Location Choice and Information Transmission

Simon Loertscher*

October 15, 2012

Abstract

Consider a location game with two agents and one decision maker who 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. This paper shows, however, that if the agents' information is precise, no equilibrium exists that satisfies the Intuitive Criterion and in which all the information is transmitted. In contrast, given sufficiently noisy information, a refined equilibrium with full information transmission exists. Though the agents play separating strategies in this equilibrium, their location choices are the same with positive probability because their information is the same with positive probability. A refined mixed strategy equilibrium in which some but not all information is transmitted always exists. However, with quadratic utility it is welfare inferior to the equilibrium in which both agents pool at the decision maker's preferred location given the prior if information is sufficiently precise.

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, Steven Callander, 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 workshop participants at AETW 2008 (Gold Coast), ESEM 2008 (Milano), AEA 2009 (San Francisco), ESAM 2009 (Canberra), APCC 2009 (Melbourne), IIOC 2010 (Vancouver), at the Australian National University, Deakin University, Queensland University of Technology, 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 such as voters or consumers will choose to locate at the bliss point of the median decision maker. If there is uncertainty about the state of the world and the agents are better informed about this state than the decision maker, will they reveal or conceal their superior private information about what is optimal for the decision maker? Will the decision maker be better informed as the agents' expertise increases? Will equilibrium welfare increase as their expertise increases?

To address questions like these, this paper studies the following 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. Second, both agents simultaneously choose locations on the real line. 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.[1] 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. Further, the DM employs a symmetric strategy.[2]

This model 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 tailor 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 her. Alternatively, consider political competition in a two-party system, in which the political candidates, whose expertise 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

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

[2] 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.

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 main results are the following. First, in any equilibrium with a symmetric strategy by the DM, 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 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 state is what the signal indicates exceeds (is less than) a half.[3] A signal being strong (weak) is equivalent to the probability, conditional on having received a specific signal, that the other signal is the same being larger (less) than a half. 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.[4] With quadratic utility, welfare in the separating PSE exceeds welfare in the pooling PSE whenever the former exists. Lastly, partly informative intuitive equilibria in which agents play a mixed strategy with binary support upon one signal always exist. However, with quadratic utility and very precise signals, welfare in the pooling PSE is larger than welfare in any of these non-degenerate mixed strategy equilibria.

The intuition for the main results, though subtle at times, can be developed rather straightforwardly. The reason that the Intuitive Criterion has no bite in refining pooling equilibria is essentially that there are no signaling costs in such an equilibrium as any

[3] 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.

[4] Appendix B shows that there is also a unique separating D1 equilibrium.

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, so that in a pooling equilibrium the DM's bliss point will be the same both on and off the equilibrium path.

The intuition why there are no intuitive separating equilibria when both signals are strong is roughly as follows. Consider a separating equilibrium. Upon receiving a strong signal, say a, each agent would be willing to deviate from the location prescribed by equilibrium if this allows him to win with certainty if the other agent has received the same signal and with probability zero otherwise. Observe that this deviation would be dominated by the expected equilibrium payoff upon receiving signal b because then the belief that the other agent has received signal a is less than a half (because both signals are strong). The Intuitive Criterion imposes the restrictions on the equilibrium locations of a separating equilibrium that each location be the maximizer of the DM's utility function given the belief that both signals have been the same for otherwise the deviation to such a location would pay off. On top of that, upon observing two different equilibrium locations the DM must be indifferent between the two locations because otherwise no agent could ever be induced to choose the location the DM likes less in this instance. But this now imposes three restrictions on the equilibrium locations of a separating equilibrium under the Intuitive Criterion that can only be satisfied non-generically.

The reason why there are separating intuitive equilibria when one signal is weak is that in this case the probability that the other agent has received the strong signal exceeds a 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.

This intuition suggests that, even if both signals are strong, mixed strategies by the agents can effectively render one signal weak in an intuitive equilibrium by making each agent's belief that the other agent chooses the same location to be less than a half upon receiving the effectively weak signal. In such a mixed strategy equilibrium, agents randomize over two locations upon the effectively weak signal and choose one of these locations with probability one upon receiving the other signal. As in the separating PSE when one signal is weak, the agents' locations diverge some but not all of the time. In the most informative of these equilibria more information is transmitted as agents' expertise increases. Interestingly, however, as the agents' signals become very precise welfare in any intuitive equilibrium in which agents employ non-degenerate mixed strategies with binary support is smaller than welfare in the unique pooling PSE if utility is quadratic. Therefore, ignorance can be a bliss.

This model exhibits the feature that separating equilibria satisfying standard refinements exist when agents' expertise is moderate but not when their expertise is strong. In the same spirit, the equilibria in which agents play non-degenerate mixed strategies and which reveal some information are welfare dominated by the equilibrium in which agents pander to the uninformed decision maker's bliss point if utility is quadratic.[5] Poole and Rosenthal (1991, 1993)'s empirical analysis of political positioning by representatives in the U.S. Congress demonstrates an increase in polarization or divergence when lawmakers' positions are pooled along party alliances over time. In light of the present model, this could be interpreted as resulting from increasingly complex policy issues.[6]

The present paper contributes to the literature on location games and to the literature on information transmission games. The standard assumption in the location games literature is that there is no uncertainty beyond the strategic uncertainty inherent to any game; 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). The cheap talk strand of the literature on information transmission games assumes that the messages agents send are intrinsically meaningless and thus costless; see, for example, Crawford and Sobel (1982), Krishna and Morgan (2001a,b), Battaglini (2002) and Ambrus and Takahashi (2008).[7]

The most important precursors to the present paper are Schultz (1996) and Callander (2008), who analyze models in which 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.[8] The main difference between Schultz's model and the present

[5]Observe that even when separating equilibria exist the locations chosen on the equilibrium path will diverge only some of the time because with positive probability the two agents will receive identical signals and therefore choose identical locations.

[6]Alternative explanations for platform divergence in political competitions have been put forth by Wittman (1977, 1983) and Calvert (1985). These require candidates to be partly motivated by policy and to be sufficiently uncertain about the median voter's preferred location.

[7]There is also a literature on models in which senders experience a cost of lying that goes back to Banks (1990) with more recent contributions by 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.

[8]Jensen (2011) analyzes a model with two possible states whose realization is known by the parties

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).[9]

Pandering by a better informed, self-interested agent to the preferences of a decision maker has recently been the focus of a number of papers 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 an incumbent in political office who is up for re-election chooses costly actions that make the median voter re-elect him more often than she would absent such wasteful behavior.[10] 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. Other contributions in this strand of literature include Laslier and Van Der Straeten (2004), who assume that the decision maker also receives a signal and show that the unique refined equilibrium outcome is separating, and Cummins and Nyman (2005), who assume that the agents' interests are partly aligned with those of the decision maker. In concurrent work, Felgenhauer (2012) introduces a third populist candidate into a setup that is otherwise identical to Heidhues and Lagerlöf's and shows that a separating equilibrium exists, and Kartik, Squintani, and Tinn (2012), who show that for a model with normally distributed states and signals and quadratic utility fully separating equilibria exist but involve inefficient, extreme policies.

Lastly, the paper also relates to the signaling and refinement literature by adapting standard refinements that have been introduced for games with one sender and one

but not the decision maker, which is similar to Schultz (1996).

[9]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 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.

[10]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.

receiver to games with multiple senders of possibly stochastic types.[11]

The remainder of this paper is organized as follows. Section 2 contains the model, preliminaries and adaption of standard equilibrium refinements to games with multiple senders. The equilibrium analysis is in Section 3. Section 4 analyzes welfare. Section 5 contains brief discussions of various extensions, and Section 6 concludes. All the proofs are in Appendix A. Appendix B analyzes equilibria satisfying D1.

2 The Model

This section introduces the basic setup, auxiliary results which the subsequent analysis builds upon, and the refinements that are used for this analysis.

2.1 Setup

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 $X = \mathbb{R}$. 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. 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 $u_2(\mu, x(\mu)) = 0$ with subscripts denoting the argument with respect to which partial derivatives are taken.[12] The assumptions on

[11] For an overview of this literature, see, for example, Riley (2001) and Sobel (2009).

[12] 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

$v(\omega, x)$ imply that $x(\mu)$ is unique and monotonically decreases in μ and that $x(1) = A$ and $x(0) = B$.[13] 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).[14] An additional restriction on $u(\mu, x)$ and its maximizer $x(\mu)$ will be imposed after updated beliefs have been introduced.

The DM's objective is to maximize her utility given her updated beliefs. Each agent's objective is to win the contest. Each agent's payoff of winning is normalized to one and the payoff of losing is assumed to be zero.

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. 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 $K \in \{A, B\}$. 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.[15] 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. A pure strategy for agent i is a location

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.

[13] 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. Totally differentiating the first order condition and using the second order condition reveals that $x'(\mu) < 0$.

[14] 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.

[15] 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 $-k \neq k$.

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. 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 this implies $\gamma(z, z) = 1/2$. 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 descriptions of competition between agents.

Beliefs and Optimal Locations Let $\lambda(x_1, x_2)$ denote the DM's posterior belief that the state is A when agents 1 and 2 choose locations x_1 and x_2, respectively. 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.[16] Analogously, let $\mu(a, b) = \mu(b, a)$ be this hypothetical belief when the agents receive divergent signals, 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

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

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 $u(\alpha, x(b, b))$

[16]The beliefs of the DM that the state is A, given locations, are denoted by λ while her – hypothetical – belief that the state is A given signals is denoted μ. Agents' beliefs about each other's signal are denoted by π.

Figure 1: Weak and strong signals.

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.

2.2 Preliminaries

Weak and Strong Signals It is useful to distinguish between what can be called weak and strong signals. Let $\pi(K|h)$ denote the probability that the state is K given signal h with $K \in \{A, B\}$ and $h \in \{a, b\}$. Signal k is quite naturally said to be strong if $\pi(K \mid k) > 1/2$ and called weak otherwise. Denote by $\pi(k|h)$ the probability that $s_j = k$ given that $s_i = h$ with $k, h \in \{a, b\}$ and $i \neq j$.

Lemma 1 *For $1-\alpha > \varepsilon$, signal b is strong, while for $\alpha > \varepsilon$, signal a is strong. Moreover, signal k is strong if and only if $\pi(k|k) > 1/2$.*

Observe also that $\pi(a|a) > 1-\pi(b|b)$ for any $(\alpha, \varepsilon) \in (0,1)\times(0,1/2)$.[17] 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 signal a is $\pi(A|a) = \frac{(1-\varepsilon)\alpha}{(1-\varepsilon)\alpha+(1-\alpha)\varepsilon}$, which is less than 1/2 since $\alpha < \varepsilon$ implies $(1-\varepsilon)\alpha < (1-\alpha)\varepsilon$.[18] According to Lemma 1, both signals are strong if and only if $\alpha \in (\varepsilon, 1-\varepsilon)$.

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.

Let x_i^k be 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

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

[18] Of course, $\pi(A|a) = \mu(a, 0)$.

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 $k \in \{a, b\}$*

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

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 selected 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 $x^a \neq x^b$.*

Separating PBE 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 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 $\pi(a|a)$ exceeding 1/2 that agent 2 received the same signal. Since on the equilibrium path 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.

2.3 Adapting Standard 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 adaptations 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.[19] 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

[19]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).

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.[20]

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.[21]

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.

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-the-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

[20] 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 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.

[21] 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.

game of the present model the PSE algorithm works as follows.[22]

First, let $\mathbf{x} = \{x_1^a, x_1^b, x_2^a, x_2^b\}$ be the set of locations whose elements the DM expects to observe with positive probability in a given PBE and assumes 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.[23] 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.

3 Equilibrium

3.1 Pooling Equilibria

The first result is a recurring theme within the literature on information transmission, namely that babbling is almost always an equilibrium that cannot be refined away.[24] Despite there being only two types, the Intuitive Criterion 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. 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-the-equilibrium-path belief

[22] 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.

[23] 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.

[24] 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 to signal their types.

satisfying the constraints imposed by PSE is $\mu = \alpha$. To see that no $x \neq x(\alpha)$ is a pooling PSE outcome, assume to the contrary that some x satisfying, say, $x < x(\alpha)$ is the outcome of a pooling PSE. As $u(\alpha, x)$ 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. 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. Note also that the unique pooling PSE outcome $x(\alpha)$ is the DM (or welfare) optimal pooling outcome.[25]

3.2 Separating Equilibria

Next consider separating equilibria.

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 $x_i^b = x^b$ for $i = 1, 2$ are strategies of a separating intuitive equilibrium only if*

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

The lemma is key. It adds two conditions on separating equilibrium locations when both signals are strong that will typically not hold simultaneously with condition (5) 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.[26]

Take any pair of separating locations $x^a, x^b \in (x(a, a), x(b, b))$. By Lemma 2 these locations must satisfy $u(\alpha, x^a) = u(\alpha, x^b)$, and each agent wins with equal probability. Consider next the beliefs of the DM following a deviation by agent 1 to a $z \in (x^b, x(b, b))$.

[25] 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.

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

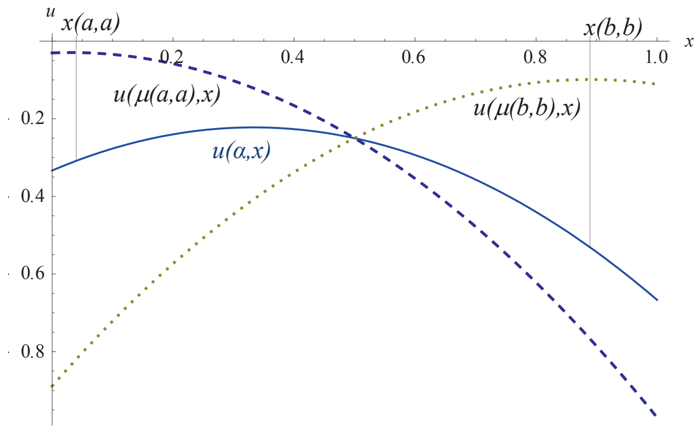

Figure 2: Necessary Conditions for Intuitive Separating Equilibria.

If agent 2 locates at x^a then even if the DM believes with probability 1 that the deviator observed signal b, her beliefs revert to the prior, and by the condition of Lemma 2 she selects agent 2. For any other beliefs over the deviator's type, the DM prefers agent 2 a fortiori. If instead agent 2 locates at x^b then the DM prefers agent 1 if she puts weight one on the deviator having received signal b. Thus, the deviating agent 1 is better off if the other agent locates at x^b and the DM believes his signal has been b and worse off if the other agent is at x^a. This gives the key observation: an agent with signal a believes the other agent is more likely to locate at x^a, as the signals are strong, and for this agent the deviation to z is not profitable. The Intuitive Criterion then requires that following the deviation z, all weight is put on the deviator having received 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.*

It is interesting that 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 $i = 1,2$ chooses $x_i^a = x(a,0)$ if $s_i = a$ and $x_i^b = x^b \neq x(a,0)$ if $s_i = b$, where $u(\alpha, x(a,0)) = u(\alpha, x^b)$. If a is the weak signal, agent $i = 1,2$ chooses $x_i^b = x(b,0)$ if $s_i = b$ and $x_i^a = x^a \neq x^b$ if $s_i = a$, where $u(\alpha, x(b,0)) = u(\alpha, x^a)$.*

Despite the fact that agents are purely motivated by winning, the model exhibits generic separating intuitive equilibria if one signal is weak. The multiplicity of separating intuitive equilibria arises because when one signal is weak both types of an agent can, a priori, benefit from a deviation given the equilibrium payoff. Therefore, CK does not impose restrictions on the beliefs over the deviating agent's types.

The construction of the unique PSE relies first, as any separating PBE, on an indifference condition of the DM in case the agents' signals differ and second on the fact that upon receiving the weak signal k an agent has not necessarily an 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.[27]

Expertise induces aggressive behavior The intuition for the non-existence of separating equilibria that satisfy the fairly minimal requirement of CK when both signals are strong and for the existence of of separating equilibria that satisfy the considerably stronger PSE requirements when one signal is weak is fairly clear. When both signals are strong, both agents are, in a sense, too well informed about the opponent's type. This makes them very aggressive and, in some sense, risk loving as they are willing to take gambles that would allow them to win only in the instance that the opponent has received the same signal. When one signal is weak, on the other hand, this aggressiveness vanishes because upon receiving a weak signal an agent would only be willing to take a gamble that allows him to win against an opponent with the opposite signal. Therefore, the presence of a weak signal eliminates one of the three constraints on equilibrium locations with strong signals – (i) $u(\alpha, x_a) = u(\alpha, x_b)$, (ii) $x^a = x(a,a)$ and (iii) $x^b = x(b,b)$ – that cannot be met generically.

This intuition suggests that mixed strategies by agents might induce some information transmission in a (refined) equilibrium even when both signals are strong because mixing upon one signal but not the other can effectively make the signal upon which mixing occurs weak. As shown next, this intuition is exactly correct. The equilibria in which agents play non-degenerately mixed strategies on the equilibrium path will be called mixed strategy equilibria.

3.3 Mixed Strategy Equilibria

Assume now that both signals are strong and, without loss of generality, that $u(\alpha, x(a,a)) > u(\alpha, x(b,b))$.[28] Let the 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 three important effects. First, upon observing (x^a, x^a) the DM's updated belief that state A has occurred, denoted $\lambda_\sigma(x^a, x^a)$, will be less than $\mu(a,a)$ for

[27] 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.

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

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 the other agent's mixed strategy is such that the probability that he chooses x^b is no more than 1/2. Thus, the Intuitive Criterion (and PSE) imposes no particular constraint on the location x^b to be chosen upon signal b. Third, upon signal a an agent's belief that the other agent plays x^a exceeds $\pi(a|a)$ for any $\sigma < 1$ because the other agent now plays x^a with probability $1 - \sigma$ even after receiving signal b.

Consider the following hybrid signal technology. If the state is A, the signal indicates a^σ with probability $1 - \varepsilon + \varepsilon(1 - \sigma) = 1 - \varepsilon\sigma$ and b^σ with probability $\varepsilon\sigma$. If the state is B, the signal indicates b^σ with with probability $(1 - \varepsilon)\sigma$ and a^σ with probability $\varepsilon + (1 - \varepsilon)(1 - \sigma)$. The constraint on the equilibrium location x^a_σ to be chosen upon signal a is then that it be $x^a_\sigma := x(\mu_\sigma(a^\sigma, 0))$, where

$$\mu_\sigma(a^\sigma, 0) := \frac{\alpha(1 - \varepsilon + \varepsilon(1 - \sigma))}{\alpha(1 - \varepsilon + \varepsilon(1 - \sigma)) + (1 - \alpha)(\varepsilon + (1 - \varepsilon)(1 - \sigma))}$$

is the DM's belief that state A has occurred if one agent's signal is a^σ while nothing is known about the other agent's signal.[29] Observe that $\mu_1(a^1, 0) = \mu(a, 0)$ and $\mu_0(a^0, 0) = \alpha$. Therefore, for any $\sigma \leq \frac{1}{2\pi(b|b)}$ one can construct such equilibria by letting

$$x^a_\sigma = x(\mu_\sigma(a^\sigma, 0)) \tag{7}$$

and by letting $x^b_\sigma \neq x^a_\sigma$ be such that

$$u(\mu_\sigma(a^\sigma, b^\sigma), x^a_\sigma) = u(\mu_\sigma(a^\sigma, b^\sigma), x^b_\sigma), \tag{8}$$

where

$$\mu_\sigma(a^\sigma, b^\sigma) = \frac{\alpha\varepsilon(1 - \varepsilon\sigma)}{\alpha\varepsilon(1 - \varepsilon\sigma) + (1 - \alpha)(1 - \varepsilon)(1 - (1 - \varepsilon)\sigma)}$$

is the belief of the DM that the state is A upon observing the signals (a^σ, b). Observe that given the strategy profile summarized by σ, $\lambda_\sigma(x^a_\sigma, x^b_\sigma) = \mu_\sigma(a^\sigma, b)$, which satisfies $\lambda_1(x^a_1, x^b_1) = \alpha$ and $\lambda_0(x^a_0, x^b_0) = \alpha\varepsilon/(\alpha\varepsilon + (1 - \alpha)(1 - \varepsilon)) < \alpha$. The most informative of these equilibria is the one with $\sigma^* := \frac{1}{2\pi(b|b)}$. Summarizing, one gets:

[29] Notice that x^a_σ 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 with $\sigma = \frac{1}{2\pi(b|b)}$ as a separate case.

Proposition 6 *There is a continuum of equilibria satisfying CK and PSE in which agents play mixed strategies when both signals are strong and* $u(\alpha, x(a,a)) \neq u(\alpha, x(b,b))$*. In these equilibria, some but not all information of the agents is transmitted to the DM. The equilibrium locations are given by (7) and (8) if* $u(\alpha, x(a,a)) > u(\alpha, x(b,b))$ *(and by analogous conditions in the converse case). Moreover, if* $u(\alpha, x(a,a)) \neq u(\alpha, x(b,b))$*,* x^b_σ *does not converge to* $x(\alpha)$ *as* σ *approaches 0. (The analogous statement is true for* x^a_σ *if* $u(\alpha, x(a,a)) < u(\alpha, x(b,b))$*). Any mixed strategy equilibrium with binary support that does not satisfy (7) and (8) does not satisfy CK.*

The last part of the proposition is proved in the appendix, which also contains the derivations of the relevant beliefs. Whether mixed strategy equilibria with supports that are not binary exist and satisfy CK (and if so, PSE) is an open technical question. The fact that x^b_σ differs from $x(\alpha)$ as σ goes to 0 has important and interesting implications for equilibrium welfare, as discussed in the next section. This discontinuity in locations arises because, even though $\mu_0(a^0, 0) = \alpha$ implies $x^a_\sigma = x(\alpha)$, $\lambda_0(x^a_0, x^b_0) < \alpha$ holds, which implies $x^b_\sigma > x(\alpha)$ even for $\sigma = 0$.

4 Welfare

Consider now equilibrium welfare. As agents do not care about locations, the DM's expected utility in equilibrium is the appropriate welfare measure. So as to have a unified expression for welfare in all three kinds of equilibria satisfying PSE – pooling, separating when one signal is weak, and mixed – assume that upon signal b agents choose x^b with probability σ and x^a with probability $1-\sigma$ while upon signal a each agent chooses x^a. Let $W(\alpha, x^a, x^b, \sigma)$ denote welfare given the prior α, equilibrium locations x^a and x^b and a mixture σ upon signal b. Welfare is

$$W(\alpha, x^a, x^b, \sigma) = \alpha W_A(x^a, x^b, \sigma) + (1-\alpha) W_B(x^a, x^b, \sigma), \tag{9}$$

where $W_K(x^a, x^b, \sigma)$ denotes expected welfare conditional on the state being $K \in \{A, B\}$ and is given as

$$\begin{aligned} W_A(x^a, x^b, \sigma) &= u(1, x^a)((1-\varepsilon)^2 + 2(1-\varepsilon)\varepsilon(1-\sigma/2) + \varepsilon^2(1-\sigma)^2 + \varepsilon^2(1-\sigma)\sigma) \\ &+ u(1, x^b)((1-\varepsilon)\varepsilon\sigma + \varepsilon^2\sigma^2 + \varepsilon^2(1-\sigma)\sigma) \quad \text{and} \\ W_B(x^a, x^b, \sigma) &= u(0, x^a)(\varepsilon^2 + 2(1-\varepsilon)\varepsilon(1-\sigma/2) + (1-\varepsilon)^2((1-\sigma)\sigma + (1-\sigma)^2)) \\ &+ u(0, x^b)((1-\varepsilon)^2\sigma^2 + (1-\varepsilon)^2(1-\sigma)\sigma + (1-\varepsilon)\varepsilon\sigma). \end{aligned}$$

Welfare in the pooling PSE, denoted W^{pool}, is now simply $W^{pool} = W(\alpha, x(\alpha), x(\alpha), \sigma)$. Welfare in the separating PSE, provided it exists, is denoted as W^{sep} and given by $W^{sep} = W(\alpha, x(a, 0), x^{b,pse}, 1)$ where $x^{b,pse} \neq x(a, 0)$ is such that $u(\alpha, x(a, 0)) = u(\alpha, x^{b,pse})$ and where $\alpha > 1/2$ is assumed, so that signal b (if any) is the weak signal. Welfare in a mixed strategy equilibrium satisfying CK and PSE is denoted W^{mse} and given as $W^{mse} = W(\alpha, x_\sigma^a, x_\sigma^b, \sigma)$ with $\sigma \leq \frac{1}{2\pi(b|b)}$. For the remainder of this section, attention is confined to the case with quadratic utility with $A = 0$ and $B = 1$, so that $u(\mu, x) = -\mu x^2 - (1-\mu)(1-x)^2$, implying $x(\mu) = 1 - \mu$.

Separating PSE are better than the pooling PSE In separating equilibria, agents' information is completely revealed. But in itself this does not imply that welfare in a separating PSE exceeds welfare in the pooling PSE because there is a priori no reason to think that the separating equilibrium locations should be optimal (and they never are, as shown below). Nonetheless, the next proposition shows that the informative, separating PSE is also welfare superior to the pooling PSE whenever a separating PSE exists.

Proposition 7 *Whenever a separating PSE exists, it generates higher welfare than the pooling PSE, that is* $W^{sep} > W^{pool}$.

Ignorance can be a bliss The expression for W^{mse} is somewhat unwieldy in general. Numerical calculations (which are not displayed) show that for a wide range of parameter values (α, ε) welfare in the informative mixed strategy equilibria satisfying CK exceeds welfare in the pooling PSE. Interestingly, however, if expertise is sufficiently large – that is, ε sufficiently small – welfare is higher in the pooling PSE than in any mixed strategy equilibrium with binary support satisfying CK:

Proposition 8 *For sufficiently small* $\varepsilon > 0$ *and any* $\sigma > 0$, $W^{pool} > W^{mse}$.

In other words, remaining ignorant can be a bliss for the DM. The intuition is as a follows. As ε goes to 0, so does $\lambda_\sigma(x_\sigma^a, x_\sigma^b)$. As the mixed strategy equilibrium requires the DM to be indifferent upon observing two different locations on the equilibrium path, the location x_σ^b will be extreme, indeed larger than B, so that both locations generate a very low payoff to the DM conditional on the information available to her in this instance. This is inefficient and makes the DM prefer the pooling PSE to any mixed strategy equilibrium with $\sigma > 0$ satisfying CK.

Second-best locations Consider now a social planner's problem who can choose locations as a function of signals and who aims at maximizing the DM's expected welfare. Assume that the planner is constrained to choose identical locations x^k upon signal k for either agent, which makes comparison with equilibrium locations possible and insightful. However, the planner does not face the constraint that the DM be indifferent between the locations x^a and x^b if the two signals differ. That is, $u(\alpha, x^a) = u(\alpha, x^b)$ is not required. Without loss of generality assume $\alpha \geq 1/2$, which will imply $u(\alpha, x^{a,sb}) \geq u(\alpha, x^{b,sb})$ for the second-best locations $(x^{a,sb}, x^{b,sb})$. The timing is as follows. The planner observes the signals, and chooses a location for each agent that only depends on the realization of this agent's signal. The DM observes the locations, updates her beliefs and chooses the agent whose location she prefers given her updated beliefs.

Second-best welfare as a function of x^a and x^b is

$$W^{SB}(\alpha, x^a, x^b) = \alpha W_A^{SB}(x^a, x^b) + (1-\alpha) W_B^{SB}(x^a, x^b), \tag{10}$$

where $W_K^{SB}(x^a, x^b)$ denotes expected welfare conditional on the state being K with $K \in \{A, B\}$, which is given as $W_A^{SB}(x^a, x^b) = u(1, x^a)((1-\varepsilon)^2 + 2(1-\varepsilon)\varepsilon) + u(1, x^b)\varepsilon^2$ and $W_B^{SB}(x^a, x^b) = u(0, x^a)(\varepsilon^2 + 2\varepsilon(1-\varepsilon)) + u(0, x^b)(1-\varepsilon)^2$.

Proposition 9 *Assume $\alpha \geq 1/2$. The second-best locations are*

$$x^{a,sb} = \frac{(1-\alpha)(2-\varepsilon)\varepsilon}{\alpha(1-2\varepsilon) + (2-\varepsilon)\varepsilon} \quad \textit{and} \quad x^{b,sb} = \frac{(1-\alpha)(1-\varepsilon)^2}{\alpha(1-2\varepsilon) + (1-\varepsilon)^2}. \tag{11}$$

Moreover, assuming signal b is weak and denoting the separating PSE locations by $x^{a,pse}$ and $x^{b,pse}$, the following holds

$$x^{a,pse} < x^{a,sb} \quad \textit{and} \quad x^{b,pse} < x^{b,sb}. \tag{12}$$

According to Proposition 9, relative to second-best the separating PSE locations are biased towards state A towards which the prior is biased.

5 Discussion

The basic model studied thus far lends itself to a number of natural extensions and variations. Some of these extensions are discussed informally in the following. The discussion is substantiated in a separate web appendix, which contains the formal models and proofs.

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.

Cheap Talk Augmented Model It may be natural to think that the set of equilibrium outcomes satisfying a given refinement changes if the agents can send cheap talk messages on top of the locations they choose because the locations, although cheap from the point of view of the agents, are not cheap from the perspective of the DM. Therefore, one may wonder whether adding an additional communication device in the form of cheap talk may have an effect on the equilibrium outcomes. The answer turns out to be no. This is almost immediate for pooling equilibria because the incentives to deviate from a given location-message pair are exactly the same as the incentives to deviate from a given location in the model without messages. For separating equilibria, the key insight underlying the invariance of equilibrium outcomes is that the only substantively novel kind of deviations in the cheap talk augmented game are deviations in messages only because if a deviation in messages and locations paid off it would pay off to choose a deviating location in the game without messages. But given that the deviation is only in the message but not in the location, the DM will correctly infer the deviator's signal, so that by construction of the equilibrium the deviation cannot pay off.

Outside Option As in parts of Krishna and Morgan (2001b) and in Gilligan and Krehbiel (1989), one could assume that there is an exogenously given outside option, such as a status quo or default location x^O, the DM can threaten to choose. It is not hard to show that $x(\alpha)$ is still the unique pooling PSE outcome. 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.[30]

[30] Interpreting the outside option as the policy proposed by a populist third party that always chooses the uninformed voter's bliss point policy, Felgenhauer (2012) shows that in a setup with binary states, signals and actions a fully informative equilibrium exists. The results above complement and extend these findings by showing that a fully informative equilibrium may come into existence even when the outside option is not perfectly populist.

Imperfect Commitment In the analysis thus far, agents have been assumed to be committed to their locations in the sense that they will not change the locations even after winning the contest. Though plausible for some applications, it seems unlikely to be strictly satisfied in, say, political economics because the winner of the contest may not necessarily implement the policy with which he campaigned. In particular, every agent knows that the policy he campaigned with is suboptimal for the DM given the information the agent has ex post, that is, after not only observing his own signal but also the location chosen by his opponent. On the other hand, political parties cannot at will and zero costs deviate from the platforms with which they campaigned because of credibility concerns and because rebuilding a (shadow) cabinet comes at some cost. This provides a motivation to analyze the model under the alternative assumption that the winning party may deviate with some probability from the location it announced during the campaign.[31] Specifically, under the assumption that the winning party implements the policy that is optimal for the DM, given the information available in the economy, with a commonly known probability that is less than one, and otherwise implements the policy with which it campaigned, one can show that the set of equilibrium locations satisfying a given refinement is invariant with respect to the probability that the party is flexible. However, welfare is of course increasing in this probability.

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. Though the location space is much larger, optimal locations, and optimal deviations, are confined to the line that connects states A and state B. Therefore, the results from the model with a one-dimensional space carry over to a multi-dimensional space without much additional ado.

Sequential Moves Though modeling agents as moving simultaneously, arguments for why sequential moves are plausible in some instances can also be made (see, for example, Krishna and Morgan (2001a,b) and Pesendorfer and Wolinsky (2003)). Interestingly, with sequential moves there are fully separating PSE regardless of the level of agents' expertise, that is for any $(\alpha, \varepsilon) \in (0,1) \times (0,1/2)$. In this PSE, each agent chooses $x(k,0)$

[31] Banks (1990), Kartik and McAfee (2007) and Callander (2008) model candidates' imperfect flexibility, which is often interpreted positively as character, in a similar way.

upon signal k on the equilibrium path with $k \in \{a, b\}$. 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 the second-mover can be preempted by the first-mover 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 the first-mover can be countered by the second-mover by playing $x(\alpha)$ upon a deviation by the first-mover, which is the DM's optimal location given the hypothesis that both agents are pooling.[32]

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. When both signals are strong, the opposite obtains in the present paper. This difference arises because in Krishna and Morgan (2001a) both agents know the state, and hence each other's type. This lack of uncertainty makes divergent actions an off-the-equilibrium path observation in Krishna and Morgan with simultaneous moves but creates strong incentives for the first-mover to induce the second-mover to "lie" in the model with sequential moves. In the present setup, on the other hand, the agents are uncertain about the state and each other's type. This makes divergence inevitable in any separating equilibrium with simultaneous moves but at the same time mitigates the Stackelberg leader incentives for the first-mover with sequential moves.

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' expertise is sufficiently low, 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). An equilibrium satisfying the Intuitive Criterion and PSE in which agents played mixed strategies and thus reveal some but not all information always exists. However, if the agents' expertise is very strong

[32] 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 because the endogenous signaling costs disappear for the second-mover. Assuming equilibrium play by the first-mover, the second -mover "knows" the first-mover's signal. Therefore, he cannot credibly reveal his signal by off equilibrium behavior.

and utility quadratic, welfare in any such mixed strategy equilibrium is lower than in the pooling PSE, in which agents pander to the decision maker's preferred location given the prior. Under these conditions, ignorance is thus a bliss.

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.

Appendix

A Proofs

Proof of Lemma 1: $\pi(A \mid a) = \frac{\alpha(1-\varepsilon)}{\alpha(1-\varepsilon)+(1-\alpha)\varepsilon}$, $\pi(B \mid b) = \frac{(1-\alpha)(1-\varepsilon)}{(1-\alpha)(1-\varepsilon)+\alpha\varepsilon}$, $\pi(B \mid a) = \frac{(1-\alpha)\varepsilon}{\alpha(1-\varepsilon)+(1-\alpha)\varepsilon}$ and $\pi(A \mid b) = \frac{\alpha\varepsilon}{\alpha\varepsilon+(1-\alpha)(1-\varepsilon)}$ by Bayes' rule. It is straightforward to verify that $\pi(A|a) > 1/2$ if and only if $\alpha > \varepsilon$ and $\pi(B|b) > 1/2$ if and only if $1-\alpha > \varepsilon$.

Conditional on $s_i = a$ the probability the $s_j = a$ with $j \neq i$ is $\pi(a \mid a) = \pi(A \mid a)(1-\varepsilon) + \pi(B \mid a)\varepsilon = \frac{\alpha(1-\varepsilon)^2+(1-\alpha)\varepsilon^2}{\alpha(1-\varepsilon)+(1-\alpha)\varepsilon}$ and conditional on signal $s_i = b$ the probability that $s_j = b$ is $\pi(b \mid b) = \pi(B|b)(1-\varepsilon)+\pi(A|b)\varepsilon = \frac{\alpha\varepsilon^2+(1-\alpha)(1-\varepsilon)^2}{\alpha\varepsilon+(1-\alpha)(1-\varepsilon)}$. Notice that $\pi(a|a)$ and $\pi(b|b)$ are weighted averages of the form $x(1-\varepsilon)+\varepsilon(1-x)$ with $x = \pi(A|a)$ and $\pi(B|b)$, respectively. As $x(1-\varepsilon)+\varepsilon(1-x)$ is increasing in x and equal to $1/2$ at $x = 1/2$, it follows that $\pi(K|k) > 1/2$ is equivalent to $\pi(k|k) > 1/2$ for $K \in \{A, B\}$ and $k \in \{a, b\}$. ■

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

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

To simplify notation, let $\theta_k \equiv \pi(k \mid k)$, $x \equiv \gamma(x_1^a, x_2^a)$, $y \equiv \gamma(x_1^a, x_2^b)$, $c \equiv 1 - \gamma(x_1^b, x_2^a)$ and $d \equiv \gamma(x_1^b, x_2^b)$.

The incentive constraints (3) can now be written as

$$\begin{aligned} U_1[x_1^a \mid s_1 = a] = \theta_a x + (1-\theta_a) y &\geq \theta_a(1-c) + (1-\theta_a)d \\ &= U_1[x_1^b \mid s_1 = a] \quad (13) \\ U_2[x_2^a \mid s_2 = a] = \theta_a(1-x) + (1-\theta_a)c &\geq \theta_a(1-y) + (1-\theta_a)(1-d) \\ &= U_2[x_2^b \mid s_2 = a] \quad (14) \\ U_1[x_1^b \mid s_1 = b] = \theta_b d + (1-\theta_b)(1-c) &\geq \theta_b y + (1-\theta_b)x \\ &= U_1[x_1^a \mid s_1 = b] \quad (15) \\ U_2[x_2^b \mid s_2 = b] = \theta_b(1-d) + (1-\theta_b)(1-y) &\geq \theta_b c + (1-\theta_b)(1-x) \\ &= U_2[x_2^a \mid s_2 = b]. \quad (16) \end{aligned}$$

Adding (13) and (14) yields $1 \leq y + c$ while adding (15) and (16) implies $1 \geq y + c$. Thus, $y + c = 1$ holds. Now,

$$U_1[x_1^a \mid s_1 = a] + U_2[x_2^a \mid s_2 = a] = \theta_a + (1-\theta_a)(y+c) = 1 \quad (17)$$

and

$$U_1[x_1^b \mid s_1 = b] + U_2[x_2^b \mid s_2 = b] = \theta_b + (1-\theta_b)(2-(y+c)) = 1, \quad (18)$$

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

$$U_1[x_2^a \mid s_1 = a] = \theta_a \gamma(x_2^a, x_2^a) + (1-\theta_a)\gamma(x_2^a, x_2^b), \quad (19)$$

or x_2^b, in which case he would get

$$U_1[x_2^b \mid s_1 = a] = \theta_a \gamma(x_2^b, x_2^a) + (1-\theta_a)\gamma(x_2^b, x_2^b). \quad (20)$$

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

$$U_1[x_2^a \mid s_1 = a] = \theta_a \frac{1}{2} + (1-\theta_a)\gamma(x_2^a, x_2^b) = \gamma(x_2^a, x_2^b) + \theta_a \left(\frac{1}{2} - \gamma(x_2^a, x_2^b)\right) \quad (21)$$

and

$$U_1[x_2^b \mid s_1 = a] = \theta_a(1 - \gamma(x_2^a, x_2^b)) + (1 - \theta_a)\frac{1}{2} = \frac{1}{2} + \theta_a\left(\frac{1}{2} - \gamma(x_2^a, x_2^b)\right). \tag{22}$$

If $\gamma(x_2^a, x_2^b) \geq \frac{1}{2}$ the expression in (21) weakly exceeds 1/2 and the expression in (22) 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 $\lambda(x_1, x) = \frac{\alpha(1-\varepsilon)}{\alpha(1-\varepsilon)+(1-\alpha)\varepsilon} = \mu(a, 0) < 1$ and $\lambda(x_1, x) = \frac{\alpha\varepsilon}{\alpha\varepsilon+(1-\alpha)(1-\varepsilon)} = \mu(b, 0) > 0$. 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 $\lambda(x^a, x^b) = \frac{\alpha(1-\varepsilon)\varepsilon}{\alpha(1-\varepsilon)\varepsilon+(1-\alpha)(1-\varepsilon)\varepsilon} = \alpha$. 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 a 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 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

$$x = \frac{A - 4A\varepsilon + 2B\varepsilon + 7A\varepsilon^2 - 5B\varepsilon^2 - 4A\varepsilon^3 + 4B\varepsilon^3}{1 - 2\varepsilon + 2\varepsilon^2}. \tag{23}$$

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

$$x = \frac{B + 2A\varepsilon - 4B\varepsilon - 5A\varepsilon^2 + 7B\varepsilon^2 + 4A\varepsilon^3 - 4B\varepsilon^3}{1 - 2\varepsilon + 2\varepsilon^2}. \tag{24}$$

For $B > A$, it is readily established that the right-hand side of (23) is less than the right-hand side of (24) 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, $\pi(b \mid b) < 1/2$ 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. ■

Proof of Proposition 6: The beliefs for the mixed strategy equilibria are

$$\lambda_\sigma(x^a_\sigma, x^a_\sigma) = \frac{\alpha[(1-\varepsilon) + \varepsilon(1-\sigma)]^2}{\alpha[(1-\varepsilon) + \varepsilon(1-\sigma)]^2 + (1-\alpha)[(1-\varepsilon)(1-\sigma) + \varepsilon]^2}. \tag{25}$$

Observe that $\lambda_1(x^a_1, x^a_1) = \mu(a, a)$ and $\lambda_0(x^a_0, x^a_0) = \alpha$ as it should. To see that $\lambda_\sigma(x^a_\sigma, x^a_\sigma)$ is as in formula (25) observe that under the stipulated strategy profile the probability of observing (x^a_σ, x^a_σ) in state A is

$$\Pr(x^a_\sigma, x^a_\sigma|\omega = A) = (1-\varepsilon)^2 + 2(1-\varepsilon)\varepsilon(1-\sigma) + (\varepsilon(1-\sigma))^2$$

while the probability of observing (x^a_σ, x^a_σ) in state B is

$$\Pr(x^a_\sigma, x^a_\sigma|\omega = B) = ((1-\varepsilon)(1-\sigma))^2 + 2(1-\varepsilon)\varepsilon(1-\sigma) + \varepsilon^2.$$

As $\lambda_\sigma(x^a, x^a) = \frac{\alpha \Pr(x^a, x^a|\omega=A)}{\alpha \Pr(x^a,x^a|\omega=A) + (1-\alpha)\Pr(x^a,x^a|\omega=B)}$, the expression follows after plugging in these probabilities. The belief $\lambda_\sigma(x^a_\sigma, x^b_\sigma)$ is

$$\lambda_\sigma(x^a_\sigma, x^b_\sigma) = \frac{\alpha\varepsilon(1-\varepsilon\sigma)}{\alpha\varepsilon(1-\varepsilon\sigma) + (1-\alpha)(1-\varepsilon)[\varepsilon + (1-\varepsilon)(1-\sigma)]} \tag{26}$$

since $\Pr(x^a_\sigma, x^b_\sigma|\omega = A) = 2(1-\varepsilon)\varepsilon\sigma + 2\varepsilon^2(1-\sigma)\sigma$ and $\Pr(x^a, x^b|\omega = B) = 2(1-\varepsilon)\varepsilon\sigma + 2(1-\varepsilon)^2(1-\sigma)\sigma$. Observe that $\lambda_1(x^a_1, x^b_1) = \alpha$ and $\lambda_0(x^a_0, x^b_0) = \frac{\alpha\varepsilon}{\alpha\varepsilon + (1-\alpha)(1-\varepsilon)} = \mu(b, 0) < \alpha$. Finally, $\lambda_\sigma(x^b_\sigma, x^b_\sigma) = \mu(b, b)$.

The last part of the proposition follows because if x^a_σ were different from $x(\mu_\sigma(a^\sigma, 0))$, an agent of type a would have an incentive to deviate to $x(\mu_\sigma(a^\sigma, 0))$, assuming $u(\alpha, x(a, a)) > u(\alpha, x(b, b))$. If the equilibrium did not satisfy (8), agents of type b could not be indifferent between choosing x^a_σ and x^b_σ. ■

Proof of Proposition 7: Tedious algebra reveals that

$$W^{sep} = -(1-\alpha)\alpha\frac{(-4\alpha^3(2\varepsilon-1)^3 + 4\alpha^2(1-2\varepsilon)^2(3\varepsilon-1) + \alpha(1-2\varepsilon(8(\varepsilon-1)\varepsilon+3)) + \varepsilon^2)}{(-2\alpha\varepsilon+\alpha+\varepsilon)^2}.$$

This is larger than $W^{pool} = -\alpha(1-\alpha)$ for any $\alpha > \frac{1-4\varepsilon}{4-8\varepsilon}$. Now α has to exceed $1-\varepsilon$ for signal b to be weak, and hence for a separating PSE to exist, and $1-\varepsilon > \frac{1-4\varepsilon}{4-8\varepsilon}$ for all $\varepsilon \in (0, 1/2)$, the result follows. ■

Proof of Proposition 8: Algebra yields

$$\begin{aligned} W^{mse} &= \alpha\Bigg(((-1+\alpha)^2)((1+(-1+\varepsilon)\sigma)^2)\Bigg\{-\frac{1}{(\sigma(-2\alpha\varepsilon+\alpha+\varepsilon-1)+1)^2} \\ &- \frac{4\alpha(\varepsilon-1)\varepsilon(2\varepsilon-1)\sigma(\varepsilon\sigma-1)}{\left(\alpha(2\varepsilon-1)(\sigma-1)+(\varepsilon-1)^2\sigma+\varepsilon-1\right)^2(\sigma(2\alpha\varepsilon-\alpha-\varepsilon+1)-1)}\Bigg\} \\ &+ (1-\alpha)\alpha(-1+\varepsilon\sigma)^2\Bigg\{-\frac{1}{(\sigma(-2\alpha\varepsilon+\alpha+\varepsilon-1)+1)^2} \\ &+ \frac{4(\alpha-1)(\varepsilon-1)\varepsilon(2\varepsilon-1)\sigma((\varepsilon-1)\sigma+1)}{\left(\alpha(2\varepsilon-1)(\sigma-1)+(\varepsilon-1)^2\sigma+\varepsilon-1\right)^2(\sigma(\alpha(2\varepsilon-1)-\varepsilon+1)-1)}\Bigg\}\Bigg), \end{aligned}$$

which is continuous in ε. At $\varepsilon = 0$,

$$W^{mse} = -\frac{(\alpha-1)\alpha((\alpha-1)(\sigma-2)\sigma-1)}{((\alpha-1)\sigma+1)^2}.$$

Therefore, at $\varepsilon = 0$

$$W^{mse} - W^{pool} = -\frac{(\alpha-1)^2\alpha^2\sigma^2}{((\alpha-1)\sigma+1)^2}, \tag{27}$$

which is negative for any $\sigma > 0$. By continuity, $W^{mse} - W^{pool} < 0$ will hold for some $\varepsilon > 0$. ■

Proof of Proposition 9: Plugging in the quadratic utility function into the formula for $W^{SB}(\alpha, x^a, x^b)$ and maximizing with respect to x^a and x^b yields the locations in (11). Algebra then yields the inequalities in (12). ■

B D1 Equilibria (not intended for publication)

This appendix analyzes equilibria that satisfy D1 (Banks and Sobel, 1987) adapted to games with multiple senders.

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.[33] 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.

Pooling D1 Equilibria For the same reasons that there is a continuum of pooling equilibria satisfying CK (Proposition 2), there is a continuum of pooling equilibria satisfying D1.

Separating D1 Equilibria

Proposition 10 *If one signal is weak, there is a unique separating equilibrium outcome satisfying D1.*

[33] See also Banks and Sobel (1987) or Sobel (2009).

Proof of Proposition 10: 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 $(v_{\theta_a}, 0)$.[34] The probabilities $v_{\theta_a}(k)$ that make type k indifferent satisfy

$$v_{\theta_a}(a)\pi(a|a) = 1/2 \Leftrightarrow v_{\theta_a}(a) = \frac{1}{2\pi(a|a)} \tag{28}$$

$$v_{\theta_a}(b)(1 - \pi(b|b)) = 1/2 \Leftrightarrow v_{\theta_a}(b) = \frac{1}{2(1 - \pi(b|b))}. \tag{29}$$

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

$$v_{\theta_b}(a)(1 - \pi(a|a)) = 1/2 \Leftrightarrow v_{\theta_b}(a) = \frac{1}{2(1 - \pi(a|a))} \tag{30}$$

$$v_{\theta_b}(b)\pi(b|b) = 1/2 \Leftrightarrow v_{\theta_b}(b) = \frac{1}{2\pi(b|b)}. \tag{31}$$

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

[34] 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.

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 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.

Calvert, R. (1985): "Robustness of the Multidimensional Voting Model: Candidate Motivation, Uncertainty, and Convergence," *American Journal of Political Science*, 29, 69 – 95.

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. (2012): "Revealing information in electoral competition," *Public Choice*, 153(1), 55 – 68.

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., and R. P. McAfee (2007): "Signaling Character in Electoral Competition.," *American Economic Review*, 97(3), 852 – 870.

Kartik, N., M. Ottaviani, and F. Squintani (2007): "Credulity, lies, and costly talk," *Journal of Economic Theory*, 134, 93–116.

Kartik, N., F. Squintani, and K. Tinn (2012): "Inormation Revelation and (Anti-)Pandering in Elections," *Mimeo.*

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.

Poole, K., and H. Rosenthal (1993): *Congress : a political-economic history of roll call voting*. Oxford University Press.

Poole, K. T., and H. Rosenthal (1991): "Patterns of Congressional Voting," *American Journal of Political Science*, 35(1), 228 – 278.

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. (2009): "Signaling Games," *Encyclopedia of Complexity and System Science (ed. by M. Sotomayor)*.

Wittman, D. A. (1977): "Candidates with Policy Preferences: A Dynamic Model.," *Journal of Economic Theory*, 14(1), 180 – 189.

——— (1983): "Candidate Motivation: A Synthesis of Alternative Theories," *American Political Science Review*, 77(1), 142 – 157.