THE UNIVERSITY OF MELBOURNE

Department of Economics

Working Paper Series

Location Choice and Information Transmission

Simon Loertscher

January 2012

Research Paper Number 1137

ISSN: 0819-2642
ISBN: 978 0 7340 4488 4

Department of Economics
The University of Melbourne
Parkville VIC 3010
www.economics.unimelb.edu.au

Location Choice and Information Transmission

Simon Loertscher*

January 13, 2012

Abstract

This paper analyzes a location game with two agents and one decision maker. The decision maker is less well informed about her bliss point location than are the agents. Therefore, the locations the agents choose can, in principle, transmit relevant information to the decision maker. However, if information is precise relative to the prior, no equilibrium exists that satisfies the Intuitive Criterion and in which all the information of the agents is transmitted. In contrast, given sufficiently noisy information, a refined equilibrium with full information transmission exists. Though the agents play pure strategies in this equilibrium, their location choices differ with positive probability because their information is different with positive probability.

Keywords: Location Games, Sender Receiver Games, Information Transmission, Signaling, Political Economics, Refinements.

JEL-Classification: C72, D72, D83

*Department of Economics, Arts West, University of Melbourne, Victoria 3010, Australia. Email: simonl@unimelb.edu.au. I want to thank Aleksander Berentsen, Esther Bruegger, Yuelan Chen, Sven Feldmann, Emmanuel Frot, Roland Hodler, Rosemary Humberstone, Navin Kartik, Pei-yu Lo, Andy McLennan, Tom Palfrey, Carlos Pimienta, Randy Silvers, Tom Wilkening and seminar participants at AETW 2008 (Gold Coast), ESEM 2008 (Milano), AEA 2009 (San Francisco), ESAM 2009 (Canberra), APCC 2009 (Melbourne), IIOC 2010 (Vancouver), at ANU, Deakin University, the Universities of Auckland, Bern, Bielefeld and Mannheim for most valuable comments and discussions. A previous version of this paper circulated under the title "Information Transmission by Imperfectly Informed Parties".

1 Introduction

Location games are fundamental to political economics and industrial organization. A central theme in the literature has been whether competing agents who cater to decision makers (like voters or consumers) will choose to locate at the bliss point of the median decision maker.[1] Departing from the standard assumption that decision makers know their bliss points with certainty, this paper asks: Will agents who compete in locations for a decision maker's favor reveal or conceal their superior private information about what is optimal for the decision maker?

The model that is set up and analyzed in this paper has a fairly broad range of applications. Consider, for example, the portfolio choice problem of a risk-averse investor faced with two portfolios offered by two competing financial advisors who can tailer the portfolio to the expected needs of the investor. Each financial advisor is paid a fixed fee if he gets the deal and nothing otherwise. The investor is less well informed than the financial advisors about the state of the economy that determines which point on the risk-return locus is optimal for the risk-averse investor. Alternatively, political candidates, whose expertise arguably exceeds that of the median voter, compete for the median voter's favor by campaigning with policy platforms that may (or in equilibrium may not) convey their superior information. Very similarly, the competing agents can be thought of as experts who participate in a debating contest whose winner is determined by a jury. The jury has a prior and picks as winner the agent whose position is closest to her bliss point given the jury's updated beliefs. The sole goal of each participant is to win the contest. Lastly, consider education programs that compete for students through the curricula they publicize. Prospective students are uncertain about what is the optimal curriculum, which depends on a state they know less about than the schools who compete for students.

The following is a description of the model. There are two agents and a decision maker (DM). First, each agent observes a private signal a or b, indicating which of the two states A or B has been realized. The signals of both agents have the same precision, and conditional on the state, each signal is correct with probability greater than a half, and incorrect otherwise. Second, both agents simultaneously choose locations

[1] More generally and more fundamentally, a question of much interest is whether this kind of competition will lead to too much, too little or the socially optimal degree of differentiation. The literature and these themes originate with Hotelling (1929), who famously but erroneously argued that two firms who first choose locations on a line and then prices will locate at the median location when consumers face linear transportation costs (d'Aspremont, Gabszewicz, and Thisse, 1979). Absent price competition, Hotelling's principle of minimum differentiation is valid.

within an interval. The objective of each agent is to be selected by the DM. Observing the locations, the DM then updates her beliefs and selects the agent whose location maximizes her expected utility, given the updated beliefs.[2] The DM's expected utility function is well-behaved and such that the associated value function is symmetric in the belief and minimized with a uniform belief. In particular, the expected utility function has a unique bliss point location that varies monotonically with her belief about the states. If the DM is indifferent between the two locations chosen by the agents, she randomizes uniformly.

The following results are obtained under the assumption that the DM employs a symmetric strategy.[3] First, in any such equilibrium, conditional on his signal each agent is selected with equal probability independently of the signal. Second, for any precision of the signals, there is a continuum of equilibria that are pooling and satisfy the Intuitive Criterion of Cho and Kreps (1987) and the D1 criterion of Cho and Kreps (1987) and Banks and Sobel (1987) appropriately extended to sender-receiver games with two senders. Among these intuitive pooling equilibria there is a unique equilibrium satisfying the PSE refinement based on Grossman and Perry (1986). In this equilibrium both agents pool at the DM's bliss point location given the prior, and thus pander to the ignorance of the decision maker. Consequently, PSE selects the welfare optimal pooling equilibrium.

A key to the equilibrium analysis is whether both signals are strong, where a signal is called strong (weak) if the probability that the other agent has received the same signal exceeds (is less than) one half.[4] In any candidate separating equilibrium this probability is the information an agent has about the likely location choice of his competitor. The third result is that when both signals are strong and the prior is non-uniform, there is no separating equilibrium that satisfies the Intuitive Criterion. Fourth, when one signal is weak, separating equilibria satisfying the Intuitive Criterion exist, and there is a unique separating PSE and a unique separating D1 equilibrium. These results are robust to the introduction of cheap talk messages and a multi-dimensional location space. However, a separating PSE always exists if agents move sequentially rather than simultaneously.

The intuition for the main results, though subtle at times, can be developed rather straightforwardly. The reason that both the Intuitive Criterion and D1 have no bite

---

[2]Throughout the paper the DM is referred to as a "she" while each agent is referred to as a "he".

[3]The DM's strategy is called symmetric if the probability that she selects agent 1 given that agent 1 chooses location x and agent 2 chooses y is equal to the probability that she selects agent 2 if 1 chooses y and 2 chooses x.

[4]A signal is strong (weak) if the probability, conditional on the state, that the signal is correct exceeds (is less than) the common prior on that state. Consequently, at least one signal is always strong.

in refining pooling equilibria is essentially that there are no signaling costs in such an equilibrium. Any agent could benefit from a deviation independently of the signal. On the other hand, the PSE refinement has bite because it requires the DM to assign prior preserving beliefs. Consequently her belief will equal the prior both on and off the equilibrium path. Thus, in a pooling equilibrium the DM's bliss point will be the same both on and off the equilibrium path.

To see why there is no intuitive separating equilibrium when both signals are strong, consider agent 1 who has received, say, signal a and assume that agent 2 plays as prescribed by the separating equilibrium under investigation. Given that a is a strong signal, agent 1 would be willing to deviate from a candidate separating equilibrium strategy if this deviation induced the DM to select him with probability one if agent 2 also received signal a and to select him with probability zero if agent 2 received signal b. The reason is simply that the probability that the other agent has received the same signal exceeds one half because the signal is strong, and that on the equilibrium path each agent wins with probability of one half. By the same token, upon signal b agent 1 has no incentive to deviate to a location that yields a win with probability one if the other agent has received signal a and with probability zero otherwise.[5] That such profitable deviations generically exist follows essentially from the assumption that the DM's utility function is well-behaved.

The reason why there are separating intuitive equilibria when one signal is weak is that now the probability that the other agent has received the strong signal exceeds one half upon receiving both the strong and the weak signal. Therefore, regardless of his own signal an agent would now benefit from the same deviation if this deviation defeats the other agent if and only if the other agent has received the strong signal. This eliminates the additional constraint that hinders the existence of separating intuitive equilibria when both signals are strong and the prior is non-uniform.

The present paper contributes to the literature on location games and to the literature on information transmission games. These two literatures have had little overlap so far, with a few notable exceptions discussed below. The standard assumption in the location games literature is that there is no uncertainty (beyond the strategic uncertainty inherent to any game).[6] The cheap talk strand of the literature on information transmission games

[5]Analogously, upon signal b agent 1 would have an incentive to deviate to a location that allows him to win with probability one if agent 2 has received signal b as well and with probability zero otherwise, and agent 1 would have no incentive to deviate to this location had he received signal a.

[6]See, for example, Hotelling (1929), Lerner and Singer (1937), Downs (1957), d'Aspremont, Gabszewicz, and Thisse (1979), Prescott and Visscher (1977), Osborne (1995), Callander (2005) and Loertscher and Muehlheusser (2008, 2011).

assumes that the messages agents send are intrinsically meaningless. That is, messages that may convey information do no affect any player's utility.[7,8]

Schultz (1996) and Callander (2008) provide models that combine location choice and information transmission. In these regards, these two papers are the most important precursors to the present paper. Schultz (1996) analyzes a model with two states in which two primarily policy motivated parties know the state (that is, whether the cost of production for a public good is high or low) while the median voter is uncertain about the state. The bliss point policies of the two parties are on different sides of median voter's bliss point for either state. Schultz finds that a separating equilibrium exists if the difference in parties' preferences is not too large.[9] Callander (2008) analyzes an agency model in which two political candidates compete for office by choosing locations on the real line. Candidates care both about policy and office. Upon being elected, a politician must exert effort to provide a public good. Candidates' bliss point policies are symmetrical on opposite sides of the median voter's bliss point. Each candidate's cost of exerting effort is his private information, and it is this private information about an essentially vertical quality that can potentially be signaled to the electorate in equilibrium.

In both Callander's and Schultz's model, policy proposals serve the dual role of potentially conveying information to the decision maker and of constituting the menu the decision maker can choose from as do locations in the present paper. The main difference between Schultz's model and the present one is that here the agents also face uncertainty about the state, and thus each other's type. This uncertainty is key for the equilibrium behavior in the present paper. In Callander's model, the decision maker faces uncertainty about candidates' costs (or motivation) and thus about which candidate is better for her. In contrast, the decision maker is uncertain about her bliss point in the present model. Moreover, agents' payoffs do not directly depend on the locations they choose in this paper whereas this is the case both in Schultz (1996) and Callander (2008).[10]

[7]See, for example, Crawford and Sobel (1982), Krishna and Morgan (2001a,b), Battaglini (2002) and Ambrus and Takahashi (2008).

[8]There is also a mostly very recent literature on models in which senders experience a cost of lying; see, for example, Banks (1990), Callander and Wilkie (2007), Kartik, Ottaviani, and Squintani (2007) and Kartik (2009). The difference between these models and the present one is that here messages are costless to agents but costly to the DM because they constrain her choice set.

[9]Jensen (2011) analyzes a similar model with two states. The state is known by the two parties, but the voters are uncertain about the state. Assuming a uniform prior and that one party is the voter's preferred government in one state and the other party in the other state, he finds that no separating equilibrium equilibrium exists that satisfies a refinement based on Bagwell and Ramey (1991).

[10]Put differently, in Schultz's model there is a state that is common for the agents and the decision maker but only known by the agents. In Callander's model, there is no such common state. The agents

Pandering by better informed, self-interested agents to the preferences of some decision maker(s) has recently been the focus of a number of papers, mostly in political economics. For example, the equilibrium behavior in Canes-Wrone, Herron, and Shotts (2001), Maskin and Tirole (2004), and Hodler, Loertscher, and Rohner (2010) is such that a single agent (that is, an incumbent in political office up for re-election) chooses costly actions that make the median voter re-elect him more often than she would absent such wasteful behavior.[11] Equilibrium pandering by two competing agents in a setup with two states and two actions has first been analyzed by Heidhues and Lagerlöf (2003). Beyond its motivation, the present paper shares with Heidhues and Lagerlöf assumptions about signal technology and the symmetry restriction for the decision maker's strategy.[12] In concurrent work, Felgenhauer (forthcoming) introduces a third populist candidate who has the same information as the decision maker and who always chooses the uninformed DM's bliss point action in a setup that is otherwise identical to Heidhues and Lagerlöf's. He shows that under these assumptions a fully separating equilibrium exists.

The paper also relates to the literature on signaling games and refinements developed for these games.[13] Its contribution to the refinements literature is of some technical interest and best explained in the terminology of this literature. The paper adapts standard equilibrium refinements that have been developed for signaling games with one sender whose type is not known by the receiver (here otherwise referred to as decision maker) and one non-stochastic receiver to a game with two senders (which are called agents in this paper) whose types are not known by the receiver (an agent's type being given by the signal he received in this paper). Because here there are two senders whose types are their private information, each sender is also uncertain about the other sender's type. The underlying idea for applying refinements defined for standard signaling games with one sender and one deterministic receiver to a game with two senders is fairly natural and as follows. For any given equilibrium, fix one sender's play to be according

are of different types and privately informed about this. In both models, the agents' payoffs depend on the state and the locations chosen, and on the identity of the winner of the contest. In the present model, there is a common state as in Schultz's model but the agents have private information about this state. Their payoffs depend only on the identity of the winner of the contest.

[11] In a cheap talk model with one agent, Che, Dessein, and Kartik (forthcoming) also obtain pandering towards the DM's preferred alternative. Pandering arises as an endogenous discrimination by the agent that is to the benefit of the alternative that is being discriminated against by making the recommendation of the discriminated alternative credible for the DM when this alternative is recommended.

[12] In contrast to the present model and the one of Heidhues and Lagerlöf (2003), Laslier and Van Der Straeten (2004) assume that the decision maker also receives a private signal. They find that the unique refined equilibrium outcome is separating. Cummins and Nyman (2005) analyze equilibrium information transmission by competing agents whose interests are partly aligned with those of the decision maker.

[13] For an overview of the signaling literature, see, for example, Riley (2001) and Sobel (forthcoming).

to his equilibrium strategy. For the other sender, say sender 1, this induces a transformed signaling game that is standard insofar as there is now only one sender – himself – facing one receiver. The only piece left that is potentially non-standard is that from the point of view of sender 1, the receiver may now have stochastic type because the information available to her will be a function of the information transmitted to her by sender 2. Hence, the receiver's beliefs and thus her preferred location will vary with her type. However, the distribution over her types is fully pinned down by sender 2's equilibrium strategy, the signal technology and the common prior. Therefore, the receiver's type in the transformed game can be interpreted as drawn by nature from a commonly known distribution. Assuming that the receiver holds the same beliefs about the deviating sender's type independently of her own type, standard refinements can then be readily applied to this transformed game.

In a standard signaling game such as Spence (1973)'s, the types of the typically single agent who can try to signal his type to a single decision maker differ with respect to their exogenously given costs in the choice variable. In equilibrium, separation can occur because of these cost differentials. There are no such exogenously given cost differences in the present model. Nevertheless, credible signaling can occur here on and off the equilibrium path because, depending on the nature of the equilibrium, the uncertainty each agent faces about the other agent's type makes some actions costlier in expectation for one type than for the other type, resulting in credible signaling on and off the equilibrium path.

The remainder of this paper is organized as follows. Section 2 introduces the basic model. Section 3 derives preliminary results, including perfect Bayesian equilibria (PBE), and introduces the refinements that are applied subsequently. The equilibrium analysis based on these refinements is performed in Section 4. Section 5 analzyes extensions, and Section 6 concludes. All the proofs are in the appendix.

2 The Model

Consider the following location game that is standard in all ways but one. There are two agents, labeled 1 and 2, and one decision maker DM. The agents' action space is an interval $X \subset \mathbb{R}$. The two agents $i = 1, 2$ simultaneously choose their locations $x_i \in X$. The novel feature is that there are two states of the world $\omega \in \{A, B\}$, which satisfy $A < B$, and that the DM's bliss point location depends on the state of the world. Both the agents and the DM are uncertain about ω, in slightly different ways as explained shortly.

Utility The DM's utility when she "consumes" a good at location x in state ω is denoted by $v(\omega, x)$. The different states are assumed to correspond to shifts in the DM's bliss point in the sense that $v(A, x) = v(B, x + B - A)$ with $A, B \in X$ representing the bliss points in state A and B, respectively. Further, $v(\omega, x)$ is assumed to be twice differentiable and strictly concave in x for either state ω and symmetric around ω. That is, $v(\omega, \omega + x) = v(\omega, \omega - x)$ for any x.

The agents and the DM have the common prior $\alpha \in (0, 1)$ that state A is true.[14] Given a belief μ that the state is A, the DM's expected utility is

$$u(\mu, x) := \mu v(A, x) + (1 - \mu)v(B, x), \tag{1}$$

whose maximizer is denoted $x(\mu)$ and satisfies $u_2(\mu, x(\mu)) = 0$.

The assumptions on $v(\omega, x)$ imply that $x(\mu)$ is unique and monotonically decreases in μ and that $x(1) = A$ and $x(0) = B$.[15] Moreover, for any $x \in [A, x(\mu))$, there is a $\overline{x}(x) \in X$ with $\overline{x}(x) > x(\mu)$ such that $u(\mu, x) = u(\mu, \overline{x}(x))$ and, analogously, that for any $x \in (x(\mu), B]$ there is a $\underline{x}(x) \in X$ satisfying $\underline{x}(x) < x(\mu)$ such that $u(\mu, x) = u(\mu, \underline{x}(x))$. Of course, $\underline{x}(\overline{x}(x)) = \overline{x}(\underline{x}(x)) = x$, $\overline{x}'(x) < 0$ and $\underline{x}'(x) < 0$ holds. A special case of the present model that is widely used in political economics and industrial organization is the model with quadratic utility $v(\omega, x) = -(\omega - x)^2$ (see, for example, d'Aspremont, Gabszewicz, and Thisse, 1979; Crawford and Sobel, 1982; Krishna and Morgan, 2001a).[16] An additional restriction on $u(\mu, x)$ and its maximizer $x(\mu)$ will be imposed below after updated beliefs and optimal locations given these updated beliefs are introduced.

[14] Under the standard assumptions that all decision makers update in the same way and that decision makers' preferences are singled peaked, this can be viewed as a shortcut to a model with many decision makers who differ with respect to some characteristic such as income, where the median decision maker would be the decisive DM. Let N_{DM} be the number of DMs and assume that N_{DM} is odd. Modify the utility function to be $u(\mu, x, \theta_D)$, where θ_D is DM D's type for $D = 1, .., N_{DM}$. Then label DMs in increasing order, so that $\theta_1 < .. < \theta_{N_{DM}}$. Assuming, as above, $u_2 > 0$ and $u_{22} < 0$, the function has a unique maximizer, denoted $x^*(\mu, \theta_D)$. The sign of $dx^*/d\mu$ and $dx^*/d\theta_D$ is the same as that of u_{12} and u_{23}, respectively. So $x^*(\mu, \theta_D)$ will be monotone in μ and θ_D if, as is assumed now, u_{12} and u_{23} have constant signs. Consequently, for any belief μ, the bliss point locations $x^*(\mu, \theta_1), .., x^*(\mu, \theta_{N_{DM}})$ can be ordered monotonically. Without loss of generality assume $x^*(\mu, \theta_1) < ... < x^*(\mu, \theta_{N_{DM}})$. If all DMs update in the same manner, the model reduces to the median decision maker model analyzed here.

[15] The first and second order conditions are $\mu v_2(A, x(\mu)) + (1 - \mu)v_2(B, x(\mu)) = 0$ and $\mu v_{22}(A, x(\mu)) + (1 - \mu)v_{22}(B, x(\mu)) < 0$, respectively, where subscripts denote the argument with respect to which partial derivatives are taken. Totally differentiating the first order condition and using the second order condition reveals that $x'(\mu) < 0$.

[16] The model with linear transportation cost, where $v(\omega, x) = -|\omega - x|$, which was first proposed by Hotelling (1929), violates the assumption of global differentiability in x. More importantly, the linear utility model fails to satisfy the property that $x(\mu)$ varies smoothly with μ since $x(\mu) \in \{A, B\}$ for almost all μ.

Information Agents are better informed than the DM in the following sense. Prior to making his location choice, each agent i receives a private signal $s_i \in \{a, b\}$ indicating, respectively, that state A or B has materialized.[17] Conditional on the state, the signal is correct with probability $1 - \varepsilon$ and incorrect with probability ε, where $0 < \varepsilon < 1/2$. For simplicity, assume that signals s_1 and s_2 are independent, conditional on the state.[18] These signals are soft information, so that they cannot be communicated directly to outsiders. All of this is common knowledge. As each agent i receives a signal $s_i \in \{a, b\}$, it is sometimes convenient to call an agent with signal $k \in \{a, b\}$ an agent of type k.

Strategies The action set for each agent is X. The DM's action set is to select one of the two agents as the winner of this contest. The DM's objective is to maximize her utility given her updated beliefs. Each agent's objective is to win the contest.

A pure strategy for agent i is a location x_i^k for each private signal $k \in \{a, b\}$. A natural equilibrium concept is Perfect Bayesian Equilibrium (PBE), and attention is restricted throughout this paper to PBE in which the DM's strategy does not depend on the labeling of the agents.[19] That is, denoting by $\gamma(y, z)$ the probability that the DM selects agent 1 when agent 1 plays y while 2 plays z, the symmetry assumption is that $\gamma(z, y) = 1 - \gamma(y, z)$. Observe that for $y = z$ this implies $\gamma(y, z) = 1/2$. Without this assumption, it is easy to construct PBE in which both agents play separating strategies, and in which the DM selects agent 1 with probability one independently of the locations chosen on the equilibrium path. These equilibria are fully revealing but arguably not very compelling as they do not reflect any notion of competition between agents. With the exception of an extension in Section 5, the focus of the paper is, for analytical tractability, on equilibria in which the agents play pure strategies.

Beliefs and Optimal Locations Let $\mu(x_1, x_2)$ denote the DM's posterior belief that the state is A when agents 1 and 2 choose locations x_1 and x_2, respectively.[20] In a slight abuse of notation let $\mu(a, a)$ and $\mu(b, b)$ be the DM's – hypothetical – belief that the state is A if both agents have received the signal a and b, respectively. Analogously, let $\mu(a, b) = \mu(b, a)$ be this hypothetical belief when the agents receive divergent signals,

[17]Throughout the paper signals are denoted with lower case and states with upper case letter. So k is the signal indicating state K is true with $k \in \{a, b\}$ and $K \in \{A, B\}$.

[18]Signals are said to be independent if conditional on state K agent i expects j to get the signal $s_j = k$ with probability $1 - \varepsilon$ and the signal $s_j = -k$ with probability ε, independent of the signal s_i i has got, with $k \in \{a, b\}$ and $-k \neq k$.

[19]See also Heidhues and Lagerlöf (2003) who impose an analogous restriction for very similar reasons.

[20]Beliefs of the DM are denoted by μ. Agents' beliefs about each other's signal, which are key to the model, are denoted by π.

and denote by $\mu(a,0) = \mu(0,a)$ and $\mu(b,0) = \mu(0,b)$ the belief on A under the hypothesis that the DM knows that one agent has received the signal a or b, respectively, without knowing the other agent's signal, which is denoted by 0. Due to the above assumptions about the signals, it is true that

$$\begin{aligned}\mu(a,a) &= \frac{\alpha(1-\varepsilon)^2}{\alpha(1-\varepsilon)^2 + (1-\alpha)\varepsilon^2} > \mu(a,0) = \frac{\alpha(1-\varepsilon)}{\alpha(1-\varepsilon) + (1-\alpha)\varepsilon} > \mu(a,b) = \alpha \\ &> \mu(b,0) = \frac{\alpha\varepsilon}{\alpha\varepsilon + (1-\alpha)(1-\varepsilon)} > \mu(b,b) = \frac{\alpha\varepsilon^2}{\alpha\varepsilon^2 + (1-\alpha)(1-\varepsilon)^2}. \end{aligned} \tag{2}$$

It is also useful to denote by $x(k,k) := x(\mu(k,k))$ the DM's preferred locations if both signals are k with $k \in \{a,b\}$. Similarly, let $x(k,0) := x(\mu(k,0))$ denote the DM's bliss point if she knows or correctly infers that one agent has received the signal $k \in \{a,b\}$ while the other agent's signal is not known. Notice that the dependence of $\mu(k,0)$ and $\mu(k,k)$ on α (and ε) has been suppressed for notational ease.

The symmetry of $v(\omega,x)$ around ω and Bayes' rule imply that $u(\alpha, x(a,a))$ intersects with $u(\alpha, x(b,b))$ at $\alpha = 1/2$ for any $\varepsilon \in (0,1/2)$. Similarly, for any $\varepsilon \in (0,1/2)$ $u(\alpha, x(a,0)) = u(\alpha, x(b,0))$ holds at $\alpha = 1/2$. Throughout this paper, it is further assumed that $\alpha = 1/2$ is the unique point of intersection of $u(\alpha, x(a,a))$ with $u(\alpha, x(b,b))$ and of $u(\alpha, x(a,0))$ with $u(\alpha, x(b,0))$ for α on $(0,1)$. This is, for example, the case with quadratic utility.

3 Preliminaries and Refinements

This section first derives preliminary results that are helpful for the subsequent analysis. Second, it adapts standard refinements that were introduced for games with one sender and one receiver to games with two senders. These adjusted refinements are the basis of the analysis in Section 4.

3.1 Preliminary Results

Weak and Strong Signals It is useful to distinguish between what can be called weak and strong signals. Denote by $\pi_i(s_j = k \mid s_i = k)$ the probability that agent j has received signal $s_j = k$ given that i has received the signal $s_i = k$ with $k \in \{a,b\}$ and $i \neq j$. Signal k is quite naturally said to be strong if $\pi_i(s_j = k \mid s_i = k) > 1/2$ and said to be weak if $\pi_i(s_j = k \mid s_i = k) < 1/2$.

Lemma 1 *For $1-\alpha > \varepsilon$, signal b is strong, while for $\alpha > \varepsilon$, signal a is strong.*

Figure 1: Weak and strong signals.

The conditions under which a signal is strong or weak are illustrated in Figure 1. To see why the relation between α and ε matters, notice that if $\alpha < \varepsilon$, signal a is more likely to be due to an error: The posterior that $\omega = A$ given $s_i = a$ is $\pi_i(\omega = A|s_i = a) = \frac{(1-\varepsilon)\alpha}{(1-\varepsilon)\alpha+(1-\alpha)\varepsilon}$, which is less than 1/2 since $\alpha < \varepsilon$ implies $(1-\varepsilon)\alpha < (1-\alpha)\varepsilon$.[21] According to Lemma 1, both signals are strong if and only if $\alpha \in (\varepsilon, 1-\varepsilon)$. Note also that $\pi_i(\omega = K \mid s_i = k) > 1/2$ if and only if k is a strong signal, where $\pi_i(\omega = K \mid s_i = k)$ is the probability that the true state is K when the signal is k. To ease the notation, let $\pi(k|k) = \pi_i(s_j = k \mid s_i = k)$. It is also useful to note that $\pi(a|a) > 1 - \pi(b|b)$ for any $(\alpha, \varepsilon) \in (0,1) \times (0, 1/2)$.[22]

Perfect Bayesian Equilibria (PBE) Unsurprisingly the present game exhibits many PBE. First, some general properties of PBE are characterized, and then it is shown that, among other types of PBE, there is a continuum of pooling PBE and a continuum of separating PBE.

Normalize each agent's payoff of winning to one and the payoff of losing to zero and denote by x_i^k the location equilibrium prescribes i to choose upon signal k with $k \in \{a, b\}$ and $i = 1, 2$. Letting $U_i[x_i^k \mid s_i = l]$ be the expected payoff to i, conditional on having received signal l, when he plays x_i^k and when j is presumed to play according to his equilibrium strategy, with $j \neq i$. By the definition of an equilibrium, the incentive constraint

$$U_i[x_i^k \mid s_i = k] \geq U_i[x_i^l \mid s_i = k] \tag{3}$$

has to hold in any PBE for else i would be better off playing x_i^l upon signal k than the prescribed action x_i^k for $l \neq k$.

An implication of the assumption that the DM's strategy is symmetric is the following:

Lemma 2 *In any PBE in which the DM's strategy is symmetric, for* $k \in \{a, b\}$

$$U_i[x_i^k \mid s_i = k] = U_j[x_j^k \mid s_j = k] = 1/2. \tag{4}$$

[21] Of course, $\pi_i(\omega = A|s_i = a) = \mu(a, 0)$.

[22] The formulas for $\pi(k|k)$ with $k \in \{a, b\}$ are in the proof of Lemma 1 in the appendix. Algebra reveals that $\pi(a|a) - (1 - \pi(b|b)) = \frac{(1-\alpha)\varepsilon(1-2\varepsilon)^2}{\alpha(1-\alpha)(1-2\varepsilon)^2+(1-\varepsilon)\varepsilon} > 0$.

A PBE is called separating if $x_i^a \neq x_i^b$ for both $i = 1, 2$. An important property of separating PBE is the following:

Lemma 3 *$x_i^a = x^a$ and $x_i^b = x^b$ for $i = 1, 2$ are strategies of a separating PBE only if*

$$u(\alpha, x^a) = u(\alpha, x^b). \tag{5}$$

Condition (5) is very similar to the arbitrage condition underlying the equilibrium construction in the cheap talk models of Crawford and Sobel (1982) and Krishna and Morgan (2001a). A subtle but important difference is that here it applies to the DM rather than the agents. The intuition is clear. Suppose (5) were violated, for example, because $u(\alpha, x^a) > u(\alpha, x^b)$. Then playing x^a independently of the signal would be a profitable deviation because it guarantees selection with probability 1 if the other agent plays x^b and with probability 1/2 if the other one plays x^a, so that overall the probability of being elected is greater than 1/2. So for agents to be willing to separate, $u(\alpha, x^a) = u(\alpha, x^b)$ has to hold. Lemma 3 is important for the results that follow which show that (5) imposes a condition that can be made to hold when one signal is weak but typically not when both signals are strong.

Proposition 1 *There is a continuum of pooling PBE, where $x_i^a = x_i^b = x$ for $i = 1, 2$. There is also a continuum of separating PBE, where $x_i^a = x^a$ and $x_i^b = x^b$ for $i = 1, 2$ and where x^a and x^b are such that $u(\alpha, x^a) = u(\alpha, x^b)$ and $x^a \neq x^b$, that is, $x^b = \overline{x}(x^a)$.*

Separating PBE do not exist in Heidhues and Lagerlöf (2003)'s model, where the agents have a binary action set. They arise here because of the continuous strategy space and the concavity properties of $u(\alpha, x)$. Proposition 1 is not a complete description of all types of PBE in the model. For example, there are also pooling PBE in which one agent plays x^a and the other one x^b, where x^a and x^b satisfy $u(\alpha, x^a) = u(\alpha, x^b)$, and there may be hybrid PBE in which one agent does not reveal his signal while the other one does. This raises the question whether some of these PBE are more plausible than others, which is addressed in the next section. Before doing so, it is useful to state and prove the following lemma.

Lemma 4 *When both signals are strong, there are no separating PBE with $x^a, x^b \notin [x(a, a), x(b, b)]$.*

Off the equilibrium path the DM's belief about the strategy of the deviating agent is not pinned down. However, since an agent's strategy is a mapping from signals to locations, there are limits to her beliefs about the state. To see this, suppose that agent 2 plays

$x^a < x(a,a)$ according to equilibrium while agent 1 plays an off the equilibrium path location $x > x^a$. Now the most favorable belief for agent 2 is that agent 1 played x if and only if his signal was a, in which case her belief about the state is $\mu(a,a)$. So by choosing $x = x(a,a)$ agent 1 has a deviation that guarantees winning with probability 1, conditional on agent 2 receiving signal a and conditional on agent 2 playing according to equilibrium. Since both signals are strong, upon signal a agent 1 has a posterior exceeding 1/2 that agent 2 received the same signal. Since on equilibrium each agent wins with probability 1/2 independently of the signal (Lemma 2), the deviation pays off. For very similar reasons, the Intuitive Criterion will impose the tighter constraint that separating equilibrium locations actually satisfy $x^a = x(a,a)$ and $x^b = x(b,b)$ when both signals are strong, as will be shown shortly.

3.2 Refinements

The present model can be interpreted as a sender-receiver game with two senders – the agents – and one receiver, the DM. As standard refinements such as those of Grossman and Perry (1986), Cho and Kreps (1987) and Banks and Sobel (1987) have been formally defined only for sender-receiver games with one sender and one receiver, some adjustments are necessary. The main idea guiding these extensions to games with multiple senders is to focus on one sender, keeping the other sender's strategy fixed according to the equilibrium under consideration and treating the effect this equilibrium strategy may have on the receiver's beliefs as moves by nature. For the purpose of defining and applying the refinements the game is thus transformed from one with two senders and one receiver of known (or fixed) type to a game with one sender facing one receiver who may be of different types.[23] If the other sender is supposed to play a separating equilibrium strategy the sender who contemplates deviation will "know" the distribution of receiver types because the probability that the other sender has received the same signal, conditional on his own signal being k, and plays the corresponding equilibrium strategy is given by $\pi(k|k)$. As sender-receiver games with multiple receiver-types are non-standard it is also necessary to be explicit about how the various types of the receiver update. The receiver will be assumed to hold the same beliefs about the deviating sender's type for each of her possible types.[24]

[23]The multiplicity of receiver types is only an issue if the equilibrium under investigation prescribes a separating strategy to the other sender: If this other sender's equilibrium strategy is pooling, the sender contemplating a deviation will correctly infer the receiver's beliefs (that is, will "know" them with certainty).

[24]Obviously, keeping the beliefs about the types of the deviating sender fixed implies that the receiver's beliefs about the state of the world will differ depending on her type. As mentioned, it is standard to

Because of the multiple senders, the receiver first needs to determine which sender has deviated after making an observation that is off the equilibrium path. The assumption underlying the procedure to identify a deviator is the hypothesis that the minimum number of deviations necessary to generate a given off the equilibrium path observation have occurred (as first proposed by Bagwell and Ramey, 1991): Suppose the receiver expects to observe both senders to play x and y with positive probability but sender 2 playing z with zero probability in a given equilibrium. Then upon observing (x, z) the receiver will identify sender 2 as the deviator. For any $\varepsilon > 0$, it will always be possible to identify a unilateral deviator in this manner given that the receiver makes an off the equilibrium path observation.[25]

Intuitive Criterion (CK) After the unilaterally deviating sender has been identified as described above, the Intuitive Criterion by Cho and Kreps (1987) (CK) can be adapted rather straightforwardly to the present setup. According to CK, the receiver (DM) must put zero probability on that (if any) type of the deviating sender whose expected equilibrium payoff exceeds the expected payoff from the deviation if the receiver plays a best response for each of her possible types for any beliefs about the deviating sender's type. Equilibria satisfying CK will also be called intuitive.

D1 Cho and Kreps (1987) call an equilibrium D1 if it is robust to deviations under the restriction that upon a deviation the receiver assigns probability zero to types who, relative to their equilibrium payoffs, benefit from strictly fewer mixed strategy best replies by the receiver than some other type.[26] The concept translates directly to the (transformed) game in the present model, in which there is only one sender and one receiver. This transformed game is a standard sender-receiver game if the equilibrium under consideration is pooling. If the equilibrium is separating, the only twist is that the receiver will be one of two types the distribution over which being given by the conditional probabilities $\pi(k|k)$ that the other sender has received the same signal as the sender who contemplates deviation. As is standard for Bayesian games, a (mixed or pure) strategy for the receiver is then a complete type-contingent plan.

assume that all receivers update in the same way in games with multiple receivers. The same requirement is imposed here on types in a game with multiple types of a single receiver.

[25]If an equilibrium is separating and prescribes playing x upon signal a and y upon b, then a sender may deviate to y upon signal a, but this will not be perceived as an off the equilibrium observation by the receiver and thus it requires no specification of updating beyond Bayes' rule.

[26]See also Banks and Sobel (1987) or Sobel (forthcoming).

PSE Grossman and Perry (1986) propose Perfect Sequential Equilibrium (PSE) as a refinement that puts restrictions on the beliefs assigned to types who potentially benefit from a deviation. PSE is very much in the spirit of CK in that it requires the receiver to put zero probability on types for whom an off equilibrium move is dominated by the payoff they get in equilibrium. In addition, PSE requires what Grossman and Perry call credible updating, that is, to assign prior preserving posterior beliefs to all types who could potentially benefit from the observed deviation. Specifically, in the transformed game of the present model the PSE algorithm works as follows.[27]

First, let $\mathbf{x} = \{x_1^a, x_1^b, x_2^a, x_2^b\}$ be the set of locations that the DM expects to observe with positive probability in a given PBE and assume that she observes a deviation by agent 2 to some $x_2 \notin \{x_2^a, x_2^b\}$. Then ask which type(s) - a, b or both - could have benefited from playing x_2 rather than the prescribed equilibrium action. Second, assign prior preserving posterior beliefs to all types who could benefit from the deviation. That is, if both types can benefit from this deviation and if μ_a and μ_b denote the priors over these types, then the posterior beliefs $\hat{\mu}_a$ and $\hat{\mu}_b$ must satisfy $\frac{\hat{\mu}_a}{\hat{\mu}_b} = \frac{\mu_a}{\mu_b}$. Third, given these updated beliefs the DM makes the choice that maximizes her expected utility. Fourth, if given this choice by the DM both types of agent 2 are no better off than when playing the equilibrium location, the deviation has not paid off. If both types benefit, the PBE fails the PSE test.[28] Last, for $\mathbf{x}$ to be a set of PSE locations there must be no deviation according to which at least one type of agent 1 or 2 could benefit, assuming that the DM goes subsequently through steps 1 to 4 of this algorithm.

4 Equilibrium

4.1 Pooling Equilibria

The first result is a recurring theme within this paper and the broader literature on information transmission, namely that babbling is almost always an equilibrium that cannot be refined away.[29] Despite there being only two types, the intuitive criterion

[27]See Grossman and Perry (1986), Riley (2001) and Hörner and Sahuguet (2007). Schultz (1996) employs a similarly strengthened version of the CK criterion. Farrell (1993)'s neologism-proofness is a closely related concept.

[28]If only one type benefits, then go through the same exercise again, but this time by assigning probability one to the type who could have benefited. The requirement is then that if probability is assigned only to one type, and if the DM takes her expected utility maximizing action given this belief, the type who is believed to have chosen this location with probability zero has indeed no incentive to make this deviation.

[29]An exception is obtained in Callander (2008)'s model, in which, for certain parameter regions, no refined pooling equilibrium exists because political candidates also care about policy, which allows them

proposed by Cho and Kreps (1987) has no bite vis-à-vis pooling equilibria.

Proposition 2 *For any $\alpha \in (0,1)$, there is a continuum of pooling equilibria that satisfy CK and D1. For any $\alpha \in (0,1)$ there is a unique pooling PSE, whose outcome is $x(\alpha)$.*

Observe first that if the DM selects the deviator with probability one, both types of an agent potentially benefit from a deviation. Thus, the off equilibrium belief satisfying the constraints imposed by PSE is $\mu = \alpha$. To see that no $x \neq x(\alpha)$ is a pooling PSE outcome, assume without loss of generality that $x < x(\alpha)$. Because $u(\alpha, x)$ has a unique maximum, which is achieved with $x(\alpha)$, the DM will prefer any $z \in (x, x(\alpha)]$ to the proposed equilibrium location. Hence, there are profitable deviations satisfying the restriction on updating imposed by PSE. Thus no $x \neq x(\alpha)$ can be a pooling PSE outcome. To see that $x(\alpha)$ is a pooling PSE outcome, it suffices to notice that the DM will strictly prefer $x(\alpha)$ to any other location as long as her beliefs are α, which as just argued they will be both on and off the equilibrium path. Notice also that, rather obviously, the unique pooling PSE outcome $x(\alpha)$ is the DM (or welfare) optimal pooling outcome.[30]

4.2 Separating Equilibria

Next consider separating equilibria, beginning with the case where both signals are strong.

Strong signals. A corollary to Lemma 3 is that $x_i^a = x^a$ and $x_i^b = x^b$ for $i = 1, 2$ are strategies in a separating intuitive equilibrium only if (5) holds. In addition, CK imposes the following restriction on separating equilibria:

Lemma 5 *If both signals are strong, $x_i^a = x^a$ and $x_i^a = x^b$ for $i = 1, 2$ are strategies of a separating intuitive equilibrium only if*

$$x^a = x(a,a) \quad \text{and} \quad x^b = x(b,b). \tag{6}$$

The lemma is key. It adds a second condition on separating equilibrium locations when both signals are strong that will typically not hold simultaneously with condition (5)

to signal their types.

[30]The expected utility of the DM is the appropriate welfare measure if it does not matter to society which agent wins the contest. Observe also that nothing in this argument hinges on the assumption that the action space X contains more than two elements. Thus, for example in the binary policy model of Heidhues and Lagerlöf (2003) PSE selects the welfare superior popular beliefs equilibrium amongst the two pooling equilibria.

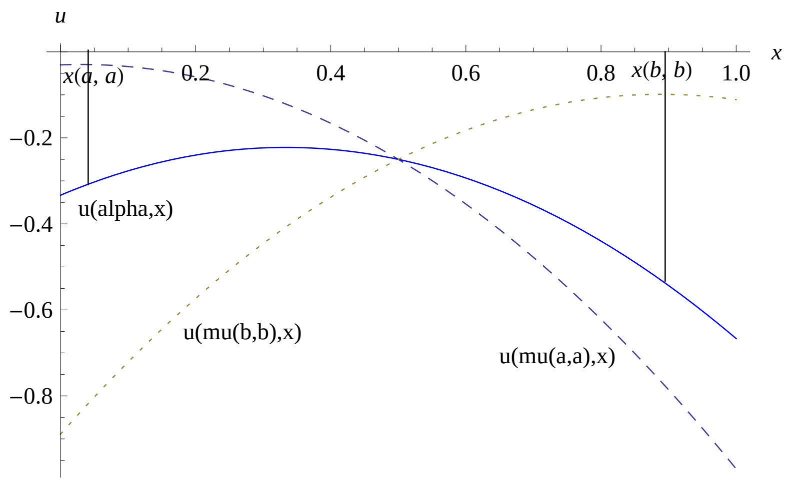

Figure 2: Necessary Conditions for Intuitive Separating Equilibria.

because $u(\alpha, x(a,a)) \neq u(\alpha, x(b,b))$ will be the case for $\alpha \neq 1/2$. Figure 2 depicts a case where $u(\alpha, x(a,a)) > u(\alpha, x(b,b))$. The solid curve is $u(\alpha, x)$. The dotted curve is $u(\mu(b,b), x)$ and the dashed curve is $u(\mu(a,a), x)$. The locations $x(a,a)$ and $x(b,b)$ are indicated by straight vertical lines.[31]

Upon getting signal k an agent's belief that the other agent has received the same signal, $\pi(k|k)$, exceeds 1/2 if signal k is strong. If he plays according to equilibrium, he wins with probability 1/2 (Lemma 2). Therefore, a deviation that guarantees victory if the other agent has received the same signal will be profitable since his expected payoff is at least $\pi(k|k) > 1/2$. If both signals are strong, and if the deviation leads to selection if and only if the other agent has received the same signal k and plays x^k, then upon getting the signal $l \neq k$ an agent will only win with probability $1 - \pi(l|l) < 1/2$. Consequently, this is not a profitable deviation. Therefore, starting from a candidate separating equilibrium, a deviator can credibly convey his signal by choosing such a location, and benefit from the deviation. Unless, that is, the prescribed equilibrium location upon signal k is $x(k,k)$. To see that the deviation leads to selection if and only if the other agent has received the same signal, recall from Lemma 4 that no separating PBE locations $x^a, x^b \notin [x(a,a), x(b,b)]$ exist. Therefore, the locations of a candidate separating intuitive equilibrium will be between $x(a,a)$ and $x(b,b)$ and satisfy $u(\alpha, x^a) = u(\alpha, x^b)$. Suppose $x^a, x^b \in (x(a,a), x(b,b))$. If only, say, type a agents benefit from the deviation to $x(a,a)$, then upon observing $(x(a,a), x^b)$ the DM will prefer x^b since her belief will be α. (And if she believes that both types of the deviating agent

[31]The figure is drawn for the model with quadratic utility under the parameterization $A = 0$, $B = 1$, $\alpha = 2/3$ and $\varepsilon = 1/5$.

have played $x(a,a)$ with positive probability, her belief that the state is A will be even less than α and she will prefer x^b a fortiori.) However, upon observing $(x(a,a), x^a)$ the DM will prefer $x(a,a)$ if she assigns probability 0 to the deviation $x(a,a)$ coming from type b. The probability that the other agent has received the same signal exceeds 1/2 for both signals when both signals are strong. Since on the equilibrium path each agent wins with probability 1/2 independently of the signal (Lemma 2), the deviation to $x(a,a)$ pays off for the agent with signal a and not for the agent with signal b. Moreover, since exactly the same logic applies for the locations upon signal b, the locations in a separating intuitive equilibrium must be $x(a,a)$ upon signal a and $x(b,b)$ upon signal b. Since for any $\alpha \neq 1/2$, $u(\alpha, x(a,a)) \neq u(\alpha, x(b,b))$ holds, the next proposition follows:

Proposition 3 *If both signals are strong and $\alpha \neq 1/2$ holds, there is no separating equilibrium that satisfies CK.*

This proposition has the following corollary:

Corollary 1 *If both signals are strong and $\alpha \neq 1/2$, the unique PSE outcome is $x(\alpha)$ when agents play pure strategies.*

Interestingly, in the present model the CK criterion has bite for separating equilibria but not for pooling equilibria. Separating intuitive equilibria when both signals are strong exist if $\alpha = 1/2$, which implies that $u(\alpha, x(a,a)) = u(\alpha, x(b,b))$. For there to be a separating PSE when separating intuitive equilibria exist, there must also be no $x \in (x(a,a), x(b,b))$ such that $u(\mu(a,0), x) > u(\mu(a,0), x(a,a))$ and $u(\mu(b,0), x) > u(\mu(b,0), x(b,b))$. The reason is that if such an x exists, the deviation to x pays off regardless of the agent's signal and induces the belief $\mu(k,0)$ if the other agent has chosen the equilibrium location prescribed upon signal k with $k \in \{a,b\}$. The conditions under which no such x exists are hard to pin down in general. However, assuming quadratic utility one can readily establish that at $\alpha = 1/2$ there is no such x.

Proposition 4 *Assume that utility is quadratic and $\alpha = 1/2$. Then there is a separating PSE.*

Weak signal. Consider now the case where one signal is weak.

Proposition 5 *If one signal is weak, then there are separating equilibria satisfying CK and there is a unique separating PSE that is generic. The PSE locations are as follows. If b is the weak signal, agent $i = 1,2$ sets $x(a,0)$ if $s_i = a$ and x^b if $s_i = b$, where $u(\alpha, x(a,0)) = u(\alpha, x^b)$. If a is the weak signal, agent $i = 1,2$ sets $x(b,0)$ if $s_i = b$ and x^a if $s_i = a$, where $u(\alpha, x(b,0)) = u(\alpha, x^a)$.*

Thus, despite the fact that agents are purely motivated by winning, the model exhibits a generic separating intuitive equilibria if one signal is weak. The construction of the unique PSE relies first, as any separating PBE, on an indifference condition of the DM in case the agents' signals disagree and second on the fact that upon receiving the weak signal k an agent has no incentive to deviate to $x(k,k)$ if the equilibrium prescribes playing $x_k \neq x(k,k)$. This is so because upon getting the weak signal the agent believes with probability less than 1/2 that the other agent has received the same signal. Therefore, he is not willing to take the gamble of defeating the opponent only in the event he has got the same signal because this event is not sufficiently likely. In contrast, upon either signal both agents would have incentives to take the gamble of defeating the opponent with certainty only in the event that the opponent has received the strong signal. Therefore, the equilibrium prescribes to play $x(\mu(l,0))$ upon $s_i = l$ if l is the strong signal.

It is also worth noting that if one signal is weak there also exists a separating equilibrium that satisfies D1. This equilibrium is qualitatively very similar to the PSE described in Proposition 5. However, because D1 constrains off equilibrium updating in ways that differ somewhat from PSE, the locations in a separating D1 equilibrium are $x(k,k)$ upon the strong signal k and $x \neq x(k,k)$ such that $u(\alpha, x(k,k)) = u(\alpha, x)$ upon the weak signal. The proof is in Appendix B. The fundamental difference is that PSE imposes prior preserving beliefs for any deviation that is interpreted as pooling (that is, as being profitable for both sender types), which imposes the restriction $x_k = x(k,0)$ for the equilibrium location upon the strong signal k. Under D1 deviations $x \in (x^a, x^b)$ are also interpreted as pooling but without imposing restrictions on the receiver's beliefs. However, if signal a is strong and $x^a > x(a,a)$, then deviations below but not too far away from x^a will be interpreted as coming from the sender with signal a only (essentially because $\pi(a|a) > 1 - \pi(b|b)$), which leads to the restriction $x^a = x(a,a)$ for x^a to be an equilibrium location.

Endogenous Signalling Costs Despite the fact that ex ante agents are identical and are, per se, willing to choose any location if it increases the chances of being selected, the model exhibits endogenous signaling costs. This contrasts with standard signaling games such as Spence's education model, where types differ with respect to their exogenously given costs of education, which allows them to separate in an intuitive equilibrium. Separation can occur here because some deviations become too costly in equilibrium, given an agent's probability assessment about the other agent's signal and hence action.

5 Extensions

Cheap Talk Augmented Model So far agents could only "communicate" through their location choices. Consider now the following cheap talk augmented model. Upon observing his signal s_i each agent i chooses a location x_i and sends a message m_i to the DM. Thus, letting $p_i = (x_i, m_i)$ be the platform or brand of agent i, the DM observes a pair (p_1, p_2), updates her beliefs and then selects the agent whose location x_i maximizes her expected utility given her thus updated beliefs. To maintain comparability with the results derived above assume as before that agents play pure strategies and that the DM plays a symmetric strategy. That is, when the platforms proposed by agent 1 and 2 are p' and p'' the DM selects agent 1 with the same probability as she selects agent 2 when the platforms of agent 1 and 2 are p'' and p', respectively.

Given the restriction to symmetric DM strategies, it still holds that in equilibrium, conditional on either signal, each agent is selected with probability 1/2. Consequently, the same indifference requirements on separating equilibrium platforms must hold as have to hold for separating equilibrium locations in the game without cheap talk messages. That is, $u(\alpha, x^a) = u(\alpha, x^b)$ has to hold, where x^k is the location equilibrium prescribes to set upon signal $k \in \{a, b\}$.

Denote by $(\overline{x}^k, \overline{m}^k)$ a prescribed equilibrium platform in the game with messages. A deviation from $(\overline{x}^k, \overline{m}^k)$ can be one of three things: (i) $(\overline{x}^k, m^k)$, (ii) $(x^k, \overline{m}^k)$ or (iii) (x^k, m^k). Obviously, deviations of type (ii) and (iii) make no difference relative to a model without cheap talk messages: Whatever the deviator manages to convey with $(x^k, \overline{m}^k)$ or (x^k, m^k) can be conveyed with a deviation x^k in the baseline model. Therefore, the only channel through which a difference between the two models can arise is via (i).

It is almost immediate that the set of pooling equilibrium locations satisfying a given refinement in the model with cheap talk messages cannot differ from those in the model without: The DM will still be indifferent between the two agents' locations given that they offer the same locations even after a deviation of type (i). Things are only slightly trickier for separating equilibria because now a priori a deviation of type (i) can pay off if it manages to affect the DM's beliefs. Suppose that the equilibrium is separating and that the deviation $(\overline{x}^a, m^a)$ shifts the beliefs towards A if the other agent chooses the equilibrium location prescribed for type b. Then the deviation would pay off since it guarantees a win in case of diverging signals and a tie in case both get signal a. But such belief shifting is not possible because the DM already correctly infers in equilibrium that the deviating agent has received the signal a. In this case she is indifferent between $\overline{x}^a$ and $\overline{x}^b$ if the other agent played $(\overline{x}^b, \overline{m}^b)$. Thus, the set of separating equilibrium locations satisfying a given refinement will not differ between the two models either.

Outside Option As in parts of Krishna and Morgan (2001b) and in Gilligan and Krehbiel (1989) assume now there is an exogenously given outside option, such as a status quo or default location x^O, the DM can threaten to choose. In political economics, this outside option can also be interpreted as the preferred and state independent policy proposed by a non-strategic third agent, who either does not receive a signal or is commonly known to be so ideological that he will ignore any piece of information. In industrial economics, the outside option may be a product offered by a firm that is not responsive to the conditions in the market under consideration. Everything else is assumed to be as in the main model.

It is not hard to show that $x(\alpha)$ is still the unique pooling PSE outcome. Another question is whether there are separating PSE once the DM has an outside option. When both signals are strong and $x^O \notin [x(a,a), x(b,b)]$, then there is no separating equilibrium satisfying CK, provided the prior is not uniform. The reason is simply that the presence of the outside option x^O is irrelevant insofar as for any admissible belief $\mu \in [\mu(a,a), \mu(b,b)]$, $x(a,a)$ or $x(b,b)$ (or both) will be preferred to x^O. Interestingly, if x^O lies between $\underline{x}(x(b,b))$ and $\overline{x}(x(a,a))$, then x^O enables the existence of a separating PSE when both signals are so strong that $\pi(k|k) > 2/3$ for $k \in \{a,b\}$.[32] To see why this is so, notice first that $x^O \in (\underline{x}(x(b,b)), \overline{x}(x(a,a)))$ implies $u(\alpha, x^O) > \max\{u(\alpha, x(a,a)), u(\alpha, x(b,b))\}$. Therefore, if $x(a,a)$ and $x(b,b)$ are the prescribed equilibrium locations upon signals a and b, respectively, the DM will choose the outside option whenever she observes $x(a,a)$ and $x(b,b)$. Keeping fixed the other agent's strategy, there is no deviation for an agent that would benefit both of his possible types. Thus, off the equilibrium path, the DM's beliefs will be $\mu(k,k)$ if both agents have received the same signal and α if they have received different signals. Assuming that $x^O \neq x(\alpha)$, a deviation to the location $x(\alpha)$ would only lead to a win in the event the competing agent has received the opposite signal, which happens with probability $1 - \pi(k|k) < 1/2$. On the equilibrium path, each agent wins with probability $\pi(k|k)/2$ upon signal k, which is greater than $1 - \pi(k|k)$ if and only if $\pi(k|k) > 2/3$.

In a model with binary states, signals and actions, Felgenhauer (forthcoming) analyzes the effect of adding a third agent who is a populist in the sense that he always

[32]Observe that $\underline{x}(x(b,b)) < \overline{x}(x(a,a))$ since by construction $\underline{x}(x(b,b)) < x(\alpha) < \overline{x}(x(a,a))$. Using the formulas in the proof of Lemma 1, one can easily show that $\pi(a|a) > 2/3$ is equivalent to $\alpha > \frac{\varepsilon(2-3\varepsilon)}{1-2\varepsilon}$ and $\pi(b|b) > 2/3$ is equivalent to $\alpha < \frac{1-4\varepsilon+3\varepsilon^2}{1-2\varepsilon}$. Since $\frac{1-4\varepsilon+3\varepsilon^2}{1-2\varepsilon} > \frac{\varepsilon(2-3\varepsilon)}{1-2\varepsilon}$ for all $\varepsilon < (3-\sqrt{3})/6$, there exist values of $\alpha \in \left(\frac{\varepsilon(2-3\varepsilon)}{1-2\varepsilon}, \frac{1-4\varepsilon+3\varepsilon^2}{1-2\varepsilon}\right)$ if and only if $\varepsilon < (3-\sqrt{3})/6$.

chooses the uninformed DM's preferred action. He shows that a fully informative equilibrium exists. Interpreting the outside option x^O as the action chosen by a third agent, the above results complement and extend these findings. They show that a fully informative equilibrium may come into existence even when the outside option is not perfectly populist, that is, even if $x^O \neq x(\alpha)$, provided it satisfies $x^O \in (\underline{x}(x(b,b)), \overline{x}(x(a,a)))$. Therefore, the DM can make an informed decision if she can credibly commit to sticking with the outside option. On the other hand, if $x^O \notin [x(a,a), x(b,b)]$ the threat to go with the outside option is empty, and therefore the outside option does not affect the equilibrium outcomes.

Imperfect Commitment So far, locations have been assumed to be fixed before the DM takes her decision. This is an appropriate assumption whenever design precedes choice. However, it may be debatable in various contexts. For example, in political economics, where the agents are political candidates who compete for office, the winner of the contest may not necessarily implement the policy with which he campaigned. So assume now that agents are political candidates and that each candidate can a priori be one of two types or qualities $q \in \{C, O\}$. These qualities matter in the following sense. A candidate of quality C is *committed* to implement the policy he proposed during the campaign. A candidate of quality O is *opportunistic* and willing to deviate from the policy he announced in the election campaign. Assume that a candidate with quality O implements the policy (or location) that maximizes the DM's expected utility given the information the agent has once he is in office.[33] Notice that candidates of quality C are the agents studied hitherto. A candidate's quality is his private information, and the probability that $q = C$ is $\beta > 0$, so that the probability of being $q = O$ is $1 - \beta$.[34] Quality is independent across candidates and independent of the signals. Regardless of his quality, each agent still cares exclusively about winning the election.[35] All the other assumptions are as in the baseline model.

Lemma 6 *Candidate quality $q \in \{O, C\}$ cannot be signalled in equilibrium.*

[33] The labels "committed" and "opportunistic" are slight misnomers because candidates of quality O are very desirable from a welfare point of view. Notice also that the qualities $q \in \{C, O\}$ are very similar to the candidates' cost of exerting effort in Callander (2008)'s model, with quality O corresponding to a low cost of effort.

[34] Notice that the model analyzed so far corresponds to the special case with $\beta = 1$ of the model sketched here.

[35] After receiving his signal s_i and before choosing his policy x_i, candidate i can therefore now be one of four types in $\{C, O\} \times \{a, b\}$.

Assume now that upon being in office a candidate learns the other candidate's signal regardless of the action chosen. This assumption simplifies the analysis.[36] Upon being in office, both politicians implement the same policy with probability $1-\beta$. Though the DM does not necessarily know what this policy is, the candidates do not differ in that respect if they are of quality O, and the DM need not worry since the policy will maximize her expected utility. So candidates and their proposals only matter if they are of quality C. Therefore, every equilibrium is completely driven by C-types. But this is the model already studied above. Consequently:

Proposition 6 *Equilibrium locations do not vary with $\beta \in (0,1]$.*

The proposition implies that whatever is an equilibrium outcome in the model with $\beta = 1$ is an equilibrium outcome for any $\beta \in (0,1)$ for any given refinement. Thus, whether or not candidates are fully committed to implement the locations they propose does not affect equilibrium behavior.[37] For any $\beta > 0$, C-types imitate O-types in any pure strategy PBE. Therefore, the presence of O-types has no impact on the locations proposed in equilibrium. Still, their presence matters as it obviously improves the DM's expected welfare.

Welfare Another interesting question is whether expected DM welfare is indeed larger in the separating PSE, as one would naturally conjecture. The expected welfare of the DM in the pooling PSE is $W^{pool} = u(\alpha, x(\alpha))$. Absent an outside option such as a status quo policy, DM welfare is the appropriate welfare measure as one of the two agents is chosen with probability one. Her expected welfare in a separating PSE with equilibrium locations x^a and x^b is[38]

$$W^{sep} = (1-\varepsilon)[\alpha u(1,x^a) + (1-\alpha)u(0,x^b)] + \varepsilon[(1-\alpha)u(0,x^a) + \alpha u(1,x^b)]. \quad (7)$$

How W^{pool} and W^{sep} compare is hard to say in general as the comparison depends both on the intricate properties of the utility function v and on the signal technology.

[36] An alternative (whose implications remain to be explored) would be to assume that once in office a politician only knows his own signal for sure and, depending on the equilibrium, may or may not infer the signal of the candidate who lost the election.

[37] The limit case with $\beta = 0$, where both candidates are opportunistic, is excluded from the proposition because it gives rise to a continuum of equilibrium outcomes even under the PSE refinement: Since both candidates implement the same policy with probability 1 upon being elected, their policy proposals do not matter. Therefore, equilibrium is indeterminate.

[38] To see that this is true, notice that $W^{sep} = \alpha(1-\varepsilon)^2 u(1,x^a) + (1-\alpha)(1-\varepsilon)^2 u(0,x^b) + \alpha(1-\varepsilon)\varepsilon[u(1,x^a) + u(1,x^b)] + (1-\alpha)(1-\varepsilon)\varepsilon[u(0,x^a) + u(0,x^b)] + \alpha\varepsilon^2 u(1,x^b) + (1-\alpha)\varepsilon^2 u(0,x^a)$. This simplifies to the expression in (7).

To simplify, assume that utility is quadratic with $A = 0$ and $B = 1$, so that $u(\alpha, x) = -\alpha x^2 - (1-\alpha)(1-x)^2$. Since $x(\alpha) = 1-\alpha$, $W^{pool} = -\alpha(1-\alpha)$. Without loss of generality, assume that signal a is strong and signal b is weak, so that a separating PSE exists for all $\alpha > 1-\varepsilon$. The equilibrium locations are $x^a = x(\mu(a,0))$ and the $x^b \neq x^a$ that solves $u(\alpha, x(\mu(a,0))) = u(\alpha, x^b)$, yielding $x^a = 1-\mu(a,0)$ and $x^b = 1-2\alpha+\mu(a,0)$. So the ex ante expected welfare in the separating PSE is

$$W^{sep} = -(1-\varepsilon)[\alpha(1-\mu(a,0))^2+(1-\alpha)(-\mu(a,0)+2\alpha)^2]-\varepsilon[\alpha(\mu(a,0)-2\alpha+1)^2+(1-\alpha)\mu(a,0)^2]. \tag{8}$$

Algebra reveals that $W^{sep} > W^{pool}$ for all $\alpha > \frac{1-4\varepsilon}{4(1-2\varepsilon)}$, which is strictly less than $1-\varepsilon$ for all $\varepsilon < 1/2$. Thus, whenever a separating PSE exists for $\alpha \neq 1/2$, it generates higher welfare than the pooling PSE.

Multi-Dimensional Space Many real-world problems that can be described as location games are naturally thought of as multi-dimensional location games. For example, a terroristic or other threat to national security may require both extending military measures and curtailing civil liberties. Under the assumption that utility is quadratic, the results from the one-dimensional model translate to an N-dimensional model straightforwardly.[39] To see this, let $X = \mathbb{R}^N$ with $N \geq 2$ be the location space and let $\boldsymbol{\omega}^K = (\omega_1^K, .., \omega_N^K)$ be the N-dimensional state K with $K \in \{A, B\}$. A location is now an N-dimensional vector, denoted $\mathbf{x} = (x_1, .., x_N)$. The DM's utility in state K when the location is $\mathbf{x}$ is $v(\boldsymbol{\omega}^K, \mathbf{x}) = -\sum_{i+1}^N (\omega_i^K - x_i)^2$. Consequently, the expected utility of the DM given belief μ is

$$u(\mu, \mathbf{x}) := -\mu \sum_{i=1}^N (\omega_i^A - x_i)^2 - (1-\mu) \sum_{i=1}^N (\omega_i^B - x_i)^2, \tag{9}$$

whose unique maximizer is $x_i(\mu) = \mu\omega_i^A + (1-\mu)\omega_i^B$ for all $i = 1, .., N$. Let $\mathbf{x}(\mu) = (x_1(\mu), .., x_N(\mu))$. Observe that $\mathbf{x}(\mu)$ is just a convex combination of the two states with weight μ on state A and $1-\mu$ in state B. The shorthand notation $\mathbf{x}(a,a), \mathbf{x}(b,b)$, and $\mathbf{x}(k,0)$ with $k \in \{a,b\}$ for the optimal locations given the updated beliefs $\mu(a,a), \mu(b,b)$ and $\mu(k,0)$, respectively, is also useful.[40]

The previous arguments can be readily used to establish that $\mathbf{x}(\alpha)$ is the unique pooling PSE location and that, if both signals are strong, $\mathbf{x}(a,a)$ and $\mathbf{x}(b,b)$ are the

[39] If utility is not quadratic but additively separable in each dimension (and well-behaved), the optimal location given belief $\mathbf{x}(\mu)$ will in general not be linear in μ. However, $\mathbf{x}(\mu)$ will still be unique for a given belief and continuous and monotone in μ. Therefore, the results can easily be generalized in this way.

[40] That is, $\mathbf{x}(k,k) := \mathbf{x}(\mu(k,k))$ and $\mathbf{x}(k,0) := \mathbf{x}(\mu(k,0))$.

prescribed equilibrium locations upon signals a and b, respectively, in any separating equilibrium satisfying CK. For any ε, $u(\alpha, \mathbf{x}(a,a)) = u(\alpha, \mathbf{x}(b,b))$ will hold for one value of α only, so that for almost all of the parameters space $(0,1) \times (0,1/2)$ there will be no fully separating equilibrium satisfying CK when both signals are strong.

Showing that a separating PSE exists when one signal is weak requires only a little bit more work. Assume that b is the weak signal. Then from the previous analysis we know that the location to be chosen upon signal a in a candidate separating equilibrium is $\mathbf{x}(a,0)$. For any $\mathbf{x} < \mathbf{x}(\alpha)$, let $\overline{\mathbf{x}}(\mathbf{x}) = (2x_1(\alpha) - x_1, .., 2x_N(\alpha) - x_N)$ be location that is symmetric to the other side of the optimal location given the prior that yields the same utility in every component as does x_i, given the prior. This implies $u(\alpha, \mathbf{x}) = u(\alpha, \overline{\mathbf{x}}(\mathbf{x}))$. Assume then that upon signal b, each agent chooses the location $\overline{\mathbf{x}}(\mathbf{x}(a,0))$. Observe that these locations are such that there is no profitable deviation for an agent of type a because any profitable deviation for type a would also be profitable for an agent of type b. That is, an agent of type b would have an incentive to choose a location that defeats a type a agent. Since $\mathbf{x}(a,0)$ is the unique utility maximizing location given the belief $\mu(a,0)$, no such deviation exists. Similarly, because $\pi(b|b) < 1/2$ when signal b is weak, an agent of type b has no incentive to deviate to a location that guarantees a win only against a type b agent. For reasons that are very similar to those in the cheap talk augmented model, the enlargement of the action space for the agents does not affect the nature of equilibria as a function of signal strengths.

Sequential Moves Assuming simultaneously moving agents may be preferable to assuming sequential moves because it presumes less detailed knowledge on behalf of the modeler. However, it is also worth investigating what happens if agents move sequentially. This alternative is considered, for example, by Krishna and Morgan (2001a,b) and Pesendorfer and Wolinsky (2003). Throughout this subsection assume agent 1 moves first and agent 2 second.

In general the assumption that senders move sequentially will require a refinement to specify how the second sender updates his beliefs. However, in the present model any sender's payoff is not directly affected by his own beliefs (or those of the other sender) but only depends on the action the receiver takes, given her beliefs.[41] In any equilibrium the second sender will correctly anticipate the receiver's beliefs both on the equilibrium path and after a deviation by sender 1. Therefore, how agent 2 updates his beliefs can be

[41] Notice the difference to the game where agents move simultaneously. There each agent's beliefs matter because, depending on the equilibrium, these beliefs can be informative (in a stochastic sense) about the action the other agent takes and therefore about the beliefs of the DM.

abstracted away from, and one only needs to be concerned about how the DM updates her beliefs. Thus, the definitions of CK, D1 and PSE, appropriately adjusted to games with multiple senders, can be directly applied to the game with sequentially moving senders.

Call an equilibrium pooling if and only if both agents choose the same location independently of their signals.

Proposition 7 *Any $x \in [x(a,a), x(b,b)]$ is the outcome of a pooling equilibrium satisfying CK and D1 with sequential moves, and no other location is the outcome of a pooling equilibrium. The unique pooling PSE outcome with sequential moves is $x(\alpha)$.*

Another question is whether there are fully separating PSE with sequential moves, that is, PSE where both players' signals are revealed on the equilibrium path. The answer is affirmative:

Proposition 8 *With sequential moves, there are fully separating PSE for any $\alpha \in (0,1)$ and $\varepsilon \in (0,1/2)$.*

In contrast to the model with simultaneous moves any deviation by any agent would be interpreted as pooling with sequential moves under the PSE refinement. Deviations from the actions prescribed by the separating equilibrium by agent 2 can be preempted by 1 by playing $x(k,0)$ upon signal $k \in \{a,b\}$, which cannot be defeated by a deviant whose behavior is interpreted as pooling. Deviations by agent 1 can be countered by agent 2 by playing $x(\alpha)$ upon agent 1's deviation, which is the DM's optimal location given the hypothesis that both agents are pooling.

A consequence of the fact that with sequential moves all deviations are pooling is that a separating equilibrium location does not have to be $x(k,k)$ if signals are strong. The reason is that the endogenous signaling costs disappear for agent 2 who, assuming equilibrium play by 1, is now certain about 1's signal. Therefore, he cannot credibly reveal his signal by off equilibrium behavior.

The difference to Krishna and Morgan (2001a) is interesting and quite striking. They find that there are fully separating PBE with simultaneous moves but not with sequential moves. For the case where both signals are strong, the opposite obtains in the present paper. The intuition for these differences is that in the setup of Krishna and Morgan (2001a) both agents know the state. Therefore, they face no uncertainty about each other's types. With simultaneous moves truth telling can therefore be induced as equilibrium strategies by credible threats in case the agents disagree. Such off the equilibrium path threats are not available in the present model since the agents will necessarily disagree some of the time because of their imperfect information about the state and each other's types. With sequential moves, any deviation from prescribed equilibrium locations by either agent would be interpreted as pooling in the present setup and can therefore be preempted through appropriately determined equilibrium locations. In contrast, small deviations from a fully separating strategy pay off for the first mover in Krishna and Morgan (2001a)'s model because they will induce the second mover to deviate as well.

Mixed Strategies Equilibria Even though there are no separating intuitive equilibria when both signals are strong and when the prior is not uniform, intuitive equilibria that are more informative than the pooling equilibria exist even in this case if agents are allowed to play mixed strategies. To see this, reconsider the basic model when both signals are strong and assume without loss of generality that $u(\alpha, x(a,a)) > u(\alpha, x(b,b))$.[42] Assume now that both agents play the following strategies: Upon signal a, they choose x^a with probability 1. Upon signal b they play x^b with probability σ and x^a with probability $1-\sigma$. This has various effects. First, upon observing (x^a, x^a) the DM's updated belief that state A has occurred, denoted $\mu_\sigma(x^a, x^a)$, will be less than $\mu(a,a)$ for any $\sigma < 1$. Second, upon receiving signal b an agent's belief that the other agent chooses x^b is now only $\pi_\sigma(b|b) := \sigma\pi(b|b)$. So if $\sigma \leq \frac{1}{2\pi(b|b)}$, signal b becomes effectively weak because of the mixed strategy. Thus, the intuitive criterion (and PSE) imposes no particular constraint on the location x^b to be played upon signal b. Third, upon signal a an agent's belief that the other agent plays x^a is even larger for any $\sigma < 1$ because the other agent now sometimes plays x^a even after receiving signal b. Let a^σ stand for the hybrid signal that is generated if a^σ is signaled with probability 1 when the true signal indicates a and with probability $1-\sigma$ if the true signal indicates b. The constraint on the equilibrium location upon signal a is then that it be $x^{a,\sigma} := x(\mu_\sigma(a^\sigma, 0))$, where $\mu_\sigma(a^\sigma, 0) := \frac{\alpha(1-\varepsilon)}{\alpha(1-\varepsilon)+(1-\alpha)(\varepsilon+(1-\varepsilon)(1-\sigma))}$ is the DM's belief that state A has occurred if one agent plays x^a according to the equilibrium strategy while nothing can be inferred about the other agent's signal.[43] Consequently, for any $\sigma \leq \frac{1}{2\pi(b|b)}$ one can construct such equi-

[42]This is without loss of generality because for $\alpha \neq 1/2$, $u(x(a,a),\alpha) \neq u(x(b,b),\alpha)$ will hold.

[43]Notice that $x^{a,\sigma}$ is the unique equilibrium location prescribed by PSE for any $\sigma < \frac{1}{2\pi(b|b)}$ because both types of the deviator would potentially benefit from the deviation because signal b is effectively weak under this condition. However, when $\sigma = \frac{1}{2\pi(b|b)}$ a type b has no strict incentive to deviate to a location that guarantees selection with probability 1 if the other agent has received the signal a and with probability 0 otherwise. So what is assumed here is that the receiver puts prior preserving beliefs on all those types who are not made worse off by a deviation. This is done for simplicity to avoid dealing

libria by letting $x^a = x^{a,\sigma}$ and $x^b \neq x^a$ such that $u(\mu_\sigma(x^a, x^b), x^{a,\sigma}) = u(\mu_\sigma(x^a, x^b), x^b)$. Clearly, the most informative of these equilibria is the one with $\sigma^* := \frac{1}{2\pi(b|b)}$. As ε approaches $1 - \alpha$ the equilibrium locations and beliefs in this mixed strategy equilibrium converge to those in the separating PSE when signal b is weak. Consequently, welfare in this mixed strategy equilibrium will be the same as in the separating PSE. Thus, equilibrium welfare will be everywhere continuous in ε, assuming the welfare optimal equilibrium is selected. With quadratic utility, welfare in the mixed strategy equilibrium is never less than in the pooling PSE when both signals are strong.[44] Observe also that σ^* goes to $1/2$ as ε approaches 0. Thus, the probability that state B will be revealed on the equilibrium path approaches $\sigma^{*2} = 1/4$ as ε approaches 0. Moreover, the DM's updated belief that the state is A given that both agents propose x^a will be strictly less than 1 even as ε goes to 0.[45]

6 Conclusions

The present paper analyzes a location choice model in which the agents' locations not only directly affect the decision maker's utility but may also convey information to the decision maker about her bliss point location. The paper shows that there is no equilibrium that satisfies the intuitive criterion and that permits the decision maker to make a fully informed decision if the agents' information is sufficiently precise. However, if the agents' information is sufficiently noisy – that is, if the signals' error probability is larger than the prior on one state – then a fully informative equilibrium exists that satisfies the intuitive criterion and the stronger PSE and D1 refinements proposed by Grossman and Perry (1986) and Banks and Sobel (1987).

Two avenues for further research seem particularly fruitful. Interpreting the model as an industrial organizations model, one could incorporate price competition between the agents after they have observed each other's locations, but not necessarily each other's signals. Particularly with the political economics interpretation in mind, one could also let agents' payoffs depend also on the decision maker's utility, and add a post election stage, where either only the winner or both the winner and the loser take some action.

with $\sigma = \frac{1}{2\pi(b|b)}$ as a separate case.

[44] Whether partially informative equilibria can be constructed that satisfy the intuitive criterion without making one of the two signals effectively weak depends on the intricacies of the utility function and the parameters. For example, with quadratic utility such equilibria do not exist for $\alpha = 2/3$ and $\varepsilon = 1/4$ but they do exist when $\varepsilon = 1/10$ (and $\alpha = 2/3$).

[45] To see this, substitute $\varepsilon = 0$ into (22) in Appendix A to get $\mu_\sigma(x^a, x^a) = \alpha/(\alpha + (1-\alpha)/4)$ at $\sigma = 1/2$.

Appendix

A Proofs

Proof of Lemma 1: Upon observing $s_i = a$ agent i's belief that A is true is $\pi_i(\omega = A \mid s_i = a) = \frac{\alpha(1-\varepsilon)}{\alpha(1-\varepsilon)+(1-\alpha)\varepsilon}$ and, consequently, i's belief that B is true given $s_i = a$ is $\pi_i(\omega = B \mid s_i = a) = \frac{(1-\alpha)\varepsilon}{\alpha(1-\varepsilon)+(1-\alpha)\varepsilon}$. So conditional on signal $s_i = a$ i's belief that j's signal is $s_j = a$ is $\pi_i(s_j = a \mid s_i = a) = \pi_i(\omega = A \mid s_i = a)(1-\varepsilon) + \pi_i(\omega = B \mid s_i = a)\varepsilon = \frac{\alpha(1-\varepsilon)^2+(1-\alpha)\varepsilon^2}{\alpha(1-\varepsilon)+(1-\alpha)\varepsilon}$ and conditional on signal $s_i = b$ i's belief that $s_j = b$ is $\pi_i(s_j = b \mid s_i = b) = \frac{\alpha\varepsilon^2+(1-\alpha)(1-\varepsilon)^2}{\alpha\varepsilon+(1-\alpha)(1-\varepsilon)}$. To see that $\pi_i(s_j = a \mid s_i = a) > \frac{1}{2}$ if and only if $\alpha > \varepsilon$ and $\pi_i(s_j = a \mid s_i = a) > \frac{1}{2}$ if and only if $\alpha < 1-\varepsilon$, notice that $\pi_i(s_j = a \mid s_i = a) = 1/2$ at $\alpha = \varepsilon$ and $\pi_i(s_j = a \mid s_i = a)$ is increasing in α for all $\varepsilon < 1/2$. Similarly, $\pi_i(s_j = b \mid s_i = b)$ decreases in α and equals $1/2$ at $\alpha = 1-\varepsilon$. ■

Proof of Lemma 2: It is first shown that $U_1[x_1^k \mid s_1 = k] + U_2[x_2^k \mid s_2 = k] = 1$ for both $k \in \{a, b\}$. Once this is established, the lemma follows rather straightforwardly because of the assumption that the DM's strategy is symmetric.

Denote by $\gamma(x_1, x_2)$ the probability that the DM select agent 1 if 1 plays x_1 and 2 plays x_2. So

$$\begin{aligned}
U_1[x_1^a \mid s_1 = a] &= \pi_1(s_2 = a \mid s_1 = a)\gamma(x_1^a, x_2^a) + (1 - \pi_1(s_2 = a \mid s_1 = a))\gamma(x_1^a, x_2^b) \\
U_2[x_2^a \mid s_2 = a] &= \pi_2(s_1 = a \mid s_2 = a)(1 - \gamma(x_1^a, x_2^a)) + (1 - \pi_2(s_1 = a \mid s_2 = a))(1 - \gamma(x_1^b, x_2^a)) \\
U_1[x_1^b \mid s_1 = b] &= \pi_1(s_2 = b \mid s_1 = b)\gamma(x_1^b, x_2^b) + (1 - \pi_1(s_2 = b \mid s_1 = b))\gamma(x_1^b, x_2^a) \\
U_2[x_2^b \mid s_2 = b] &= \pi_2(s_1 = b \mid s_2 = b)(1 - \gamma(x_1^b, x_2^b)) \\
&+ (1 - \pi_2(s_1 = b \mid s_2 = b))(1 - \gamma(x_1^a, x_2^b)).
\end{aligned}$$

Notice that $\pi_1(s_2 = k \mid s_1 = k) = \pi_2(s_1 = k \mid s_2 = k)$ for $k \in \{a, b\}$. To simplify notation, let $\theta_k \equiv \pi_1(s_2 = k \mid s_1 = k)$, $x \equiv \gamma(x_1^a, x_2^a)$, $y \equiv \gamma(x_1^a, x_2^b)$, $c \equiv 1 - \gamma(x_1^b, x_2^a)$ and $d \equiv \gamma(x_1^b, x_2^b)$.

The incentive constraints (3) can now be written as

$$\begin{aligned}
U_1[x_1^a \mid s_1 = a] = \theta_a x + (1-\theta_a)y &\geq \theta_a(1-c) + (1-\theta_a)d \\
&= U_1[x_1^b \mid s_1 = a] \qquad (10) \\
U_2[x_2^a \mid s_2 = a] = \theta_a(1-x) + (1-\theta_a)c &\geq \theta_a(1-y) + (1-\theta_a)(1-d) \\
&= U_2[x_2^b \mid s_2 = a] \qquad (11) \\
U_1[x_1^b \mid s_1 = b] = \theta_b d + (1-\theta_b)(1-c) &\geq \theta_b y + (1-\theta_b)x \\
&= U_1[x_1^a \mid s_1 = b] \qquad (12) \\
U_2[x_2^b \mid s_2 = b] = \theta_b(1-d) + (1-\theta_b)(1-y) &\geq \theta_b c + (1-\theta_b)(1-x) \\
&= U_2[x_2^a \mid s_2 = b]. \qquad (13)
\end{aligned}$$

Adding (10) and (11) yields $1 \leq y + c$ while adding (12) and (13) implies $1 \geq y + c$. Thus, $y + c = 1$ holds. Now,

$$U_1[x_1^a \mid s_1 = a] + U_2[x_2^a \mid s_2 = a] = \theta_a + (1-\theta_a)(y+c) = 1 \qquad (14)$$

and

$$U_1[x_1^b \mid s_1 = b] + U_2[x_2^b \mid s_2 = b] = \theta_b + (1-\theta_b)(2-(y+c)) = 1, \qquad (15)$$

where the second equalities hold because $y + c = 1$. Thus, the first part of the proof is complete.

To see that the first equality in (4) holds, suppose to the contrary that it does not. Without loss of generality, assume $U_1[x_1^a \mid s_1 = a] < U_2[x_2^a \mid s_2 = a]$. Since $U_1[x_1^a \mid s_1 = a] + U_2[x_2^a \mid s_2 = a] = 1$, this implies $U_1[x_1^a \mid s_1 = a] < 1/2$. But now, upon $s_1 = a$ agent 1 could play x_2^a instead of the prescription x_1^a, in which case he would get

$$U_1[x_2^a \mid s_1 = a] = \theta_a \gamma(x_2^a, x_2^a) + (1-\theta_a)\gamma(x_2^a, x_2^b), \qquad (16)$$

or x_2^b, in which case he would get

$$U_1[x_2^b \mid s_1 = a] = \theta_a \gamma(x_2^b, x_2^a) + (1-\theta_a)\gamma(x_2^b, x_2^b). \qquad (17)$$

Due to the assumption that the DM's strategy must not depend on the agents' labels, $\gamma(x_2^a, x_2^a) = \gamma(x_2^b, x_2^b) = \frac{1}{2}$ and $\gamma(x_2^b, x_2^a) = 1 - \gamma(x_2^a, x_2^b)$. Thus, these two equations simplify to

$$U_1[x_2^a \mid s_1 = a] = \theta_a \frac{1}{2} + (1-\theta_a)\gamma(x_2^a, x_2^b) = \gamma(x_2^a, x_2^b) + \theta_a \left(\frac{1}{2} - \gamma(x_2^a, x_2^b)\right) \qquad (18)$$

and

$$U_1[x_2^b \mid s_1 = a] = \theta_a(1 - \gamma(x_2^a, x_2^b)) + (1 - \theta_a)\frac{1}{2} = \frac{1}{2} + \theta_a \left(\frac{1}{2} - \gamma(x_2^a, x_2^b)\right). \quad (19)$$

If $\gamma(x_2^a, x_2^b) \geq \frac{1}{2}$ the expression in (18) weakly exceeds 1/2 and the expression in (19) is (strictly) larger than 1/2 otherwise. Since agent 1 can either play x_2^a or x_2^b, it has to be the case that his expected equilibrium payoff weakly exceeds 1/2. And since exactly the same argument applies for agent 2, it follows that indeed $U_i[x_i^k \mid s_i = k] = U_j[x_j^k \mid s_j = k] = \frac{1}{2}$ for $k \in \{a, b\}$ as claimed. ■

Proof of Lemma 3: Assume (5) does not hold, for example because $u(\alpha, x^a) < u(\alpha, x^b)$, yet x^a and x^b are set in a separating PBE. But then the deviation to play x^b when the signal is a pays off when the other agent plays x^a by increasing the probability of winning from 1/2 to 1 and when the other other agent plays x^b by increasing the probability of winning from 0 to 1/2.■

Proof of Proposition 1: If both agents play x independently of their signals, the DM is indifferent between the two agents and randomizes. On the equilibrium path, no information is transmitted and the DM's posterior equals her prior α. If, say, agent 1 deviates to some $x_1 \neq x$, then the DM must choose 1 with a probability smaller than 1/2. For this to be sequentially rational, her off equilibrium belief $\mu(x_1, x)$ must be such that her expected utility of selecting agent 2 exceeds her utility of selecting the deviating agent 1. Though PBE does not restrict the off equilibrium beliefs of the DM about the strategy played by the deviating agent, it still imposes bounds on her beliefs $\mu(x_1, x)$: Given the observation (x_1, x) the hypothesis that implies the largest probability on A is that the DM assumes agent 1 plays the strategy "x_1 if $s_1 = a$ and x otherwise" while the hypothesis that minimizes the belief that A is true is that 1 plays "x_1 if $s_1 = b$ and x otherwise". These hypotheses imply, respectively, the updated beliefs $\mu(x_1, x) = \frac{\alpha(1-\varepsilon)}{\alpha(1-\varepsilon)+(1-\alpha)\varepsilon} = \mu(a, 0) < 1$ and $\mu(x_1, x) = \frac{\alpha\varepsilon}{\alpha\varepsilon+(1-\alpha)(1-\varepsilon)} = \mu(b, 0) > 0$. The DM's preferred locations, given these beliefs, are $\underline{x}(\alpha) \equiv x(\mu(a, 0)) < \overline{x}(\alpha) \equiv x(\mu(b, 0))$. Hence, for any "prescribed" equilibrium location $x \in [\underline{x}(\alpha), \overline{x}(\alpha)]$ there are beliefs that make it rational not to choose the deviating agent: Just choose the off equilibrium beliefs $\mu(x_1, x)$ so that $x = x(\mu(x_1, x))$.

As for the separating PBE, notice first that one agent choosing x^a and the other one x^b is an on the equilibrium path observation. Using Bayes' rule, the DM updates her beliefs to $\mu(x^a, x^b) = \frac{\alpha(1-\varepsilon)\varepsilon}{\alpha(1-\varepsilon)\varepsilon+(1-\alpha)(1-\varepsilon)\varepsilon} = \alpha$. Hence, the DM will be indifferent between the two. If both agents choose the same location, she will also be indifferent between the two. In either case, she randomizes uniformly. If one agent deviates and chooses an off equilibrium location x, she must not select the deviating agent with probability larger than one half. Upon observing (x_1, x^a) where x_1 is the off equilibrium observation generated by agent 1 and x^a is the on equilibrium location agent 2 plays upon receiving $s_2 = a$ the DM's hypothesis that puts most probability on state A is that 1 plays the strategy "x_1 if and only if $s_1 = a$". Consequently, $\max \mu(x_1, x^a) = \frac{\alpha(1-\varepsilon)^2}{\alpha(1-\varepsilon)^2+(1-\alpha)\varepsilon^2} = \mu(a,a)$. The hypothesis that puts the least probability onto state A is "x_1 if and only if $s_1 = b$", yielding $\min \mu(x_1, x^a) = \frac{\alpha\varepsilon(1-\varepsilon)}{\alpha\varepsilon(1-\varepsilon)+(1-\alpha)(1-\varepsilon)\varepsilon} = \alpha$. Similarly, upon observing (x_1, x^b), the hypothesis that puts most probability onto state A is that 1 plays "x_1 if and only if $s_1 = a$", yielding $\max \mu(x_1, x^b) = \frac{\alpha\varepsilon(1-\varepsilon)}{\alpha\varepsilon(1-\varepsilon)+(1-\alpha)(1-\varepsilon)\varepsilon} = \alpha$, and hypothesis with the smallest probability on state A is "x_1 if and only if $s_1 = b$", yielding $\min \mu(x_1, x^b) = \frac{\alpha\varepsilon^2}{\alpha\varepsilon^2+(1-\alpha)(1-\varepsilon)^2} = \mu(b, b)$.

Assume $x^a, x^b \in [x(a,a), x(b,b)]$, where $u(\alpha, x^a) = u(\alpha, x^b)$. For any off equilibrium x, that is, for any $x \notin \{x^a, x^b\}$ there are beliefs that make it rational not to select the deviating agent. Without loss of generality assume that the agent that plays on equilibrium plays x^b. If $x > x^b$, the DM can choose the belief α, so that $u(\alpha, x^b) > u(\alpha, x)$. For $x < x^b$, she can choose the belief $\mu(b,b)$ so that $u(\mu(b,b), x^b) > u(\mu(b,b), x)$. ■

Proof of Lemma 4: Suppose to the contrary that there is a separating PBE where, say, $x^a < x(a,a)$, so that $x(a,a)$ is an off equilibrium observation. Now upon observing one agent playing x^a and the other one the off equilibrium $x(a,a)$, the DM's belief that is worst for the agent playing off equilibrium is $\mu(a,a)$ as may be recalled from the proof of Proposition 1. But with this belief the DM prefers $x(a,a)$ to x^a and so she will prefer the off equilibrium location $x(a,a)$ to x^a for any belief $\mu \leq \mu(a,a)$ that is more favorable for the deviating agent. An analogous argument applies for x^b and $x(b,b)$. Finally, the deviation to $x(k,k)$ upon signal k with $k \in \{a,b\}$ pays off for an agent when both signals are strong: Since on the equilibrium path an agent wins with probability 1/2 independently of his signal, the probability of winning upon receiving signal k and playing $x(k,k)$ exceeds 1/2. Thus, the deviation is profitable. ■

Proof of Lemma 5: For $\alpha \in (\varepsilon, 1-\varepsilon)$ both signals are strong. On the equilibrium path, i wins with probability 1/2 independently of his signal. Therefore, if there is a deviation that allows i to win with certainty against x_k and to lose with certainty against x_l with

$k \neq l$, i wants to play this deviation upon signal $s_i = k$ and not upon signal $s_i = l$. From Lemma 4 it is known that $x(a,a) \leq x^a$ and $x^b \leq x(b,b)$. It is now argued that for $x(a,a) < x^a$ and $x^b < x(b,b)$ such a deviation exists, this deviation being $x(k,k)$ upon signal k.

If only an agent with signal a (type a for short) benefits from playing $x(a,a)$, the DM's belief upon observing $(x^a, x(a,a))$ is $\mu(a,a)$, in which case she strictly prefers $x(a,a)$ to x^a by construction of $x(a,a)$. Hence the deviation pays off for type a if only type a benefits from it. To see that the latter is indeed true, notice that if both types benefit from playing $x(a,a)$ the DM's belief upon observing $(x^b, x(a,a))$ is $\mu(b,0)$ because the deviating agent's behavior is not informative. But recall now from Lemma 3 that $u(\alpha, x^a) = u(\alpha, x^b)$. Therefore, upon $(x^b, x(a,a))$ and having belief $\mu(b,0) < \alpha$, the DM strictly prefers x^b to $x(a,a)$ since $x(a,a) < x^a$. Therefore, type b's payoff from the deviation $x(a,a)$ is $(1 - \pi(b|b))$, which is strictly less than his equilibrium payoff of $1/2$ since b is a strong signal. Thus, it is not possible that both types benefit from the deviation $x(a,a)$. (And indeed, if only type a benefits, the DM's belief upon observing $(x^b, x(a,a))$ is $\mu(b,a) = \alpha$, in which case she strictly prefers x^b.) Completely analogous reasoning applies for $x^b < x(b,b)$, Therefore, for $x^a > x(a,a)$ ($x^b < x(b,b)$) playing $x(a,a)$ upon signal A ($x(b,b)$ upon signal b) is a deviation that pays off. ■

Proof of Proposition 3: Consider a candidate separating equilibrium. Then upon a strong signal k it must be the case that agents play $x(k,k)$. Because for $\alpha \in (\varepsilon, 1-\varepsilon)$ both signals are strong, separation in a PSE requires conditions (5) and (6) to hold, which is excluded by assumption for $\alpha \neq 1/2$. So for $\alpha \in (\varepsilon, 1-\varepsilon) \setminus \{1/2\}$ there is no PSE where both agents separate. ■

Proof of Corollary 1: A pure strategy for an agent can be either pooling, that is he sets the same location independently of his signal, or separating, that is he sets different locations as a function of his signal. If the DM updates as required by Grossman and Perry (1986), the (or a) best response against an agent who pools is $x(\alpha)$, and $x(\alpha)$ is the unique best response to itself, as shown in Proposition 2.

Since every PSE is an intuitive equilibrium, Proposition 3 implies that there are no separating PSE when both signals are strong and $\alpha \neq 1/2$. Moreover, since $x(a,0)$ and $x(b,0)$ are unique, there are no hybrid PSE where one agent separates by playing $x(\mu(k,0))$ upon signal k and the other agent pools by playing some x independently of his signal since the DM will always strictly prefer $x(\mu(k,0))$ to $x \notin \{x(a,0), x(b,0)\}$. It can also not be the case that one agent pools at, say, $x = x(a,a)$ and the other one separates because the separating agent would have an incentive to play $x(\alpha)$, whereby he would win for sure. Thus, for $\alpha \in (\varepsilon, 1-\varepsilon)$ the unique PSE outcome is $x(\alpha)$, provided $\alpha \neq 1/2$. ■

Proof of Proposition 4: Given the arguments made in the text, all that is left to be shown is that there is no $x \in (x(a,a), x(b,b))$ such that $u(\mu(a,0),x) > u(\mu(a,0),x(a,a))$ and $u(\mu(b,0),x) > u(\mu(b,0),x(b,b))$. The maximizer of the quadratic utility function $u(\mu,x) = -\mu(A-x)^2 - (1-\mu)(B-x)^2$ is $x(\mu) = B-(B-A)\mu$. Therefore, the point of intersection of $u(\mu(a,0),x)$ with $u(\mu(a,0),x(a,a))$ that is greater than $x(a,a)$ occurs at

$$x = \frac{A - 4A\varepsilon + 2B\varepsilon + 7A\varepsilon^2 - 5B\varepsilon^2 - 4A\varepsilon^3 + 4B\varepsilon^3}{1 - 2\varepsilon + 2\varepsilon^2}. \tag{20}$$

Similarly, the point of intersection of $u(\mu(b,0),x)$ with $u(\mu(b,0),x(b,b))$ that is smaller than $x(b,b)$ occurs at

$$x = \frac{B + 2A\varepsilon - 4B\varepsilon - 5A\varepsilon^2 + 7B\varepsilon^2 + 4A\varepsilon^3 - 4B\varepsilon^3}{1 - 2\varepsilon + 2\varepsilon^2}. \tag{21}$$

For $B > A$, it is readily established that the right hand side of (20) is less than the right hand side of (21) for all $\varepsilon \in (0,1/2)$. Because at $\alpha = 1/2$, $u(\mu(a,0),x(a,a)) = u(\mu(b,0),x(b,b))$ due to symmetry, the result follows. ■

Proof of Proposition 5: Notice first that because $\varepsilon < 1/2$ either $\alpha > \varepsilon$ or $\alpha < 1-\varepsilon$ holds. Therefore, there will be exactly one strong signal, a in the former, b in the latter case. Recall that upon receiving a strong (weak) signal an agent has a posterior exceeding (less than) 1/2 that the other agent has received the same signal. For the sake of the argument, suppose a is the strong signal, that is $\alpha > \varepsilon$. A necessary condition for (x^a, x^b) to be part of a separating equilibrium is that they satisfy $u(x^a,\alpha) = u(x^b,\alpha)$; see Lemma 3. It is now shown that (x^a,x^b) with $u(x^a,\alpha) = u(x^b,\alpha)$ can be part of a separating PSE if and only if $x^a = x(a,0)$ holds. To see necessity, notice that agent i can potentially benefit from a deviation both after $s_i = a$ and $s_i = b$ if upon the deviation he is selected with probability one if the other agent plays x^a. So upon seeing (x^a, x'), where x' is a deviation by i, the DM's belief, updated according to PSE, is $\mu(a,0)$. So unless $x^a = x(a,0)$, the DM prefers $x' = x(a,0)$ to x^a. So on top of $u(x^a,\alpha) = u(x^b,\alpha)$, PSE requires $x_k = x(\mu(k,0))$, where $k \in \{a,b\}$ is the strong signal. Notice that there is now only one constraint, namely the one imposed by the strong signal, whereas in Lemma 5 there were two constraints that have to hold simultaneously on top of $u(x^a,\alpha) = u(x^b,\alpha)$. That such (x^a,x^b) exist is guaranteed since $x(a,0) > x(1)$ and $x(b,0) < x(0)$.

To show sufficiency, maintain the assumption that signal a is strong. The last thing to show is that upon $s_i = b$, i has no incentive to deviate from x^b, provided $x^a = x(a, 0)$ and $u(\alpha, x(a, 0)) = u(\alpha, x^b)$. By on equilibrium play, i wins with probability 1/2 upon either signal. By construction of $x(a, 0)$ there is no deviation that both types can beneficially play. Therefore, the only deviation that potentially benefits an agent with signal b is one that the DM prefers if the other agent plays x^b. However, $\pi_i(s_j = b \mid s_i = b) < 1/2$ because b is the weak signal, so that the expected payoff of the deviation is less than 1/2. Hence, there is no deviation that benefits only agent b. Thus, no profitable deviation exists.

Since every PSE is intuitive, it follows that there are intuitive separating equilibria when one signal is weak. ■

Sketch of Proof of Lemma 6: Suppose candidate i of quality O can successfully signal his quality (and benefit from doing so). Then if his quality were C it would benefit from behaving as if it were of quality $q = O$, and vice versa. ■

Proof of Proposition 7: To see that $x \notin [x_a^*, x_b^*]$ is not the outcome of a pooling equilibrium, it suffices to notice that agent 2 could play $x(\mu(a, a))$ (or $x(\mu(b, b))$) after agent 1 played the prescribed equilibrium location $x < x(\mu(a, a))$ (or $x > x(\mu(b, b))$). Since there are no beliefs that make it sequentially rational to prefer $x < x(\mu(a, a))$ over $x(\mu(a, a))$ (or to prefer $x > x(\mu(b, b))$ over $x(\mu(b, b))$) it follows that such a x is not an equilibrium location.

To see that any $x \in [x_a^*, x_b^*]$ is the outcome of a pooling equilibrium satisfying CK, recall first that on equilibrium each agent is selected with probability 1/2. Thus, no deviation is equilibrium payoff dominated. Thus, CK does not pin down off equilibrium beliefs. Second, assume that the equilibrium strategies are such that agent 2 plays x no matter what agent 1 played. Since there are beliefs that make it sequentially rational to prefer $x \in [x_a^*, x_b^*]$ to any other location, it follows that x is a pooling equilibrium outcome that satisfies CK. Since no type of any agent has a differential gain from a deviation from $x \in [x_a^*, x_b^*]$, it also follows that any such x is the outcome of a pooling equilibrium satisfying D1.

Last, consider the PSE equilibrium outcome $x(\alpha)$. Given that agent 1 pools at $x(\alpha)$ any deviation by the other agent from the prescribed location $x(\alpha)$ would be pooling and would hence be defeated. Agent 1 on the other hand has no incentive to deviate either because agent 2 can simply play $x(\alpha)$ and guarantee that he wins, given that 1's deviation would be pooling. Uniqueness follows along the previous lines: Any deviation by 2 would be pooling. So 2 could profitably deviate if the prescribed location were not $x(\alpha)$. ■

Proof of Proposition 8: Let 1 play x_k, $x_k = x(\mu(k,0))$ upon signal $s_1 = k$ with $k \in \{a,b\}$. If his signal is k as well, 2 plays $x(\mu(k,0))$ as well. So upon observing 1 and 2 play $x(\mu(k,0))$ the DM holds the belief $\mu(k,k)$ and randomizes uniformly between the two locations. If his signal is not k, 2 plays x'_k upon signal k, with $x'_a := \overline{x}(x(\mu(a,0)))$ and $x'_b : \underline{x}(x(\mu(b,0)))$. Upon observing $(x(\mu(k,0)), x'_k)$ the DM's belief is α and she is, again, indifferent between the two locations.

Now 2 has no incentives to deviate as any of his deviations would be pooling (as both of his types can potentially benefit from the deviation) and thus induce the DM to have the belief $\mu(k,0)$, so that the DM strictly prefers 1's proposal. Similarly, but slightly more complicatedly, 1 has no incentive to deviate either as his deviations would be pooling as well. The best response of 2 would thus be to play $x(\alpha)$ independent of his signal and get selected with probability of, at least, 1/2 (exactly 1/2 if 1 deviated to $x(\alpha)$ and 1 otherwise). ■

Beliefs for Mixed Strategy Equilibria: The beliefs for the mixed strategy equilibria are $\mu_\sigma(x^a, x^a) =$

$$\frac{\alpha(1-\varepsilon)^2 + 2\alpha(1-\varepsilon)\varepsilon(1-\sigma) + \alpha\varepsilon^2(1-\sigma)^2}{\alpha(1-\varepsilon)^2 + 2\alpha(1-\varepsilon)\varepsilon(1-\sigma) + \alpha\varepsilon^2(1-\sigma)^2 + (1-\alpha)(1-\varepsilon)^2(1-\sigma)^2 + 2(1-\alpha)(1-\varepsilon)\varepsilon(1-\sigma)}, \tag{22}$$

$\mu_\sigma(x^a, x^b) = \frac{\alpha\varepsilon}{\alpha\varepsilon + (1-\alpha)[\varepsilon + (1-\varepsilon)(1-\sigma)]}$, $\mu_\sigma(x^b, x^b) = \mu(b,b)$ and $\pi_\sigma(b|b) = \sigma\pi(b|b)$.

B Separating D1 Equilibrium (not intended for publication)

Proposition 9 *If one signal is weak, there is a unique separating equilibrium outcome satisfying D1.*

Proof of Proposition 9: Recall that (i) $\pi(a|a) > 1 - \pi(b|b)$ and assume that (ii) $\pi(a|a) > 1/2 > \pi(b|b)$. The latter is without loss of generality insofar as D1 can only have additional bite compared to CK when one signal is weak.

Let $\theta \in \{\theta_a, \theta_b\}$ be the types of the receiver, where θ_k is the type that occurs whenever the other (non-deviating) sender plays the strategy he is supposed to play upon signal k.

Denote by $k \in \{a, b\}$ the deviating sender's types. It is first shown that rather trivially the set inclusion of mixed strategies does not eliminate any type if the receiver plays an unconditional mixed strategy, that is, if she selects the deviator with probability $v \in (0, 1)$ independently of her own type. This is trivially true because the probability $v(k)$ that makes type k indifferent between deviating and not is $v(a) = v(b) = 1/2$ for the simple reason that in equilibrium both are selected with probability 1/2.

Upon a given deviation x let $(v_{\theta_a}, v_{\theta_b})$ be the strategy of the receiver, where v_{θ_a} is the probability she votes for the deviator if her type is θ_a and v_{θ_b} is the corresponding probability when her type is θ_b.

In a separating equilibrium there is a unique deviation location $\tilde{x} \in (x^a, x^b)$ such that the receiver is indifferent between x^a and $\tilde{x}$ when of type θ_a and between $\tilde{x}$ and x^b when of type θ_b, keeping her beliefs about the deviator's type fixed. To see this, let $\rho_{\theta_k}(x)$ be the probability that the sender is of type a when making the deviation x such that the receiver is indifferent between x_k and x when she is of type θ_k, that is $\rho_{\theta_k}(x)$ is such that $u(\mu_{\theta_k}, x_k) = u(\mu_{\theta_k}, x)$. Notice that for $x \in (x^a, x^b)$, $\rho_{\theta_a}(x)$ is monotonically and continuously decreasing and satisfies $\lim_{x \to x^a} = 1$ and $\lim_{x \to x^b} = 0$. Analogously, $\rho_{\theta_b}(x)$ is a monotonically and continuously increasing function of x for $x \in (x^a, x^b)$ and satisfies $\lim_{x \to x^a} = 0$ and $\lim_{x \to x^b} = 1$. Consequently, there exists exactly one $\tilde{x}$ such that $\rho_{\theta_a}(\tilde{x}) = \rho_{\theta_b}(\tilde{x})$. The sender of type a will be indifferent between the deviation $\tilde{x}$ and equilibrium play if $v_{\theta_a}\pi(a|a) + v_{\theta_b}(1 - \pi(a|a)) = 1/2$. The largest mixture in v_{θ_a} (and smallest in v_{θ_b}) that keeps him indifferent is $\left(\frac{1}{2\pi(a|a)}, 0\right)$ and the smallest mixture in v_{θ_a} (and largest in v_{θ_b}) that keeps him indifferent is $\left(\frac{2\pi(a|a)-1}{2\pi(a|a)}, 1\right)$. Analogously, for type b the indifference condition is $v_{\theta_a}(1 - \pi(b|b)) + v_{\theta_b}\pi(b|b) = 1/2$, so that the largest and smallest mixtures that keep type b indifferent are, respectively, $\left(\frac{1}{2(1-\pi(b|b))}, 0\right)$ and $\left(\frac{1-2\pi(b|b)}{2(1-\pi(b|b))}, 1\right)$. Fact (i) implies that $\frac{1}{2\pi(a|a)} < \frac{1}{2(1-\pi(b|b))}$ and $\frac{1-2\pi(b|b)}{2(1-\pi(b|b))} < \frac{2\pi(a|a)-1}{2\pi(a|a)}$. So neither set includes the other one. Consequently, no type $k \in \{a, b\}$ can be deleted upon the deviation $\tilde{x}$, and the receiver is free to chose her beliefs in this instance.

Next consider any other deviation $x \in (x^a, x^b) \setminus \tilde{x}$. These deviations are such that the receiver can randomize for at most one of her types. That leaves the following four cases: 1. $(1, v_{\theta_b})$, 2. $(0, v_{\theta_b})$, 3. $(v_{\theta_a}, 1)$ and 4. $(v_{\theta_a}, 0)$. Since both sender types will benefit from a deviation if the receiver plays $(1, 0)$, case 1 does not eliminate any sender type. Notice also that for a given deviation x there are (different) beliefs about the deviating sender's types μ_{θ_a} and μ_{θ_b} that make it sequentially rational for the receiver to play $(v_{\theta_a}, 0)$ and $(0, v_{\theta_b})$, respectively, if and only if the deviation x satisfies $x^a < x < x^b$. The reason is simply that the receiver's indifference between x^a and x^b given the prior α makes it impossible for her to be indifferent between, say, x^a and $x > x^b$ because α will be the most favorable belief she can have for x when the non-deviating sender plays x^a (that is when she is of type θ_a).

Consider first the strategy $(v_{\theta_a}, 0)$.[46] The probabilities $v_{\theta_a}(k)$ that make type k indifferent satisfy

$$v_{\theta_a}(a)\pi(a|a) \;=\; 1/2 \Leftrightarrow v_{\theta_a}(a) = \frac{1}{2\pi(a|a)} \tag{23}$$

$$v_{\theta_a}(b)(1-\pi(b|b)) \;=\; 1/2 \Leftrightarrow v_{\theta_a}(b) = \frac{1}{2(1-\pi(b|b))}. \tag{24}$$

Fact (i) implies $v_{\theta_a}(b) > v_{\theta_a}(a)$. So the set of mixed strategies that makes type b better off than in equilibrium is a strict subset of the corresponding strategies for type a.

Consider now the strategy $(0, v_{\theta_b})$. The probabilities $v_{\theta_b}(k)$ that make type k indifferent are given by

$$v_{\theta_b}(a)(1-\pi(a|a)) \;=\; 1/2 \Leftrightarrow v_{\theta_b}(a) = \frac{1}{2(1-\pi(a|a))} \tag{25}$$

$$v_{\theta_b}(b)\pi(b|b) \;=\; 1/2 \Leftrightarrow v_{\theta_b}(b) = \frac{1}{2\pi(b|b)}. \tag{26}$$

Consequently, $v_{\theta_b}(b) < v_{\theta_b}(a)$. So the set of mixed strategies that makes type a better off than in equilibrium is a strict subset of the corresponding strategies for type b. (Similar computations can be done for $(v_{\theta_a}, 1)$ but adding these will not exclude any additional type since the set of types who can benefit 'most' is already maximal.) Consequently, for deviations $x \in (x^a, x^b)$ D1 has no bite in that it does not eliminate any types. Thus, for such deviations the receiver is free to choose her beliefs, and thus such deviations can be deterred without further ado.

Consider now a deviation $x < x^a$ and assume $x^a > x(a,a) \equiv x(\mu(a,a))$. This is for example the case if $x^a = x(a,0) \equiv x(\mu(a,0))$ as in the separating PSE. (The logic is quite similar for deviations $x > x^b$ but somewhat less important, b being the weak signal.) The fact that $x^a > x(a,a)$ implies that there is a location $x' < x(a,a)$ such that $u(\mu(a,a), x') = u(\mu(a,a), x^a)$. By continuity for any $x'' \in (x', x(a,a))$ there exist beliefs μ_{θ_a} that put positive probability on each of the deviating sender's types (a,b) that make the receiver of type θ_a indifferent and hence make selecting the deviator with probability v_{θ_a} sequentially rational. Because $\pi(a|a) > 1-\pi(b|b)$ such deviations are infinitely more likely to arise from a type a than from a type b (see the derivations above). Consequently, the unique best reply by the receiver will be $v_{\theta_a} = 1$ given x'' and consequently a necessary

[46] Observe that both types $k \in \{a,b\}$ will benefit strictly if the receiver plays $v_{\theta_a} = 1$ independently of v_{θ_b} because signal a is strong.

condition for x^a to be a separating D1 equilibrium location is $x^a \leq x(a, a)$. (For otherwise the deviation to $x(a, a)$ pays off for the type a.) Analogously, upon a deviation $x_b^* > x^b$ the optimal sequentially rational choice for the type θ_b receiver is to select the deviator with probability 1, that is $v_{\theta_b} = 1$. But since $\pi(b|b) < 1/2 < \pi(a|a)$ such a deviation does not pay off for either type of the sender.

Last notice that $x^a = x(a, a)$ and $x^b \neq x(a, a)$ such that $u(\alpha, x(a, a)) = u(\alpha, x^b)$ is also sufficient for these to be equilibrium locations in a D1 equilibrium: Deviations inside $(x(a, a), x^b)$ can be deterred as argued above. Deviations below $x(a, a)$ are interpreted as stemming from type a only, so that the receiver prefers $x(a, a)$ to the deviation (and a fortiori, if available, she prefers x^b to the deviation) while deviations to $x > x^b$ are interpreted as stemming from type b only, so that no special deterrence is required for these because being chosen with probability 1 if the opponent has received the weak signal and with probability 0 otherwise is not a profitable deviation.

Summarizing, we have a separating D1 equilibrium if and only if $x^a = x(a, a)$ and $x^b \neq x^a$ is such that $u(\alpha, x(a, a)) = u(\alpha, x^b)$. (The only if part follows from the insight that no $x < x(a, a)$ can be supported as an equilibrium location if signal a is strong.) ■

References

Ambrus, A., and S. Takahashi (2008): "Multi-sender cheap talk with restricted state spaces," *Theoretical Economics*, 3, 1–27.

Bagwell, K., and G. Ramey (1991): "Oligopoly Limit Pricing," *RAND Journal of Economics*, 22(2), 155–172.

Banks, J. S. (1990): "A Model of Electoral Competition with Incomplete Information," *Journal of Economic Theory*, 50, 309–325.

Banks, J. S., and J. Sobel (1987): "Equilibrium Selection in Signaling Games," *Econometrica*, 55(3), 647–661.

Battaglini, M. (2002): "Multiple Referrals and Multidimensional Cheap Talk," *Econometrica*, 70(4), 1379–1401.

Callander, S. (2005): "Electoral Competition in Heterogenous Districts," *Journal of Political Economy*, 113(5), 1116–1145.

Callander, S. (2008): "Political Motivations," *Review of Economic Studies*, 75, 671–697.

Callander, S., and S. Wilkie (2007): "Lies, Damned Lies and Politcal Campaigns," *Games and Economic Behavior*, 60, 262–286.

Canes-Wrone, B., M. C. Herron, and K. W. Shotts (2001): "Leadership and Pandering: A Theory of Executive Policymaking," *American Journal of Political Science*, 45(3), 532–550.

Che, Y., W. Dessein, and N. Kartik (forthcoming): "Pandering to Persuade," *American Economic Review.*

Cho, I.-K., and D. M. Kreps (1987): "Signaling Games and Stable Equilibria," *Quarterly Journal of Economics*, 102, 179–221.

Crawford, V., and J. Sobel (1982): "Strategic Information Transmission," *Econometrica*, 50(6), 1431–51.

Cummins, J. G., and I. Nyman (2005): "The Dark Side of Competitive Pressure," *RAND Journal of Economics*, 36(2), 361–377.

D'ASPREMONT, C., J. J. GABSZEWICZ, AND J.-F. THISSE (1979): "On Hotelling's 'Stability in Competition'," *Econometrica*, 47(5), 1145–50.

DOWNS, A. (1957): *An Economic Theory of Democracy.* Harper Collins.

FARRELL, J. (1993): "Meaning and Crdibility in Cheap-Talk Games," *Games and Economic Behvaior*, 5, 514–531.

FELGENHAUER, M. (forthcoming): "Revealing information in electoral competition," *Public Choice.*

GILLIGAN, T., AND K. KREHBIEL (1989): "Asymmetric Information and Legislative Rules with a Heterogenous Committee," *American Journal of Political Science*, 41(3), 459–490.

GROSSMAN, S. J., AND M. PERRY (1986): "Perfect Sequential Equilibrium," *Journal of Economic Theory*, 39, 97–119.

HEIDHUES, P., AND J. LAGERLÖF (2003): "Hiding information in electoral competition," *Games and Economic Behavior*, 42, 48–74.

HODLER, R., S. LOERTSCHER, AND D. ROHNER (2010): "Inefficient Policies and Incumbency Advantage," *Journal of Public Economics*, 94, 761–767.

HÖRNER, J., AND N. SAHUGUET (2007): "Costly Signalling in Auctions," *Review of Economic Studies*, 74, 173–206.

HOTELLING, H. (1929): "Stability in Competition," *Economic Journal*, 39, 41–57.

JENSEN, T. (2011): "Elections, Information, and State-Dependent Candidate Quality," *Mimeo, University of Copenhagen.*

KARTIK, N. (2009): "Strategic Communication with Lying Costs," *Review of Economic Studies*, 76(4), 1359–1395.

KARTIK, N., M. OTTAVIANI, AND F. SQUINTANI (2007): "Credulity, lies, and costly talk," *Journal of Economic Theory*, 134, 93–116.

KRISHNA, V., AND J. MORGAN (2001a): "A Model of Expertise," *Quarterly Journal of Economics*, pp. 747–775.

——— (2001b): "Asymmetric Information and Legislative Rules: Some Amendments," *American Political Science Review*, 95(2), 435–452.

Laslier, J.-F., and K. Van Der Straeten (2004): "Electoral competition under imperfect information," *Economic Theory*, 24, 419–46.

Lerner, A., and H. Singer (1937): "Some notes on duopoly and spatial competition," *Journal of Political Economy*, 45, 145–186.

Loertscher, S., and G. Muehlheusser (2008): "Global and local players in a model of spatial competition," *Economics Letters*, 98(1), 100–106.

——— (2011): "Sequential Location Games," *RAND Journal of Economics*, 42(3), 639–663.

Maskin, E., and J. Tirole (2004): "The Politician and the Judge: Accountability in Government," *American Economic Review*, 94(4), 1034–1054.

Osborne, M. J. (1995): "Spatial Models of Political Competition under Plurality Rule," *Canadian Journal of Economics*, 28(2), 261–301.

Pesendorfer, W., and A. Wolinsky (2003): "Second Opinions and Price Competition," *Review of Economic Studies*, 70(2), 417–438.

Prescott, E. C., and M. Visscher (1977): "Sequential location among firms with foresight," *Bell Journal of Economics*, 8(2), 378–393.

Riley, J. G. (2001): "Silver Signals: Twenty-Five Years of Signalling," *Journal of Economic Literature*, 39(2), 432–78.

Schultz, C. (1996): "Polarization and Inefficient Policies," *Review of Economci Studies*, 63(2), 331–44.

Sobel, J. (forthcoming): "Signaling Games," *Encyclopedia of Complexity and System Science (ed. by M. Sotomayor).*

Spence, M. (1973): "Job Market Signaling," *Quarterly Journal of Economics*, 87, 355–374.