4,s_5),(s_7,s_2),(s_6,s_3),(s_5,s_4)\}$;

$S_9 = \{(s_2,s_8),(s_3,s_8),(s_4,s_8),(s_8,s_2),(s_8,s_3),(s_8,s_4)\}$; $S_{10} = \{(s_5,s_5),(s_6,s_6),(s_7,s_7)\}$;

$S_{11} = \{(s_5,s_6),(s_5,s_7),(s_6,s_7),(s_6,s_5),(s_7,s_5),(s_7,s_6)\}$;

$S_{12} = \{(s_5,s_8),(s_6,s_8),(s_7,s_8),(s_8,s_5),(s_8,s_6),(s_8,s_7)\}$; $S_{13} = \{(s_8,s_8)\}$.

Assume first that $\beta \in \left(\frac{2}{3},1\right)$. We find the measure μ on the set of beliefs W_s that can support strategy profile to sustain an equilibrium for each class in which candidates promise zero, two and three public goods respectively (and in case $k=2$ for each $i=1,2,3$ there exists J such that $s_i^J = 1$).

S_1 : An equilibrium with candidate strategy pair $(s^A,s^B) = (s_1,s_1)$ holds with any beliefs ω such that $\theta_i^J(\omega) \geq 3\beta - 1$. Whereas, if $\theta_i^J(\omega) < 3\beta - 1$, candidate J can deviate to s_2 and win the election. Thus, $\mu(W_{s_1}) = 3\beta - 1$, where abusing notation we write $\mu(W_{s_1})$ to denote the measure for any strategy profile in S_1.

S_{11} : An equilibrium with candidate strategy pair $(s^A,s^B) = (s_5,s_6)$ holds with any beliefs such that $\theta_i^J(\omega) > 1$. Whereas, if $\theta_i^J(\omega) \leq 1$, candidate A can deviate to s_8 so that if full information is not revealed, voter 1 abstains, voter 2 votes for A, and 3 abstains or votes for A. Thus, $\mu(W_{s_{11}}) = 1$

S_{13} : An equilibrium with candidate strategy pair $(s^A,s^B) = (s_8,s_8)$ holds with any beliefs such that $\theta_i^J(\omega) \geq 3(1-\beta)$. Whereas, if $\theta_i^J(\omega) < 3(1-\beta)$, candidate J can deviate to s_1 and win the election. Thus, $\mu(W_{s_{13}}) = 3(1-\beta)$.

Since $3(1-\beta) < 1 < 3\beta - 1$ for any $\beta \in \left(\frac{2}{3},1\right)$, it follows that (s_8,s_8) is the candidates' strategy profile supported by a largest set of beliefs.

Assume next that $\beta \in \left(\frac{1}{3}, \frac{2}{3}\right)$. There are two additional classes of equilibria. In both $k = 1$ for both candidates. In the class S_5 both candidates promise the public good to the same voter, while in class S_6 they promise the public good to a different voter.

S_1 : An equilibrium with candidate strategy pair $(s^A, s^B) = (s_1, s_1)$ holds with any beliefs such that $\theta_i^J(\omega) \geq 3\beta - 1$. Whereas, if $\theta_i^J(\omega) < 3\beta - 1$, candidate J can deviate to s_2 and win the election. Thus, $\mu(W_{s_1}) = 3\beta - 1$.

S_5 : An equilibrium with candidate strategy pair $(s^A, s^B) = (s_2, s_2)$ holds with any beliefs such that $\theta_i^J(\omega) \geq 3\beta$. Whereas, if $\theta_i^J(\omega) < 3\beta$, candidate J can deviate to s_5 and win the election. Thus, $\mu(W_{s_5}) = 3\beta$.

S_6 : An equilibrium with candidate strategy pair $(s^A, s^B) = (s_2, s_3)$ holds with any beliefs such that $\theta_i^J(\omega) \geq 3\beta$. Whereas, if $\theta_i^J(\omega) < 3\beta$, candidate A can deviate to s_6 and win the election. Thus, $\mu(W_{s_6}) = 3\beta$.

S_{11} : An equilibrium with candidate strategy pair $(s^A, s^B) = (s_5, s_6)$ holds with any beliefs such that $\theta_i^J(\omega) > 1$. Whereas, if $\theta_i^J(\omega) \leq 1$, candidate A can deviate to s_8 so that if full information is not revealed, voter 1 abstains, voter 2 votes for A, and 3 abstains or votes for A. Thus, $\mu(W_{s_{11}}) = 1$

S_{13} : An equilibrium with candidate strategy pair $(s^A, s^B) = (s_8, s_8)$ holds with any beliefs such that $\rho^{i,J}(0) \geq 3(1-\beta)$. Whereas, if $\rho^{i,J}(1) < 3(1-\beta)$, candidate J can deviate to s_1 and win the election. Thus, $\mu(W_{s_{13}}) = 3(1-\beta)$

Since $3\beta - 1 < \min\{3\beta, 1, 3(1-\beta)\}$ for any $\beta \in \left(\frac{1}{3}, \frac{2}{3}\right)$, (s_1, s_1) is the candidates' strategy profile supported by a largest set of beliefs. ■

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] In-Koo Cho and David M. Kreps. 1987. "Signaling Games and Stable Equilibria." *Quarterly Journal of Economics* 102(2): 179-222.

[3] Jon X. Eguia and Antonio Nicolò. 2012. "On the Efficiency of Partial Information in Elections." SSRN Working paper 1931040.

[4] Urs Fischbacher. 2007. "z-Tree - Zurich Toolbox for Readymade Economic Experiments", *Experimental Economics*, 10: 171-178.

[5] 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.

[6] Alessandro Gavazza and Alessandro Lizzeri. 2009. "Transparency and Economic Policy." *Review of Economic Studies* 76 (3): 1023-1048.

[7] Oliver Hart and Jean Tirole. 1990. "Vertical Integration and Market Foreclosure," *Brookings Papers on Economic Activity*, Special Issue, 205–276.

[8] David M. Kreps and Robert Wilson. 1982. "Sequential Equilibria," *Econometrica* 50: 863-894.

[9] Stephen Martin, Hans-Theo Normann and Christopher M. Snyder. 2001. "Vertical Foreclosure in Experimental Markets." *RAND Journal of Economics* 32(3): 466-496.

[10] R. Preston McAfee and Marius Schwartz. 1994. "Opportunism in Multilateral Vertical Contracting: Nondiscrimination, Exclusivity, and Uniformity." *American Economic Review* 84: 210-230.

[11] Daniel P. O'Brien and Greg Shaffer. 1992. "Vertical Control with Bilateral Contracts." *RAND Journal of Economics* 23(3): 299-308.

[12] Marco Pagnocci and Salvatore Piccolo. 2011. "Vertical Separation with Private Contracts." *Economic Journal* 122: 173-207.

[13] Patrick Rey and Jean Tirole. 2007. "A Primer on Foreclosure" in *Handbook of Industrial Organization* (eds. M. Amstrong and R. Porter) vol. 3: 2145-2220.

[14] Patrick Rey and Thibaud Vergé. 2004. "Bilateral Control with Vertical Contracts." *RAND Journal of Economics* 35(4): 728-746.

[15] Ilya Segal. 1999. "Contracting with Externalities." *Quarterly Journal of Economics* 114: 337-388.

[16] Ilya Segal and Michael D. Whinston. 2003. "Robust Predictions for Bilateral Contracting with Externalities." *Econometrica* 71(3): 757–791.