Information Transmission in Electoral Competition

Simon Loertscher*

July 15, 2009

Abstract

This paper analyzes a model where two office motivated parties first receive noisy but informative private signals about which of two states is true and then simultaneously propose policies. The voter receives no signal and her bliss point policy varies monotonically with her belief about the two states. Observing the policies, she updates her beliefs and votes for the party whose policy maximizes her expected utility. If the prior on one state is smaller than the signal's error probability, the model exhibits a fully separating refined equilibrium. However, as the signals' error probability goes to zero, in the generically unique refined equilibrium both parties pool at the voter's uninformed bliss point policy. Therefore, the better the parties' information the less likely will this information be transmitted in equilibrium. The model thus predicts convergence to the median voter's uninformed bliss point policy on issues that are well understood and the possibility of divergent and state dependent policy proposals when parties' understanding of the issues is less advanced e.g. because of their novelty.

Keywords: Information Transmission, Sender Receiver Games, Political Economics.

JEL-Classification: C72, D72, D83

*Department of Economics, Economics & Commerce Building, University of Melbourne, Victoria 3010, Australia. Email: simonl@unimelb.edu.au. I want to thank Yuelan Chen, Emmanuel Frot, Roland Hodler, Andrew McLennan, Carlos Pimienta and Tom Wilkening and seminar participants at AETW 2008 (Gold Coast), ESEM 2008 (Milano), AEA 2009 (San Francisco), ESAM 2009 (Canberra), Australian National University and at the University of Auckland for most valuable comments and discussions and an anonymous referee for a suggestion that ultimately led to the present paper. This paper has grown out of joint work with Aleksander Berentsen and Esther Bruegger, whose comments on a much earlier version I gratefully acknowledge. Any errors are mine.

1 Introduction

In democracies typically poorly informed voters choose in low frequency amongst policies proposed by competing parties,[1] i.e. by professional organizations that are permanently engaged in political debates and political decision making. Therefore, the expertise of political parties is likely to exceed the one of the voters. An important question is thus whether voters can rely on information acquisition and information transmission by political parties in order to make informed decisions.

To address this question, this paper analyzes a model where two office motivated parties are better informed than the median voter about which of two possible states of nature is true. The policy space is continuous and the voter's unique bliss point policy varies monotonically with her belief.[2] Each party first observes a private signal that indicates the true state with probability $1-\varepsilon$ and the wrong state with probability $\varepsilon < 1/2$. Second, the two parties simultaneously propose their policies, which they are committed to implement upon being elected into office. Observing these policies, the voter updates her beliefs and votes for the party whose policy maximizes her expected utility.

Focusing on equilibria where the voter's strategy does not depend on parties' labels and that survive a refinement based on Grossman and Perry (1986), this paper obtains the following answers. First, a generic separating equilibrium exists if and only if one signal is weak in the sense that conditional on this signal the probability that the other party's signal is the same is less than one half.[3] Second, pooling at the voter's uninformed bliss point policy is (almost)[4] always an equilibrium. For quadratic utility, voter welfare in the separating equilibrium is larger than in the pooling equilibrium. Third, more precise signals increase the range of parameters for which the pooling equilibrium is the generically unique equilibrium. Therefore, as ε goes to zero, the pooling equilibrium is generically unique for any prior. Consequently, the better the parties' information the less likely will this information be transmitted in equilibrium. Policy convergence to the median voter's optimal policy absent additional information occurs whenever parties are well informed, and divergence may occur when parties' expertise is small.

[1] This is true for for representative or indirect democracies as well as for systems of direct democracy because even in the latter policy proposals are made by political parties or similarly well organized large interest groups.

[2] For example, the two states may be that another country has weapons of mass destruction or that it has not. The policy space is defense expenditure.

[3] Accordingly, a signal is called strong if this belief exceeds 1/2. The signal for state K is strong if the prior on K exceeds ε. So at least one signal is always strong.

[4] The qualification refers to an extension where the voter gets some additional information with positive probability. If this probability is sufficiently large, then there is no pooling equilibrium.

From a normative point of view, the voter may paradoxically be better off and make more informed decisions when parties' expertise, and the information overall in the economy, is relatively small.

Extensions reveal that these results do not crucially depend on parties being perfectly committed to implement the policies proposed once in office. The main results are also robust to the introduction of signals that are positively correlated conditional on the state.[5] Moreover, if the voter gets additional information with positive, but sufficiently small probability, pooling at the voter's uninformed bliss point policy is still an equilibrium, and generically unique as the error probability goes to zero.

Political parties in the present model are imperfectly informed parties, or experts, whose *messages* to the electorate take the form of policy proposals. Therefore, the paper is related to the cheap talk literature initiated by Crawford and Sobel (1982); see also Gilligan and Krehbiel (1989), Austen-Smith (1990) and Krishna and Morgan (2001b,a). Indeed, there is no cost inherent in messages, or policy proposals, to the parties other than the indirect cost that policy proposals affect the voter's decision. However, messages, or policy proposals, are not costless to the voter since they constrain her choice set. A second important difference from the typical cheap talk model is that here parties (or senders) are not assumed to be perfectly informed about the state of the world. Krishna and Morgan (2001b, p.769) acknowledge that "[i]n practice, the information of experts is neither perfect nor identical".[6] The present framework provides a tractable departure from the assumption of perfect and identical information. The results in the model with imperfectly informed parties contrast sharply with those when parties are perfectly informed as a separating equilibrium generically exists with $\varepsilon = 0$ but not for any $\varepsilon \to 0$ (Proposition 5).

Notwithstanding these differences, the arbitrage condition underlying the equilibrium construction in Crawford and Sobel (1982) and Krishna and Morgan (2001b) relies on the same concavity property of the utility function as the separating equilibrium here. Consequently, the informative equilibrium in the present paper, if it exists, rests on a very similar ground. In a subtle but important difference, however, here the arbitrage condition applies to the *receiver* (voter) rather than to the *sender(s)* (party or parties) as in Crawford and Sobel (1982) and

[5]Not surprisingly, this has qualitatively the same effect as decreasing ε and decreases the range for which a separating equilibrium exists.

[6]Perfect information is also assumed by Crawford and Sobel (1982) and Battaglini (2002). The assumption is relaxed by Battaglini (2004).

Krishna and Morgan (2001b).[7]

The paper also relates to the vast literature on signalling models. In a standard signalling game such as Spence (1973)'s education model, types differ with respect to their exogenously given costs of education, and this allows them to separate in equilibrium. Here, in contrast, there are no such exogenously given differences. Therefore, one might think that separation is not possible in equilibrium. But this is not generally true since equilibrium separation can occur. The reason is an expected cost induced by the probability of the other party's signal and the corresponding belief about its equilibrium action. Thus, the model exhibits endogenous signalling costs.[8]

Bernhardt, Duggan, and Squintani (2007) provide an alternative model that generates policy divergence and that rests on the assumption that parties are uncertain about the median voter's location. Schultz (1996) analyzes a political economy model of information transmission that is complementary to the present study as the parties in his model are mainly (i.e. lexicographically) policy motivated and perfectly informed about the state of the world.The present paper is also close to Heidhues and Lagerlöf (2003) and Laslier and Van Der Straeten (2004), who both analyze information transmission from parties to a voter in setup with two states and two policies. The difference between the two is that Laslier and Van Der Straeten (2004) assume the voter receives an independent signal whereas in Heidhues and Lagerlöf (2003) she does not. Therefore, the main model in this paper is closer to the one by Heidhues and Lagerlöf. In an extension of my model, where the voter learns the truth independently of the actions the parties take with positive probability, the result of Laslier and Van Der Straeten is corroborated, and complemented, by showing that the pooling equilibrium exists only if this probability is substantially less than one. Interestingly, in Heidhues and Lagerlöf (2003) a welfare superior (mixed strategy) equilibrium that involves some separation exists if and only if both signals are strong. Here, in contrast, the separating equilibrium exists if and only if one signal is weak. Callander (2008)'s model is related to the present one in that here and there parties' policy proposals serve the dual role of, potentially, conveying information to the voter

[7]There is also a mostly very recent literature on models where senders experience a cost of lying; see e.g. Banks (1990), Callander and Wilkie (2007), Kartik, Ottaviani, and Squintani (2007) and Kartik (2008). The difference between these models and the present one is that here messages are costless to parties but costly to the voter because they constrain her choice set.

[8]Since the parties who compete for a voter choose locations on a line, the paper also relates and contributes to the literature on Hotelling (1929) location games by adding the twist that locations may reveal or conceal information.

and of constituting the voter's choice set.[9]

The remainder of this paper is organized as follows. Section 2 introduces the model. Section 3 discusses preliminaries and derives perfect Bayesian equilibria (PBE). Equilibrium analysis is performed in Section 4. Section 5 presents a short welfare analysis. The model is extended in various ways in Section 6. Section 7 concludes. All the proofs are in the appendix. Appendix B shows, in turn, that the main results still hold when signals are positively correlated or of different precision.

2 The Model

Consider the following model of political competition with two parties that are better informed than the single voter. The policy space is $T \equiv [\underline{t}, \overline{t}] \subseteq \mathbb{R}$ and parties 1 and 2 compete by simultaneously announcing platforms $\tau_i \in T$ for $i = 1, 2$, which they are by assumption committed to implement if elected. One way to think of τ is as overall tax burden, zero meaning anarchy and one meaning totalitarianism, in which this case $T = [0, 1]$. On the other hand, if τ measures effort or government expenditure, then $T = [0, \infty)$.

Parties solely care for being in office. Parties and the voter are uncertain about which of two possible states $\omega \in \{A, B\}$ is true. Denote by $\alpha \in (0, 1)$ the common prior that state A is true, and denote by $u(\alpha, \tau)$ the expected utility of the voter when her belief is α and the implemented policy is τ. I assume that for any $\alpha \in (0, 1)$, $u(\alpha, \tau)$ has a unique maximizer $\tau(\alpha)$, which is monotone in α.[10] Without loss of generality, I assume that $\tau(\alpha)$ decreases in α, which corresponds to $u_{12} < 0$. Additionally, I assume that for any $\tau \in [\tau(1), \tau(\alpha))$ there is a $\tau' \in T$ with $\tau' > \tau(\alpha)$ such that $u(\alpha, \tau) = u(\alpha, \tau')$ and, analogously, that for any $\tau \in (\tau(\alpha), \tau(0)]$ there is a $\tau^0 \in T$ with $\tau^0 < \tau(\alpha)$ such that $u(\alpha, \tau) = u(\alpha, \tau^0)$.

For example, the utility function $u(\alpha, \tau) = (1 - \alpha) \ln(\tau) - \tau$ satisfies these properties, where $(1 - \alpha)$ may be thought of as the believed productivity of the tax (or treatment or effort) and $-\tau$ is the disutility associated with it. Another example that is a special case of the present

[9] Less closely but still related are political economy models with experimentation and/or try-and-error politics when decision makers face uncertainty about how the economy functions such as Laslier, Trannoy, and Van Der Straeten (2003), Majumdar and Mukand (2004), Berentsen, Bruegger, and Loertscher (2008), Strulovici (2008) and Callander (2009).

[10] Sufficient conditions for these properties are $u_2(\alpha, \underline{t}) > u_2(\alpha, \overline{t})$ and $u_{22}(\alpha, \tau) < 0$, and that the sign of $u_{12}(\alpha, \tau)$ is constant for all $\alpha \in (0, 1)$, where subscripts denote partial derivatives. These assumptions are reminiscent of, but different from those in supermodular games (see e.g. Vives, 2005). To see that they imply monotonicity of $\tau(\alpha)$ in α, observe that $u_2(\alpha, \tau(\alpha)) = 0$ by definition of $\tau(\alpha)$. Now if $u_{12} > 0$ holds, then $u_2(\overline{\alpha}, \tau(\alpha)) > 0$ for any $\overline{\alpha} > \alpha$. But $u_{22} < 0$ then implies that $\tau(\overline{\alpha}) > \tau(\alpha)$. The argument for the case $u_{12} < 0$ being symmetric, monotonicity follows.

model is the quadratic utility model often used in the cheap talk literature.[11] In state ω the voter's utility is $-(\omega-\tau)^2$ when the policy is τ. Letting $A=0$ and $B=1$ the expected utility function is thus $u(\alpha,\tau)=-\alpha\tau^2-(1-\alpha)(1-\tau)^2$. For both the quadratic and the logarithmic examples, $\tau(\alpha)=1-\alpha$.[12]

The idea that parties are better informed than the voter is captured in the following way. Each party i receives a private signal $s_i \in \{a,b\}$ indicating that state A or B has materialized before choosing its policy τ_i.[13] 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. All of this is common knowledge. Appendix B.1 replace the independence assumption with the alternative that signals are positively correlated, conditional on the state and shows that such correlation makes the range of the prior where pooling is the generically unique equilibrium outcome larger.[14]

When parties 1 and 2 choose policies τ_1 and τ_2, the voter's posterior belief that A is true is denoted $\mu(\tau_1,\tau_2)$. Some notation for the voter's belief that A is true under the hypothesis that she knew the signal of one or both parties is also useful. In slight abuse of notation let $\mu(a,a)$ and $\mu(b,b)$ be her belief that A is true if both parties have received the signal a and b, respectively. Analogously, let $\mu(a,b)=\mu(b,a)$ be this hypothetical belief when the parties receive divergent signals, and denote by $\mu(a,0)=\mu(0,a)$ and $\mu(b,0)=\mu(0,b)$ the belief that A is true under the hypothesis that the voter knows that one party has received the signal a or b, respectively, without knowing the other party's signal, which is denoted by 0. Due to the above assumptions about the signals, it is true that $\mu(a,b)=\alpha$, $\mu(a,a)=\frac{\alpha(1-\varepsilon)^2}{\alpha(1-\varepsilon)^2+(1-\alpha)\varepsilon^2}$ and $\mu(b,b)=\frac{\alpha\varepsilon^2}{\alpha\varepsilon^2+(1-\alpha)(1-\varepsilon)^2}$, and $\mu(a,0)=\frac{\alpha(1-\varepsilon)}{\alpha(1-\varepsilon)+(1-\alpha)\varepsilon}$ and $\mu(b,0)=\frac{\alpha\varepsilon}{\alpha\varepsilon+(1-\alpha)(1-\varepsilon)}$. These beliefs satisfy $\mu(a,a)>\mu(a,0)>\alpha>\mu(b,0)>\mu(b,b)$.

[11] See e.g. Crawford and Sobel (1982) or Krishna and Morgan (2001b).

[12] Under the standard assumptions that all voters update in the same way and that voter preferences are singled peaked, this can be viewed as shortcut to a model with many voters who differ with respect to some characteristic such as income, where the median voter would be the decisive voter. Let N_V be the number of voters and assume that N_V is odd. Modify the utility function to be $u(\alpha,\tau,\theta_v)$, where θ_v is voter v's type for $v=1,..,N_V$. Label voters in increasing order, so that $\theta_1<..<\theta_{N_V}$. Assuming, as above, $u_2>0$ and $u_{22}<0$, the function has a unique maximizer, denoted $\tau^*(\alpha,\theta_v)$. The sign of $d\tau^*/d\alpha$ and $d\tau^*/d\theta_v$ is the same as that of u_{12} and u_{23}, respectively. So $\tau^*(\alpha,\theta_v)$ will be monotone in α and θ_v if, as is assumed now, u_{12} and u_{23} have constant signs. Consequently, for any belief α, the bliss point policies $\tau^*(\alpha,\theta_1),..,\tau^*(\alpha,\theta_{N_V})$ can be ordered monotonically. Without loss of generality assume $\tau^*(\alpha,\theta_1)<...<\tau^*(\alpha,\theta_{N_V})$. Under the standard assumption in the literature (see e.g. Callander, 2008) that all voters update in the same manner, the model reduces to the median voter model analyzed here.

[13] Throughout I denote signals 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 $K\in\{A,B\}$.

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

A pure strategy for party i is a policy τ_i that depends on the private signal s_i. The policy party i plays upon signal $s_i = k$ is denoted τ_i^k with $i = 1, 2$ and $k = a, b$. A natural equilibrium concept is Perfect Bayesian Equilibrium (PBE), and I restrict attention to PBE where the voter's strategy does not depend on the labelling of the parties.[15] That is, denoting by $\gamma(\tau', \tau'')$ the probability that the voter elects party 1 when party 1 plays τ' while 2 plays τ'', the symmetry assumption is that $\gamma(\tau'', \tau') = 1 - \gamma(\tau', \tau'')$. Observe that for $\tau' = \tau''$ this implies $\gamma(\tau', \tau'') = 1/2$. Without this assumption, it is easy to construct PBE where the voter elects party 1 with probability one independent of the policy proposals and where both parties play separating strategies. These equilibria are fully revealing but arguably not very compelling as they do not reflect any notion of political competition. For analytical tractability I also focus on equilibria where the parties play pure strategies.

The model is fairly general and permits a variety of interpretations. First, as in Berentsen, Bruegger, and Loertscher (2008) uncertainty may pertain to the shape of the production function for a public good. Second, uncertainty may concern the military or terroristic threat level to which the country is exposed and accordingly to the optimal amount of expenditure on national security. Third, when policy consists of redistributing income via a linear tax and the (median) voter's income is smaller than the average income, the shape of the Laffer curve may be uncertain. Observe also that this setup contains Schultz (1996)'s model with respect to policy space and states as a special case, where ω is the cost of production of a public good and τ the expenditure for the public good, so that $u(\alpha, \tau) = F(\tau) - \alpha A - (1 - \alpha)B$, where the public good's production function $F(.)$ satisfies Inada conditions and where $A > B > 0$.

Throughout most of the paper, I focus on the positive question how much information is transmitted in equilibrium. The question in which equilibrium the voter's expected utility is largest is tackled in Section 5 for the case with quadratic utility. The reason is that the fairly general form of the model permits a general answer to the former, but not to the latter question.

3 Preliminaries

This section first distinguishes between weak and strong signals, which are key to the equilibrium behavior, and then shows that there are different policies that make the voter indifferent. It also shows briefly that there are both many pooling PBE and many separating PBE.

[15] See also Heidhues and Lagerlöf (2003).

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

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

The lemma implies that both signals are strong if and only if $\alpha \in (\varepsilon, 1 - \varepsilon)$. Note also that $\mu_i(\omega = K \mid s_i = k) > 1/2$ if and only if k is a strong signal, where $\mu_i(\omega = K \mid s_i = k)$ is the probability that the true state is K when the signal is k. To ease the notation, I sometimes denote $\mu(k \mid k) \equiv \mu_i(s_j = k \mid s_i = k)$.

(In)Different Policies The assumptions I made above imply that for every $\tau_a < \tau(\alpha)$ close enough to $\tau(\alpha)$ there is a $\tau_b > \tau(\alpha)$ such that

$$u(\alpha, \tau_a) = u(\alpha, \tau_b). \tag{1}$$

As $u(\alpha, \tau)$ is strictly increasing in τ for $\tau < \tau(\alpha)$ and decreasing in τ for $\tau > \tau(\alpha)$, it follows that for all pairs (τ_a, τ_b) satisfying (1) $d\tau_a/d\tau_b < 0$ holds.

Perfect Bayesian Equilibria (PBE) Unsurprisingly the present game exhibits many PBE. I first characterize some general properties of PBE and then show that, among other types of PBE, there is a continuum of pooling PBE and a continuum of separating PBE.

Normalize the payoff of winning the election to one and the payoff of losing to zero and denote by $E_i[\tau_i^k \mid s_i = l]$ the expected payoff to i when playing τ_i^k after signal $l \in \{a, b\}$, where τ_i^k is the action equilibrium prescribes i to play upon $s_i = k$. Now by the definition of an equilibrium, the incentive constraint

$$E_i[\tau_i^k \mid s_i = k] \geq E_i[\tau_i^l \mid s_i = k] \tag{2}$$

has to hold in any PBE for else i would be better off playing τ_i^l after signal k than the prescribed action τ_i^k.

Next I state and prove an implication of the assumption that the voter's strategy does not depend on the parties' labels.

Lemma 2 *In any PBE where the voter's strategy does not depend on the parties' labels,*

$$E_i[\tau_i^k \mid s_i = k] = E_j[\tau_j^k \mid s_j = k] = 1/2 \tag{3}$$

for $k \in \{a, b\}$.

Lemma 2 may appear obvious, but it is not insofar as it does not extend to models where the voter receives some independent information with positive probability (see Section 6.4 below). I now establish an important property of separating PBE. A PBE is called separating if $\tau_i^a \neq \tau_i^b$ for both $i = 1, 2$.

Lemma 3 *$\tau_i^a = \tau_a$ and $\tau_i^b = \tau_b$ for $i = 1, 2$ are strategies of a separating PBE only if*

$$u(\alpha, \tau_a) = u(\alpha, \tau_b). \tag{4}$$

Condition (4) is very similar to the arbitrage condition underlying the equilibrium construction in Crawford and Sobel (1982) and Krishna and Morgan (2001b). A subtle but important difference is that here it applies to the voter rather than the party(s). Its intuition is clear: Suppose (4) were violated e.g. because $u(\alpha, \tau_a) > u(\alpha, \tau_b)$. Then playing τ_a independently of the signal would be a profitable deviation because it guarantees election with probability 1 if the other party plays τ_b and with probability 1/2 if the other one plays τ_a, so that overall the probability of being elected is greater than 1/2. So for parties to be willing to separate, $u(\alpha, \tau_a) = u(\alpha, \tau_b)$ has to hold. Lemma 3 is important for the results that follow which show that (4) imposes a condition that can be made to hold if one signal is weak but generically not if both signals are strong.

Proposition 1 *There is a continuum of pooling PBE, where $\tau_i^a = \tau_i^b = \tau$ for $i = 1, 2$. There is also a continuum of separating PBE, where $\tau_i^a = \tau_a$ and $\tau_i^b = \tau_b$ for $i = 1, 2$ and where τ_a and τ_b satisfy (1).*

The separating PBE arise here because of the continuous strategy space and the concavity properties of $u(\alpha, \tau)$. Proposition 1 is not a complete description of all types of PBE in the model. For example, there are also pooling PBE where one party plays τ_a and the other one τ_b, where τ_a and τ_b satisfy (1), and there may be hybrid PBE where one party does not reveal its signal while the other one does. This raises the question whether some of these PBE are more plausible than others, which is what I address in the next section. Before doing so, it is useful to state and prove the following lemma. Denote by $\tau_a^* \equiv \tau(\mu(a, a))$ and $\tau_b^* \equiv \tau(\mu(b, b))$ the voter's preferred policies if both signals are a and b, respectively.

Lemma 4 *When both signals are strong, there are no separating PBE where $\tau_a, \tau_b \notin [\tau_a^*, \tau_b^*]$.*

Intuitive Criterion and PSE for Games with Multiple Senders In the present model each party i can be of two types, type a after receiving $s_i = a$ or type b upon $s_i = b$. Despite there being only two types, the spirit of the probably most widely used refinement, the intuitive criterion proposed by Cho and Kreps (1987), has no bite in the current context. To see this, consider any PBE described in Proposition 1. Though some of these rest on beliefs that are not very plausible, no equilibrium policy is dominated by equilibrium payoffs because on equilibrium every party is elected with probability 1/2 and could hence potentially benefit if, after an off equilibrium move, it were elected with a larger probability.[16] Cho and Kreps (1987) only require to put zero probability on types who could not possibly benefit from the off equilibrium move. It puts no restriction on the beliefs assigned to types who could potentially benefit.

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 Cho and Kreps (1987) in that it requires the voter to put zero probability on types for whom an off equilibrium move is dominated by the payoff they get in equilibrium. In addition to the Cho-Kreps criterion, PSE requires credible updating, i.e. to assign prior preserving posterior beliefs to all types who could potentially benefit from the observed deviation. Specifically, in the present model the PSE algorithm works as follows:[17]

First, let $\mathbf{t} = (\tau_1^a, \tau_1^b, \tau_2^a, \tau_2^b)$ be the set of policies that the voter expects to observe with positive probability in a given PBE and assume that he observes a deviation by party 2 to some $\tau_2 \notin \mathbf{t}$. Then ask which type(s) - a, b or both - could have benefitted from playing τ_2 rather than the equilibrium policy. 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 p_A and p_B denote the priors, then the posterior beliefs $\hat{p}_A$ and $\hat{p}_B$ must satisfy $\frac{\hat{p}_A}{\hat{p}_B} = \frac{p_A}{p_B}$. Third, given these updated beliefs the voter makes the choice that maximizes her expected utility. Fourth, if given this choice by the voter, both types of party 2 are no better off than when playing the equilibrium policy, the deviation has not paid off. If both types benefit, the PBE fails the PSE test.[18] Last, for $\mathbf{t}$ to be a set of PSE polices there must be no deviation

[16] In contrast to the beer-quiche game of Cho and Kreps (1987) there is no PBE where a type gets his first-best (quiche and no fight or beer and no fight) in the present model.

[17] See Grossman and Perry (1986), Riley (2001) and Hörner and Sahuguet (2007).

[18] If only one type benefits, then go through the same exercise again, but this time by assigning probability

where at least one type of party 1 or 2 could benefit, assuming that the voter goes subsequently through steps 1 through 4 of this algorithm.

4 Equilibrium Analysis

In this section I focus on PSE. The first result is a recurring theme within this paper and in the broader literature on information transmission, namely that babbling is (almost) always an equilibrium that cannot be refined away.

Proposition 2 *For any $\alpha \in (0,1)$, there is a unique pooling PSE, whose outcome is $\tau(\alpha)$.*

Any deviation has two effects. It changes the menu of policies the voter can choose from and it affects the voter's posterior. PSE's power stems from the latter: Since any deviation can come from either type of the deviating party, all deviations from a pooling equilibrium are considered pooling. Therefore, if the PBE under consideration is pooling, the voter's posterior belief equals her prior both on and off equilibrium. Unless the equilibrium policy is $\tau(\alpha)$, either party can thus profitably deviate to $\tau(\alpha)$ instead of the prescribed equilibrium policy and win the election with probability one. Analogously, if the prescribed equilibrium policy is $\tau(\alpha)$ no party can profitably deviate to some $\tau \neq \tau(\alpha)$ since the voter's off equilibrium belief will be α, so that she prefers $\tau(\alpha)$ to any alternative.

Every PSE is a PBE since PSE only puts additional restrictions on the off equilibrium beliefs. Whatever beliefs are consistent with a PSE will therefore be consistent with a PBE. Consequently, a corollary to Lemma 3 is that $\tau_i^a = \tau_a$ and $\tau_i^b = \tau_b$ for $i = 1, 2$ are strategies in a separating PSE only if

$$u(\alpha, \tau_a) = u(\alpha, \tau_b). \tag{5}$$

In addition, PSE imposes the following restriction on separating equilibria:

Lemma 5 *If both signals are strong, $\tau_i^a = \tau_a$ and $\tau_i^a = \tau_b$ for $i = 1, 2$ are strategies of a separating PSE only if*

$$\tau_a = \tau_a^* \quad \textit{and} \quad \tau_b = \tau_b^*. \tag{6}$$

The lemma is key. It adds a second condition on equilibrium policies when both signals are strong that will generically not hold simultaneously with the condition (5). Figure 1 illustrates

one to the type who could have benefitted. The requirement is then that if probability is assigned only to one type, and if the voter takes her expected utility maximizing action given this belief, the type who is believed to have chosen this policy with probability zero has indeed no incentive to make this deviation.

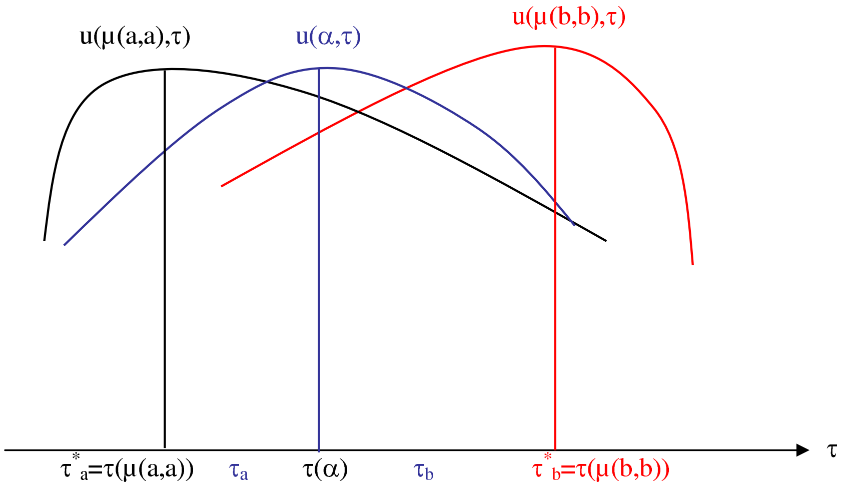

Figure 1: PSE Conditions: The generic case.

the generic case. The black and red curve depict, respectively, $u(\mu(a,a),\tau)$ and $u(\mu(b,b),\tau)$. The blue curve is $u(\alpha,\tau)$. Generically, $u(\alpha,\tau_a^*) \neq u(\alpha,\tau_b^*)$ will be the case. The non-generic case is illustrated in Figure 2.

The intuition is quite clear. Upon getting signal k a party's belief that the other party has got the same signal, $\mu(k \mid k)$, exceeds 1/2 if signal k is strong. If it plays on equilibrium, it wins with probability 1/2 (Lemma 2). Therefore, a deviation that guarantees victory if the other party has received the same signal will be profitable since its expected payoff is at least $\mu(k \mid k) > 1/2$. Now if both signals are strong, and if the deviation leads to election if and only if the other party has got the same signal k, then upon getting the signal $l \neq k$ the other party has indeed no incentive to play this deviation as it guarantees victory only with probability $1 - \mu(l \mid l) < 1/2$. Therefore, starting from a candidate separating equilibrium a deviator can credibly convey its signal by choosing such a policy, and benefit from the deviation. Unless, that is, the prescribed equilibrium policy upon signal k is τ_k^*. To see that the deviation leads to election iff the other party has received the same signal, recall from Lemma 4 that no separating PBE policies $\tau_a, \tau_b \notin [\tau_a^*, \tau_b^*]$ exist. Therefore, the policies of a candidate separating PSE will be between τ_a^* and τ_b^* and satisfy $u(\alpha,\tau_a) = u(\alpha,\tau_b)$. If only one type of a party benefits from the deviation to, say, τ_a^*, then upon observing (τ_a^*, τ_b) the voter will prefer τ_b since her belief will be α, which is consistent with the hypothesis that only one type benefits. On the other hand, if both types benefitted, then upon observing (τ_a^*, τ_b) her belief would be $\mu(b,0) < \alpha$, so that she prefers τ_b to τ_a^* a fortiori, which contradicts that both types benefit from the deviation.

Proposition 3 *If both signals are strong, the generically unique PSE outcome is* $\tau(\alpha)$.

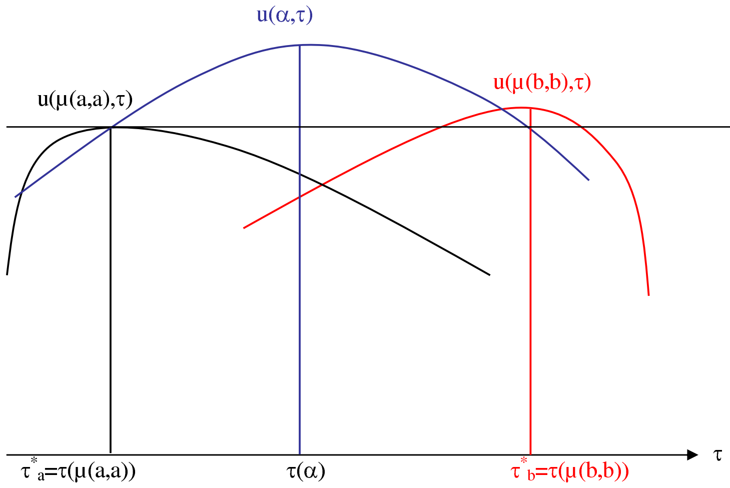

Figure 2: PSE Conditions: The non-generic case.

Non-generic separating PSE with both signals strong may exist if by a fluke it so happens that $u(\alpha, \tau_a^*) = u(\alpha, \tau_b^*)$.[19] On the other hand, if one signal is weak, there are generic separating PSE. Let $\hat{\tau}_k \equiv \arg\max_\tau u(\mu(k, 0), \tau)$ with $k \in \{a, b\}$. That is, $\hat{\tau}_k = \tau(\mu(k, 0))$. Observe that $\hat{\tau}_k$ maximizes the voter's expected utility if she knows or correctly infers that one party has received the signal k while the other party's signal is not known.

Proposition 4 *If one signal is weak, then there are separating PSE that are generic. If b is the weak signal, party $i = 1, 2$ sets $\hat{\tau}_a$ if $s_i = a$ and τ_b if $s_i = b$, where $u(\alpha, \hat{\tau}_a) = u(\alpha, \tau_b)$. If a is the weak signal, party $i = 1, 2$ sets $\hat{\tau}_b$ if $s_i = b$ and τ_a if $s_i = a$, where $u(\alpha, \tau_a) = u(\alpha, \hat{\tau}_b)$.*

Thus, despite the fact that parties are purely office motivated, the model exhibits generic separating PSE. Their construction relies first, as any separating PBE, on an indifference condition of the voter in case the parties disagree and second on the fact that upon receiving the weak signal k a party has no incentive to deviate to τ_k^* if the equilibrium prescribes playing $\tau_k \neq \tau_k^*$. This is so because upon getting the weak signal the party believes with probability less than 1/2 that the other party has received the same signal. Therefore, it is not willing to take the gamble of defeating the opponent only in the event it has got the same signal, which is not sufficiently likely. In contrast, upon getting a strong signal both parties would have incentives to take this gamble. Therefore, the equilibrium must prescribe to play $\hat{\tau}_l$ upon $s_i = l$ if l is the strong signal.

[19] In addition, there must be no $\tau \in (\tau_a^*, \tau_b^*)$ such that $u(\mu(a, 0), \tau) > u(\mu(a, 0), \tau_a^*)$ and $u(\mu(b, 0), \tau) > u(\mu(b, 0), \tau_b^*)$. The conditions under which this does or does not hold are somewhat obscure.

Despite the fact that ex ante parties are identical and are, per se, willing to propose any policy if that increases the chances of being elected, the model exhibits endogenous signalling costs. This contrasts with standard signalling 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 equilibrium. Separation can occur here because some deviations become too costly in equilibrium, given a party's probability assessment about the other party's signal and hence action.

A separating PSE exists if the signals' error probability is larger than the prior on one of the two states. An alternative and evidently equivalent way of expressing this is that a separating PSE exists if the prior is large, or extreme. So, according to this model, Obama's dissent from supporting a war on Iraq in 2002 may have been an equilibrium behavior exactly because the common prior that Saddam Hussein's regime was a threat to the U.S. was extreme one year after the terroristic attacks on New York and Washington.

Proposition 3 has the following corollary:

Corollary 1 *As signals become arbitrarily precise (i.e. as $\varepsilon \to 0$), the generically unique PSE outcome is the outcome of the pooling PSE, where $\tau_i^k = \tau(\alpha)$ for $i = 1, 2$ and both signals.*

A few farther remarks are in order. The result that in the (generically) unique PSE outcome no information is transmitted as parties' signals become perfectly accurate may appear counterintuitive. However, the reason for this is not that the voter trusts the parties less as their signals become more accurate. Rather, parties trust their source of information too much for a separating PSE to be supported: Since both signals are strong, party i believes with probability larger than 1/2 that party j has received the same signal upon receiving $s_i = a$ or $s_i = b$. In order for i not to deviate from the prescribed policy τ_a (or τ_b) it must be the case that $\tau_a = \tau_a^*$ (and $\tau_b = \tau_b^*$). But generically, $u(\alpha, \tau_a^*) \neq u(\alpha, \tau_b^*)$ will be the case. As ε goes to zero the interval $(\varepsilon, 1 - \varepsilon)$ coincides with $(0, 1)$. Consequently, as signals become perfectly accurate, both signals are strong for any interior prior. Thus, the generically unique PSE outcome is the pooling outcome $\tau(\alpha)$.

It is also noteworthy that the model exhibits a discontinuity at $\varepsilon = 0$.[20] That is:

Proposition 5 *A generic separating PSE exists for $\varepsilon = 0$.*

[20] This observation is due to Yuelan Chen.

A formal proof is skipped because the result is easily understood. Observe first that for $\varepsilon = 0$ both signals are perfectly correlated, i.e. $\mu(k \mid k) = 1$ for $k \in \{a, b\}$ and any $\alpha \in (0, 1)$. So as before it has to be the case in the separating PSE that $\tau_k = \tau_k^*$. However, because signals are perfectly correlated the observation where one party plays τ_a and the other one τ_b is now off equilibrium, with the subtle and important twist that the voter does not know which party has deviated. To deter parties from deviating, it must be the case that the voter is indifferent between τ_a and τ_b. But because this observation is off equilibrium and because PSE does not restrict the beliefs about which *player* deviated, the voter is free to choose her belief to be $\tilde{\mu}$, which is such that $u(\tilde{\mu}, \tau_a^*) = u(\tilde{\mu}, \tau_b^*)$.[21] Therefore, $\tau_a = \tau_a^*$ and $\tau_b = \tau_b^*$ are now the only restrictions on equilibrium policies. Consequently, a generic separating PSE exists in the, rather non-generic, case with $\varepsilon = 0$. This suggests that information transmission models where senders are perfectly informed may be a limit case that is not robust to some small noise in the senders' information.

5 Welfare Analysis

An interesting question is whether expected voter welfare is indeed larger in the separating PSE, as one would naturally conjecture. The expected welfare of the voter in the pooling PSE is $W^{pool} = u(\alpha, \tau(\alpha))$. Absent an outside option such as a status quo policy, voter welfare is the appropriate welfare measure as one of the two parties is chosen with probability one. Her expected welfare in a separating PSE with equilibrium policies τ_a and τ_b is[22]

$$W^{sep} = [u(1, \tau_a)\alpha + u(0, \tau_b)(1 - \alpha)](1 - \varepsilon) + [u(0, \tau_a)(1 - \alpha) + u(1, \tau_b)\alpha]\varepsilon. \qquad (7)$$

Observe that one effect of a separating equilibrium is that the informational content of one of the two signals is lost. In addition, separating PSE also imposes constraints of the proposed policies, as already observed.

How W^{pool} and W^{sep} compare is hard to say in general as the comparison depends both on the intricate properties of the utility function $u(.)$ and on the signal technology. To simplify, assume therefore that utility is quadratic with $u(\alpha, \tau) = -\alpha\tau^2 - (1 - \alpha)(1 - \tau)^2$. Since $\tau(\alpha) = 1 - \alpha$, $W^{pool} = -\alpha(1 - \alpha)$. Without loss of generality, assume that signal a is strong and signal b is weak, so that a separating PSE exists for all $\alpha > 1 - \varepsilon$. The equilibrium policies are

[21] It can easily be shown that such beliefs always exist.

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

$\tau_a = \tau(\mu(a,0))$ and the $\tau_b \neq \tau_a$ that solves $u(\alpha, \tau(\mu(a,0))) = u(\alpha, \tau_b)$, yielding $\tau_a = 1 - \mu(a,0)$ and $\tau_b = 1 - 2\alpha + \mu(a,0)$. So the ex ante expected welfare in the separating PSE is

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

Tedious algebra reveals that $W^{sep} > W^{pool}$ for all $\alpha > \frac{1-4\varepsilon}{4(1-2\varepsilon)}$, which is strictly less than $1-\varepsilon$ for all $\varepsilon < 1/2$. Thus, whenever a generic separating PSE exists, it generate higher welfare than the pooling PSE. Among other things, this implies that in the neighborhoods of $\alpha = 1-\varepsilon$ and $\alpha = \varepsilon$ the voter would actually prefer parties to be less well informed, i.e. to have somewhat larger ε because that induces them to play the separating PSE (assuming that the welfare superior equilibrium is played in case of multiple equilibria).

Information Sharing In light of the above an important question is what can be done to increase the amount of information transmitted in equilibrium from the parties to the receiver. To that end, assume now both parties are committed to share the information that is available to them, so that party 1 knows party 2's signal and vice versa. Denote by $I \in \{(a,a),(b,b),(a,b)\}$ the shared information available to both parties and, in slight abuse of notation, by $\mu(I)$ the belief that A is true conditional on information I. Then:

Proposition 6 *The following is a PSE (leaving the voter's belief unspecified): Upon information I, player i plays $\tau_i^I = \arg\max_\tau u(\mu(I), \tau)$. On equilibrium, the two parties always propose identical policies. If it is clear off equilibrium which party deviated, the voter elects the one who did not deviate with probability 1. Otherwise, i.e. if both propose different policies each of which would be played with positive probability in equilibrium, she randomizes uniformly between the two parties.*

A formal proof is skipped as the intuition is very much the same as for the PSE with $\varepsilon = 0$ (Proposition 5). There is no profitable deviation as the information is still contained in the policy by the non-deviating player and this policy maximizes voter welfare conditional on the available information. Like in the pooling PSE of the model with independent information parties still always propose identical policies. Importantly, however, the policies they propose now vary with the available information and therefore, probabilistically, with the state.

6 Extensions

Various extensions of the basic model are worthwhile being considered. This section first analyzes sequential moves by parties. Second, a status quo policy is introduced the voter can choose if not satisfied with the parties' proposals. Interestingly, a status quo policy is very similar to having more than two parties, which is the third extension briefly considered. Further extensions allow, in turn, the voter to receive independent feedback, parties to be imperfectly committed to implement the policies they propose and introduce more than two states. Appendix B.2 shows that the conclusion that pooling is the unique PSE outcome when signals become very precise does not depend on the assumption that parties' signals are of exactly the same quality, provided the voter adheres to a symmetric strategy.

6.1 Sequential Moves

The alternative of sequential moves by parties/experts is considered e.g. by Krishna and Morgan (2001b,a) and Pesendorfer and Wolinsky (2003). This relates to the previous discussion about endogenous signalling costs. These disappear with sequential moves, as will be shown shortly. Throughout this subsection assume party 1 moves first and party 2 second. It is easy to see that:

Proposition 7 *$\tau_i^k = \tau(\alpha)$ for $i = 1, 2$ and $k = a, b$ is a PSE outcome with sequential moves.*

Given that one party pools at $\tau(\alpha)$ any deviation by the other one from the prescribed policy $\tau(\alpha)$ would be pooling and would hence be defeated. Another question is whether there are fully separating PSE with sequential moves, i.e. PSE where both players' signals are revealed on the equilibrium path. The answer is affirmative:

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

In contrast to the model with simultaneous moves any deviation by any player would be pooling with sequential moves. Deviations from the actions prescribed by the separating equilibrium by player 2 can be preempted by 1 by playing $\hat{\tau}_k$ upon signal $k = a, b$, which cannot be defeated by a pooling deviant. Deviations by player 1 can be countered by player 2 by playing $\tau(\alpha)$, which is the optimal policy given pooling behavior.

A consequence of the fact that with sequential moves all deviations are pooling is that, unlike with simultaneous moves, a separating equilibrium policy does not have to satisfy τ_k^*

if signals are strong. The reason is that the endogenous signalling costs disappear for player 2 who, assuming equilibrium play by 1, is now certain about 1's signal. Therefore, he cannot credibly reveal his signal by off equilibrium behavior.

The difference to Krishna and Morgan (2001b) is interesting and quite striking. They find that there are fully separating PBE with simultaneous moves but not with sequential moves. For the case where both signals are strong (which corresponds to the case analyzed by Krishna and Morgan who focus on perfectly informed parties), the opposite result obtains in the present paper.

6.2 Status Quo Policy

As in (parts of) Krishna and Morgan (2001a) and Gilligan and Krehbiel (1989) assume there is a status quo policy τ^O the voter can threaten to choose. This status quo policy can also be interpreted as the preferred and state independent policy proposed by a third party, which would for instance be the case if the third party is ideological and does either not receive a signal or is known to ignore it.

Proposition 9 $\tau(\alpha)$ *is still the unique pooling PSE outcome.*

Another, and perhaps more interesting or relevant question is whether there are separating PSE once the voter has an outside option.

Proposition 10 *If both signals are strong and* $\tau^O \notin [\tau_a^*, \tau_b^*]$*, then there is generically no separating PSE.*

6.3 More than Two Parties

Assume now that there are $N > 2$ parties who simultaneously propose policies after receiving independent private signals.[23] The first result is very general and holds even when different parties have signals of different precision.

Proposition 11 $\tau_i^k = \tau(\alpha)$ *is a PSE outcome for* $k = a, b$*,* $i = 1, .., N$*, and it is the unique pooling PSE outcome.*

It is noteworthy that the result does depend on the voter's strategy being symmetric. Suppose, for whatever reason, that party i is chosen with probability $\gamma_i < 1/N$ when proposing $\tau(\alpha)$.

[23]It is well known that in location games with more than two strategically locating players and many voters (or consumers) pure strategy equilibria need not exist. Therefore, this extension somewhat critically depends on the assumption that there is only one voter.

Both of i's types would want to deviate if upon deviation i were elected with probability larger than γ_i. Thus, the deviation would be pooling.

Proposition 11 motivates to see whether there are other generic PSE for $N > 2$, which is the content of the remainder of this subsection. For that purpose, assume from now on that all parties have symmetric signals with error probability $\varepsilon < 1/2$. An albeit only partial answer is given in the following for $N = 3$.

Now the on equilibrium probability of winning upon signal a is $\mu(a \mid a)^2/3 + 2\mu(a \mid a)(1 - \mu(a \mid a))/2$, assuming that $u(\tau_a, \mu(2,1)) > u(\tau_b, \mu(2,1))$, where in slight abuse of notation $\mu(3-x, x)$ is the belief that $\omega = A$ given that $3 - x$ parties have got the signal a and x have got the signal b with $x = 0, 1, 2, 3$.[24] The off equilibrium payoff of going to the middle is $2(1 - \mu(a \mid a))\mu(a \mid a)$. Observe, that upon signal a deviation pays off if (and only if) $\mu(a \mid a)^2/3 + 2\mu(a \mid a)(1 - \mu(a \mid a))/2 < 2(1 - \mu(a \mid a))\mu(a \mid a) \Leftrightarrow \mu(a \mid a) < 3/4$. Since under the maintained hypothesis the equilibrium payoff upon $s_i = b$ is only $\mu(b \mid b)^2/3$, the deviation of going to the middle pays off whenever $\mu(b \mid b)^2/3 < 2(1 - \mu(b \mid b))\mu(b \mid b) \Leftrightarrow \mu(b \mid b) < 6/7$. Thus, a sufficient condition for there to be no separating PSE is $\mu(k \mid k) < 3/4$. Now $\mu(a \mid a) < 3/4 \Leftrightarrow \alpha < \frac{\varepsilon(3-4\varepsilon)}{1-2\varepsilon}$ and $\mu(b \mid b) < 3/4 \Leftrightarrow \alpha > \frac{\varepsilon(1-5\varepsilon+5\varepsilon^2)}{1-2\varepsilon}$. Since $\frac{\varepsilon(1-5\varepsilon+5\varepsilon^2)}{1-2\varepsilon} < \frac{\varepsilon(3-4\varepsilon)}{1-2\varepsilon} \Leftrightarrow \varepsilon > \frac{1}{2} - \frac{1}{\sqrt{2}} \approx 0.14$, it follows:

Proposition 12 *For $\varepsilon \in \left(\frac{1}{2} - \frac{1}{\sqrt{2}}, \frac{1}{2}\right)$ there is no separating PSE with $N = 3$ for any $\alpha \in \left(\frac{\varepsilon(1-5\varepsilon+5\varepsilon^2)}{1-2\varepsilon}, \frac{\varepsilon(3-4\varepsilon)}{1-2\varepsilon}\right)$.*

Preliminary results indicate that otherwise separating PSE with $N = 3$ exist quite generally.

6.4 Voter Gets Feedback

Assume that there are two parties whose signals are independent, conditional on the state and one voter. In contrast to the analysis so far, with positive probability the voter now gets feedback about the state that is independent of, and in addition to, the information conveyed in the parties' proposals. The voter's feedback is modelled in two different, yet largely equivalent ways.

Feedback rule I: With probability λ the voter learns the truth irrespective of what the parties do. With probability $1 - \lambda$ the voter does not get any additional information.

[24] Notice that this is without loss of generality since $u(\tau_a, \mu(2,1)) \neq u(\tau_b, \mu(2,1))$ or/and $u(\tau_a, \mu(1,2)) \neq u(\tau_b, \mu(1,2))$ will hold. Therefore, upon one signal the expected equilibrium payoff will be larger, which is assumed to be signal a.

This is a particularly simple way of introducing feedback. However, it is somewhat at odds with the signal structure of the experts as it implies for λ close to 1 that the voter has more expertise than the parties. A rule more in line with the spirit of rest of the paper is:

Feedback rule II: With probability λ the voter gets a signal s_V that conditional on the state is independent of the signals of the parties. The signal takes on the value $s_V = k$ in state K with probability $1 - \varepsilon$ and the value $j \neq k$ with probability ε with $K = A, B$ and $k = a, b$. With probability $1 - \lambda$ the voter gets no additional information.

Under either feedback rule, a voter is said to be, or remain, ignorant, if she does not get additional information (i.e. in the event that occurs with probability $1 - \lambda$). The assumption that the signal has the same quality (i.e. the same ε) as the parties' is made merely for simplicity. What really matters is that the voter remains ignorant, beyond whatever she can learn from the parties' behavior, with positive probability. Typically the interest will be on the case where λ is small.[25]

Two questions of interest are whether the possibility of feedback destroys the pooling PSE or allows for separating PSE. A partial answer to the latter is:

Lemma 6 *For either feedback rule, there is no separating PBE where $\tau_k^i = \tau_k$ for $i = 1, 2$ and $k = a, b$ with $\tau_a < \tau_b$ and $u(\alpha, \tau_a) \neq u(\alpha, \tau_b)$ if $\lambda \leq 1/2$.*

Consider a candidate separating equilibrium that prescribes $\tau_k \neq \tau_k^*$ upon signal $k \in \{a, b\}$ when both signals are strong. In the model without feedback, a party of type k who deviates to τ_k^* wins independently of the true state of world if the other party plays τ_k. In the model with feedback, this is not necessarily true as the voter may learn what the truth is and therefore may prefer τ_k to τ_k^* in that instant. Therefore, the conditions under which deviations to τ_k^* upon signal k pay off become slightly more stringent though not completely different if the voter receives feedback. Before tackling the problem of separating PSE with feedback, it is useful to state the following result, which is the analog under feedback rule I to Lemma 4:

Lemma 7 *Assume $\alpha \in (\varepsilon, 1 - \varepsilon)$. Under feedback rule I there exists $\overline{\lambda}(\varepsilon, \alpha) > 0$ satisfying $\lim_{\varepsilon \to 0} \overline{\lambda}(\varepsilon, \alpha) = 1/2$ such that there is no separating PBE with equilibrium policies $\tau_a, \tau_b \notin [\tau_a^*, \tau_b^*]$.*

[25] An alternative modelling assumption, which is made by Laslier and Van Der Straeten (2004), is to set $\lambda = 1$, so that the voter becomes essentially as good of an expert as the parties even under feedback rule II. Which assumption is more appropriate is an empirical matter. It seems that typically voters are indeed less well informed than political parties, which corresponds to $\lambda < 1$.

The lemma allows us to state:

Lemma 8 *For $\lambda < \overline{\lambda}(\varepsilon, \alpha)$ and $\alpha \in (\varepsilon, 1-\varepsilon)$ $\tau_a = \tau_a^*$ and $\tau_b = \tau_b^*$ are necessary conditions for a separating PSE under feedback rule I.*

Lemmas 6 and 8 imply:

Proposition 13 *For $\lambda < \overline{\lambda}(\varepsilon, \alpha)$, there is generically no separating PSE under feedback rule I if both signals are strong.*

The equivalent statements to Proposition 13 and Lemmas 7 and 8 for feedback rule II remain to be worked out. A corollary to Proposition 13 and Lemma 8 is:

Corollary 2 *As $\varepsilon \to 0$ there is generically no separating PSE under feedback rule I for any $\lambda < 1/2$.*

An answer to the first question raised above is:

Proposition 14 *Under feedback rule I, there is a unique pooling PSE, whose outcome is $\tau(\alpha)$, if and only if $\lambda < \min\left\{\frac{1}{2\mu_i(\omega=A|s_i=a)}, \frac{1}{2\mu_i(\omega=B|s_i=b)}\right\}$. Under feedback rule II, there is a $\hat{\lambda}(\alpha, \varepsilon) \in (1/2, 1)$ such that there is a unique pooling PSE, whose outcome is $\tau(\alpha)$, if and only if $\lambda < \hat{\lambda}(\alpha, \varepsilon)$.*

Observe that the bounds in Proposition 14 always exceed 1/2. Notice also that the proposition relates to, and is consistent with, Laslier and Van Der Straeten (2004), who find that there is no pooling equilibrium satisfying a stability criterion if the voter gets a signal with probability one.

An open issue is what happens under, say, feedback rule I, if λ is so large that no pooling PSE exists. In the limit when $\lambda = 1$, there is a separating PSE whenever both signals are strong, where upon signal $s_i = a$ parties play $\tau(1)$ and upon signal $s_i = b$ they play $\tau(0)$. No indifference of the form $u(\alpha, \tau(1)) = u(\alpha, \tau(0))$ is required. By continuity this separating PSE will still exist for λ's somewhat smaller than 1 for most (α, ε)'s such that both signals are strong. However, a complete answer for all values of λ, α and ε remains to be worked out.

6.5 Imperfect Commitment

So far parties have been assumed to somewhat stubbornly implement the policies with which they campaign. This assumption is now relaxed: Each party can a priori be one of two types or qualities $q \in \{C, O\}$, which matter in the following sense. A party of quality C is *committed* to implement the policy it proposed during the campaign. A party of quality O is *opportunistic* and willing to deviate from the policy it announced in the election campaign. For simplicity I assume that if of quality O either party implements the policy that maximizes the voter's expected utility given the information the party has once it is in office.[26] Notice that parties of quality C are the parties studied hitherto. A party's quality is its private information, and the probability that $q = C$ is $\beta > 0$, so that the probability of being $q = O$ is $1 - \beta$.[27] Quality is independent across parties and independent of the signals. Regardless of its quality,[28] each party still cares exclusively about winning the election. All the other assumptions are as in the baseline model ($\lambda = 0$).

Lemma 9 *Party quality $q \in \{O, C\}$ cannot be signalled in equilibrium.*

Assume now that upon being in office a party learns the other party's signal regardless of the action chosen. This assumption is made for simplicity. An alternative that remains to be explored is that a party in office only knows its own signal for sure and, depending on the equilibrium, may or may not infer the other party's signal. Given this simplifying assumption, upon being in office both parties implement the same policy with probability β. Though the voter does not necessarily know what this policy is, the parties do not differ in that respect if they are of quality O, and the voter need not worry since the policy will maximize her expected utility. So parties and their proposals only matter if they are of quality C. Therefore, every PSE is completely driven by C-types. But this is the model already studied above. Consequently:

Proposition 15 *PSE policies do not vary with $\beta \in (0, 1]$.*

The proposition implies that whatever is a PSE outcome in the model with $\beta = 1$ is a PSE outcome for any $\beta \in (0, 1)$. Thus, whether or not parties are fully committed to implement the policies they propose does not affect equilibrium behavior. The limit case with $\beta = 0$, where both parties are opportunistic, is excluded from the Proposition because it gives raise to a continuum of PSE outcomes: As both parties implement the same policy with probability 1 upon being elected, their policy proposals do not matter. Therefore, equilibrium is indeterminate.

With $\beta > 0$, C-types imitate O-types in any pure strategy PBE. Therefore, the presence of O-types has no impact on the policies proposed in equilibrium. Still, their presence matters as

[26]Therefore, the labels "committed" and "opportunistic" are slight misnomers.

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

[28]Notice that after receiving its signal s_i and before choosing its policy τ_i party i can now be one of four types in $\{C, O\} \times \{a, b\}$.

it obviously improves the voter's expected welfare because it improves the expected quality of policies that are ultimately implemented.

6.6 More than Two States

Suppose now that there are $S \geq 2$ states ω, but otherwise the conditions of the baseline model hold (two parties, independent signals, no voter feedback). Denote by α_k the common prior on state $K, K \in \{1, .., S\}$ and assume that two parties receive signals that are correct conditional on the state being true with probability $1 - \varepsilon$. With probability $\varepsilon/(S-1)$ the signal indicates falsely each of the other states. Let $\alpha = (\alpha_1, .., 1 - \sum_{k \neq S} \alpha_k)$ be the vector of the prior distribution. As before, let $\tau(\alpha)$ be the maximizer of $u(\alpha, \tau)$ and assume that it is unique. Also, assume that u is strictly concave in τ for any belief α.[29]Identical arguments as before establish:

Proposition 16 *With $S \geq 2$ states the unique pooling PSE outcome is $\tau(\alpha)$.*

It is also easy to see that signal k is strong iff $\alpha_k > \varepsilon$. Therefore, it seems that a necessary condition for a separating PSE when all signals are strong is $\tau_k = \tau_k^* \equiv \tau(\mu(k, k))$. Whether this can be shown and, if true, whether this can be used to derived conditions that are not generically satisfied is an open question.

7 Conclusions

The empirical content of the theory proposed in this paper is that one should observe policy convergence in policy areas where political expertise is advanced (ε small) while policy can diverge if parties' expertise is moderate or limited. This is fairly intuitive as it seems natural to expect more variance in opinions, and accordingly in policy proposals, when expertise is moderate. The subtle twist here is that with excellent expertise policy converges to an uninformative, conventional wisdom type of policy that does not vary with the available information.[30] Thus, parties appear to be experimenting when uncertainty is significant but not when they are almost certain what the truth is.

One way to test this empirically would be to check whether the electorate's policy preferences differ before and after parties formulate their policies and depending on whether these

[29]Observe that now monotonicity of $\tau(\alpha)$ can only be defined in terms of first order stochastic dominance. That is, considering two distributions α and α' and assuming α' first order stochastically dominates α, the natural monotonicity assumption would be $\tau(\alpha) < \tau(\alpha')$.

[30]As the mechanism that yields to policy convergence rests on (the likelihood of) ties between signals, the result seems unlikely to carry over to setup where signals are continuous.

areas are characterized by excellent or moderate political expertise. In that vein, the U.S. decision whether or not to invade Iraq could be perceived as an issue with limited political expertise that allowed for divergent policy proposals (e.g. by then Illinois state senator Obama) while the question whether or not to bail out General Motors is one where parties' expertise is quite developed and the general consensus is in favor of a bail-out.

Wittman (1989) has argued forcefully that political competition would result in efficient political outcomes. The present paper provides a model where political competition cannot be relied upon to generate efficient outcomes. Callander (2008, p.680) notes that "standard intuition about the median voter theorem doesn't extend to incomplete information environments." To the extent that there is another non-Downsian equilibrium whenever one signal is weak the present paper corroborates Callander's findings. However, as parties' signals become very precise, the generically unique equilibrium involves pooling at the voter's uninformed bliss point policy, which corresponds to the median voter theorem.

Appendix

A Proofs

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

Proof of Lemma 2: I first show that $E_1[\tau_1^k \mid s_1 = k] + E_2[\tau_2^k \mid s_2 = k] = 1$ for both $k = a, b$. Once this is established, the lemma follows rather straightforwardly because of the assumption that the voter's strategy does not depend on party labels.

Recall that $\gamma(\tau_1, \tau_2)$ denotes the probability that the voter elects party 1 if 1 plays τ_1 and

2 plays τ_2. So

$$\begin{aligned}
E_1[\tau_1^a \mid s_1 = a] &= \mu_1(s_2 = a \mid s_1 = a)\gamma(\tau_1^a, \tau_2^a) + (1 - \mu_1(s_2 = a \mid s_1 = a))\gamma(\tau_1^a, \tau_2^b) \\
E_2[\tau_2^a \mid s_2 = a] &= \mu_2(s_1 = a \mid s_2 = a)(1 - \gamma(\tau_1^a, \tau_2^a)) + (1 - \mu_2(s_1 = a \mid s_2 = a))(1 - \gamma(\tau_1^b, \tau_2^a)) \\
E_1[\tau_1^b \mid s_1 = b] &= \mu_1(s_2 = b \mid s_1 = b)\gamma(\tau_1^b, \tau_2^b) + (1 - \mu_1(s_2 = b \mid s_1 = b))\gamma(\tau_1^b, \tau_2^a) \\
E_2[\tau_2^b \mid s_2 = b] &= \mu_2(s_1 = b \mid s_2 = b)(1 - \gamma(\tau_1^b, \tau_2^b)) \\
&+ (1 - \mu_2(s_1 = b \mid s_2 = b))(1 - \gamma(\tau_1^a, \tau_2^b)).
\end{aligned}$$

Notice that $\mu_1(s_2 = k \mid s_1 = k) = \mu_2(s_1 = k \mid s_2 = k)$ for $k = a, b$. To simplify notation, I write $\theta_k \equiv \mu_1(s_2 = k \mid s_1 = k)$, $x \equiv \gamma(\tau_1^a, \tau_2^a)$, $y \equiv \gamma(\tau_1^a, \tau_2^b)$, $c \equiv 1 - \gamma(\tau_1^b, \tau_2^a)$ and $d \equiv \gamma(\tau_1^b, \tau_2^b)$.

The incentive constraints (2) can now be written as

$$\begin{aligned}
E_1[\tau_1^a \mid s_1 = a] = \theta_A a + (1 - \theta_a)y &\geq \theta_a(1 - c) + (1 - \theta_a)d \\
&= E_1[\tau_1^b \mid s_1 = a] \qquad (9) \\
E_2[\tau_2^a \mid s_2 = a] = \theta_a(1 - x) + (1 - \theta_a)c &\geq \theta_a(1 - y) + (1 - \theta_a)(1 - d) \\
&= E_2[\tau_2^b \mid s_2 = a] \qquad (10) \\
E_1[\tau_1^b \mid s_1 = b] = \theta_b d + (1 - \theta_b)(1 - c) &\geq \theta_b y + (1 - \theta_b)x \\
&= E_1[\tau_1^a \mid s_1 = b] \qquad (11) \\
E_2[\tau_2^b \mid s_2 = b] = \theta_b(1 - d) + (1 - \theta_b)(1 - y) &\geq \theta_b c + (1 - \theta_b)(1 - x) \\
&= E_2[\tau_2^a \mid s_2 = b]. \qquad (12)
\end{aligned}$$

Adding (9) and (10) yields $1 \leq y + c$ while adding (11) and (12) implies $1 \geq y + c$. Thus, $y + c = 1$ holds.

Now,

$$E_1[\tau_1^a \mid s_1 = a] + E_2[\tau_2^a \mid s_2 = a] = \theta_a + (1 - \theta_a)(y + c) = 1 \qquad (13)$$

and

$$E_1[\tau_1^b \mid s_1 = b] + E_2[\tau_2^b \mid s_2 = b] = \theta_b + (1 - \theta_b)(2 - (y + c)) = 1, \qquad (14)$$

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 (3) holds, suppose to the contrary that it does not. Without loss of generality, assume $E_1[\tau_1^a \mid s_1 = a] < E_2[\tau_2^a \mid s_2 = a]$. Since $E_1[\tau_1^a \mid s_1 = a] + E_2[\tau_2^a \mid s_2 = a] = 1$, this implies $E_1[\tau_1^a \mid s_1 = a] < 1/2$. But now, upon $s_1 = a$ party 1 could play τ_2^a instead of the prescription τ_1^a, in which case it would get

$$E_1[\tau_2^a \mid s_1 = a] = \theta_a \gamma(\tau_2^a, \tau_2^a) + (1 - \theta_a)\gamma(\tau_2^a, \tau_2^b), \qquad (15)$$

or τ_2^b, in which case it would get

$$E_1[\tau_2^b \mid s_1 = a] = \theta_a \gamma(\tau_2^b, \tau_2^a) + (1-\theta_a)\gamma(\tau_2^b, \tau_2^b). \tag{16}$$

Due to the assumption that the voter's strategy must not depend on the parties' labels, $\gamma(\tau_2^a, \tau_2^a) = \gamma(\tau_2^b, \tau_2^b) = \frac{1}{2}$ and $\gamma(\tau_2^b, \tau_2^a) = 1 - \gamma(\tau_2^a, \tau_2^b)$. Thus, these two equations simplify to

$$E_1[\tau_2^a \mid s_1 = a] = \theta_a \frac{1}{2} + (1-\theta_a)\gamma(\tau_2^a, \tau_2^b) = \gamma(\tau_2^a, \tau_2^b) + \theta_a \left(\frac{1}{2} - \gamma(\tau_2^a, \tau_2^b)\right) \tag{17}$$

and

$$E_1[\tau_2^b \mid s_1 = a] = \theta_a(1 - \gamma(\tau_2^a, \tau_2^b)) + (1-\theta_a)\frac{1}{2} = \frac{1}{2} + \theta_a \left(\frac{1}{2} - \gamma(\tau_2^a, \tau_2^b)\right). \tag{18}$$

The expression in (17) weakly exceeds 1/2 if $\gamma(\tau_2^a, \tau_2^b) \geq \frac{1}{2}$ and the expression in (18) is (strictly) larger than 1/2 otherwise. Since party 1 can either play τ_2^a or τ_2^b, it has to be the case that its expected equilibrium payoff weakly exceeds 1/2. And since exactly the same argument applies for party 2, it follows that indeed $E_i[\tau_i^k \mid s_i = k] = E_j[\tau_j^k \mid s_j = k] = \frac{1}{2}$ for $k \in \{a, b\}$ as claimed. ■

Proof of Lemma 3: Assume (4) does not hold, e.g. because $u(\alpha, \tau_a) < u(\alpha, \tau_b)$, yet τ_a and τ_b are set in a separating PBE. But then the deviation to play τ_b when the signal is a pays off when the other party plays τ_a by increasing the probability of winning from 1/2 to 1 and when the other other party plays τ_b by increasing the probability of winning from 0 to 1/2.■

Proof of Proposition 1: If both parties play τ independently of their signals, the voter is indifferent between the two parties and randomizes. On equilibrium, no information is transmitted and the voter's posterior equals her prior α. If e.g. party 1 deviates to some $\tau_1 \neq \tau$, then the voter must vote for 1 with a probability smaller than $\frac{1}{2}$. For this to be sequentially rational, her off equilibrium belief $\mu(\tau_1, \tau)$ must be such that her expected utility of voting for party 2 exceeds her utility of voting for the deviating party 1. Though PBE does not restrict the off equilibrium beliefs of the voter about the strategy played by the deviating party, it still imposes bounds on her beliefs $\mu(\tau_1, \tau)$: Given the observation (τ_1, τ) the belief most favorable for A is that the voter assumes party 1 plays the strategy "τ_1 if $s_1 = a$ and τ otherwise" while the least favorable belief is that 1 plays "τ_1 if $s_1 = b$ and τ otherwise". These assumptions imply, respectively, the updated beliefs $\mu(\tau_1, \tau) = \frac{\alpha(1-\varepsilon)}{\alpha(1-\varepsilon)+(1-\alpha)\varepsilon} = \mu(a, 0) < 1$ and $\mu(\tau_1, \tau) = \frac{\alpha\varepsilon}{\alpha\varepsilon+(1-\alpha)(1-\varepsilon)} = \mu(b, 0) > 0$. The voter's preferred policies, given these beliefs, are $\underline{\tau}(\alpha) \equiv \tau(\mu(a, 0)) < \overline{\tau}(\alpha) \equiv \tau(\mu(b, 0))$. Hence, for any "prescribed" equilibrium policy

$\tau \in [\underline{\tau}(\alpha), \overline{\tau}(\alpha)]$ there are beliefs that make it rational not to vote for the deviating party: Just choose the off equilibrium beliefs $\mu(\tau_1, \tau)$ so that $\tau = \tau(\mu(\tau_1, \tau))$.

As for the separating PBE, notice first that one party choosing τ_a and the other one τ_b is an on equilibrium observation. Using Bayes' rule, the voter updates her beliefs to $\mu(\tau_a, \tau_b) = \frac{\alpha(1-\varepsilon)\varepsilon}{\alpha(1-\varepsilon)\varepsilon+(1-\alpha)(1-\varepsilon)\varepsilon} = \alpha$. Hence, the voter will be indifferent between the two. If both parties choose the same policy, he will also be indifferent between the two. In either case, she randomizes uniformly. If one party deviates and plays an off equilibrium policy τ, she must not vote for the deviating party with probability larger than one half. Upon observing (τ_1, τ_a) where τ_1 is the off equilibrium observation generated by party 1 and τ_a is the on equilibrium policy party 2 plays upon receiving $s_2 = a$ the voter's hypothesis that is most favorable for state A is that 1 plays the strategy "τ_1 if and only if $s_1 = a$". Consequently, $\max \mu(\tau_1, \tau_a) = \frac{\alpha(1-\varepsilon)^2}{\alpha(1-\varepsilon)^2+(1-\alpha)\varepsilon^2} = \mu(a, a)$. The least favorable hypothesis is "τ_1 if and only if $s_1 = b$", yielding $\min \mu(\tau_1, \tau_a) = \frac{\alpha\varepsilon(1-\varepsilon)}{\alpha\varepsilon(1-\varepsilon)+(1-\alpha)(1-\varepsilon)\varepsilon} = \alpha$. Similarly, upon observing (τ_1, τ_b) the most favorable hypothesis for state A is that 1 plays "τ_1 if and only if $s_1 = a$", yielding $\max \mu(\tau_1, \tau_b) = \frac{\alpha\varepsilon(1-\varepsilon)}{\alpha\varepsilon(1-\varepsilon)+(1-\alpha)(1-\varepsilon)\varepsilon} = \alpha$, and the least favorable hypothesis is "τ_1 if and only if $s_1 = b$", yielding $\min \mu(\tau_1, \tau_b) = \frac{\alpha\varepsilon^2}{\alpha\varepsilon^2+(1-\alpha)(1-\varepsilon)^2} = \mu(b, b)$.

Assume $\tau_a, \tau_b \in [\tau_a^*, \tau_b^*]$, where $u(\alpha, \tau_a) = u(\alpha, \tau_b)$. For any off equilibrium τ, i.e. for any $\tau \neq \tau_a, \tau_b$ there are beliefs that make it rational not to vote for the deviating party. Without loss of generality assume that the party that plays on equilibrium plays τ_b. If $\tau > \tau_b$, the voter can choose the belief α, so that $u(\alpha, \tau_b) > u(\alpha, \tau)$. For $\tau < \tau_b$, she can choose the belief $\mu(b, b)$ so that $u(\mu(b, b), \tau_b) > u(\mu(b, b), \tau)$. ■

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

Proof of Proposition 2: Observe first that if the voter elects the deviator with probability one both types of a party potentially benefit from a deviation. Thus, the off equilibrium belief consistent with PSE is $\mu = \alpha$. Next I show that no $\tau \neq \tau(\alpha)$ is a pooling PSE outcome. Assume without loss of generality that $\tau < \tau(\alpha)$. Because $u(\alpha, \tau)$ has a unique maximum it follows that the voter will prefer any $\tau' \in (\tau, \tau(\alpha)]$ to the proposed equilibrium policy. Hence, there are profitable deviations and thus no $\tau \neq \tau(\alpha)$ can be a pooling PSE outcome. To see that $\tau(\alpha)$ is a pooling PSE outcome, notice that $\tau(\alpha)$ is the unique maximizer of $u(\alpha, \tau)$. Hence, the voter will strictly prefer $\tau(\alpha)$ to any other policy as long as her beliefs are α, which they will be, even off equilibrium, as just argued. ■

Proof of Lemma 5: Recall that for $\alpha \in (\varepsilon, 1-\varepsilon)$ both signals are strong. On equilibrium, i wins with probability 1/2 independently of its signal. Therefore, if there is a deviation that allows i to win with certainty against τ_k and to lose with certainty against τ_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 $\tau_a^* \leq \tau_a$ and $\tau_b \leq \tau_b^*$. I will now argue that for $\tau_a^* < \tau_a$ and $\tau_b < \tau_b^*$ such a deviation exists, this deviation being τ_k^* upon signal k.

If only a party with signal a (type a for short) benefits from playing τ_a^*, the voter's belief upon observing (τ_a, τ_a^*) is $\mu(a, a)$, in which case she strictly prefers τ_a^* to τ_a by construction of τ_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 τ_a^* the voter's belief upon observing (τ_b, τ_a^*) is $\mu(b, 0)$ because the deviating party's behavior is not informative. But recall now from Lemma 3 that $u(\alpha, \tau_a) = u(\alpha, \tau_b)$. Therefore, upon (τ_b, τ_a^*) and having belief $\mu(b, 0) < \alpha$, the voter strictly prefers τ_b to τ_a^* since $\tau_a^* < \tau_a$. Therefore, type b's payoff from the deviation τ_a^* is $(1 - \mu(b \mid b))$, 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 τ_a^*. (And indeed, if only type a benefits, the voter's belief upon observing (τ_b, τ_a^*) is $\mu(a, b) = \alpha$, in which case she strictly prefers τ_b.) Completely analogous reasoning applies for $\tau_b < \tau_b^*$, Therefore, for $\tau_a > \tau_a^*$ ($\tau_b < \tau_b^*$) playing τ_a^* upon signal A (τ_b^* upon signal b) is a deviation that pays off. ■

Proof of Proposition 3: A pure strategy for a party can be either pooling, i.e. it sets the same policy independently of its signal, or separating, i.e. it sets different policies as a function of its signal. If the voter updates as required by Grossman and Perry (1986), the (or a) best response against a party who pools is $\tau(\alpha)$, and $\tau(\alpha)$ is the unique best response to itself, as shown in Proposition 2.

Consider now a candidate separating equilibrium. Then upon a strong signal k it must be the case that parties play τ_k^*. Because for $\alpha \in (\varepsilon, 1-\varepsilon)$ both signals are strong, separation in a PSE requires conditions (5) and (6) to hold. These impose two independent restrictions on τ_a and τ_b that in general will not hold simultaneously. So for $\alpha \in (\varepsilon, 1-\varepsilon)$ there is generically no PSE where both parties separate.

Moreover, since $\hat{\tau}_a$ and $\hat{\tau}_b$ are unique, there are no hybrid PSE where one party separates by playing $\hat{\tau}_k$ upon signal k and the other party pools by playing some τ independently of its signal since the voter will always strictly prefer $\hat{\tau}_k$ to $\tau \neq \hat{\tau}_a, \hat{\tau}_b$. It can also not be the case that one party pools at, say, $\tau = \tau_a^*$ and the other one separates because the separating party would have an incentive to play $\tau(\alpha)$, whereby it would win for sure. Thus, for $\alpha \in (\varepsilon, 1-\varepsilon)$ the generically unique PSE outcome is $\tau(\alpha)$. ■

Proof of Proposition 4: Notice first that 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 a party has a posterior exceeding (less than) 1/2 that the other party has received the same signal. For the sake of the argument, suppose a is the strong signal, i.e. $\alpha > \varepsilon$, and assume $\tau_a \neq \hat{\tau}_a$. A necessary condition for (τ_a, τ_b) to be part of a separating PSE is that they satisfy $u(\tau_a, \alpha) = u(\tau_b, \alpha)$; see Lemma 3. I am now going to show that unless $\tau_a = \hat{\tau}_a$ holds, (τ_a, τ_b) cannot be part of a separating PSE. To see this, notice that party i can potentially benefit from a deviation both after $s_i = A$ and $s_i = b$ if upon the deviation it is elected with probability one if the other party plays τ_a. So upon seeing (τ_a, τ'), where τ' is a deviation by i, the voter's belief, updated according to PSE, is $\mu(a, 0)$. So unless $\tau_a = \hat{\tau}_a$, the voter prefers $\tau' = \hat{\tau}_a$ to τ_a. So on top of $u(\tau_a, \alpha) = u(\tau_b, \alpha)$ PSE requires $\tau_k = \hat{\tau}_k$, 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 $u(\tau_a, \alpha) = u(\tau_b, \alpha)$. That such (τ_a, τ_b) exist is guaranteed by the symmetry assumption since $\hat{\tau}_a > \tau(1)$ and $\hat{\tau}_b < \tau(0)$.

To conclude the argument, 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 τ_b, provided $\tau_a = \hat{\tau}_a$ and

$u(\alpha, \hat{\tau}_a) = u(\alpha, \tau_b)$. By on equilibrium play i wins with probability 1/2 upon either signal. By construction of $\hat{\tau}_a$ there is no deviation that both types can beneficially play. Therefore, the only deviation that potentially benefits a party with signal b is one that the voter prefers if the other party plays τ_b. However, $\mu_i(s_j = b \mid s_i = b) < 1/2$ because b is the weak signal, so that the expected subjective payoff of the deviation is less than 1/2. Hence, there is no deviation that benefits only party b. Thus, no profitable deviation exists. ■

Proof of Lemma 6: Consider feedback rule I first and assume $u(\alpha, \tau_a) > u(\alpha, \tau_b)$. The expected payoff of playing τ_b upon $s_i = b$ is

$$E_i[\tau_b \mid s_i = b] \quad = \quad \mu_i(b \mid b)/2 + \lambda(1 - \mu_i(b \mid b))(1 - \alpha),$$

where $\mu_i(b \mid b)$ is shorthand for $\mu_i(s_j = b \mid s_i = b)$. The expression $E_i[\tau_b \mid s_i = b]$ is explained as follows. The first term captures the event when both parties get the same signal and thus play the same action. In this event, i wins with probability 1/2 irrespective of whether the voter remains ignorant or not. The only chance for i to win if the opponent gets signal $s_j = a$ is that the voter learns the truth when the truth is B. Since the other party has received a contradicting signal, the correct belief that A is true in this contingency is α. So the probability that B is true in this event is $1 - \alpha$. Now a sufficient condition for $E_i[\tau_b \mid s_i = b] < 1/2$ is $\lambda \leq 1/2$.

On the other hand the expected payoff of playing τ_a upon $s_i = b$ is

$$E_i[\tau_a \mid s_i = b] \quad = \quad (1 - \mu_i(b \mid b))/2 + \mu_i(b \mid b)[\lambda\mu_i(\omega = A \mid s_i = b, s_j = b) + (1 - \lambda)].$$

With probability $(1 - \mu_i(b \mid b))$ the opponent has received the signal $s_j = a$ and plays τ_a as well, in which case i wins with probability 1/2. With probability $\mu_i(b \mid b)$ j gets signal b and plays τ_b. Now i wins with probability 1 whenever the voter remains ignorant, which happens with probability $1 - \lambda$ and if she learns that the truth is A, which happens with probability $\lambda\mu_i(\omega = A \mid s_i = b, s_j = b)$, given that both parties have got the signal b. Obviously, $E_i[\tau_a \mid s_i = b] > (1 - \mu_i(b \mid b)/2 + \mu_i(b \mid b)(1 - \lambda)$, which is (weakly) greater than 1/2 for $\lambda \leq 1/2$. Thus, the deviation pays off.

Consider feedback rule II now. The expected payoff of playing τ_b upon signal $s_i = b$ is

$$E_i[\tau_b \mid s_i = b] \quad = \quad \mu_i(b \mid b)/2 + \lambda(1 - \mu_i(b \mid b))[\alpha\varepsilon + (1 - \alpha)(1 - \varepsilon)],$$

which is less than 1/2 for $\lambda \leq 1/2$ since $\alpha\varepsilon + (1-\alpha)(1-\varepsilon) < 1$.[31] The expected payoff of playing τ_a is

$$E_i[\tau_a \mid s_i = b] \;=\; (1-\mu_i(b \mid b))/2 + \mu_i(b \mid b)[\lambda\mu_i(s_V = a \mid s_i = b, s_j = b) + (1-\lambda)],$$

which is strictly more than $(1-\mu_i(b \mid b))/2 + \mu_i(b \mid b)(1-\lambda)$, which in turn exceeds 1/2 for $\lambda < 1/2$. Thus, the deviation pays off under feedback rule II as well. ■

Proof of Lemma 7: Let $\sigma \in [0,1]$ be the probability that the voter elects the party proposing τ_b if the other one proposes τ_a. A necessary condition for (τ_a, τ_b) to be equilibrium policies is

$$\begin{aligned}
E_i[\tau_a \mid s_i = a] &= \mu(a \mid a)/2 + (1-\mu(a \mid a))[(1-\lambda)(1-\sigma) + \lambda\alpha] \\
&\geq \mu(a \mid a)[(1-\lambda)\sigma + \lambda(1-\mu(a,a))] + (1-\mu(a \mid a))/2 = E_i[\tau_b \mid s_i = a] \\
E_i[\tau_b \mid s_i = b] &= \mu(b \mid b)/2 + (1-\mu(b \mid b))[(1-\lambda)\sigma + \lambda(1-\alpha)] \\
&\geq \mu(b \mid b)[(1-\lambda)(1-\sigma) + \lambda\mu(b,b)] + (1-\mu(b \mid b))/2 = E_i[\tau_a \mid s_i = b].
\end{aligned}$$

Substituting and rearranging reveals that the first inequality holds if and only if

$$\sigma \leq \frac{2\alpha\lambda + (1-2\lambda)(\alpha+\varepsilon-2\alpha\varepsilon)}{2(1-\lambda)(\alpha+\varepsilon-2\alpha\varepsilon)} := \overline{\sigma} \tag{19}$$

and that the second holds if and only if

$$\sigma \geq \frac{2\alpha\varepsilon\lambda + (1-2\lambda)(1-\varepsilon-\alpha+2\alpha\varepsilon)}{2(1-\lambda)(1-\varepsilon-\alpha+2\alpha\varepsilon)} := \underline{\sigma}. \tag{20}$$

These bounds will prove helpful shortly.[32] Notice also that for $\lambda = 0$, $\overline{\sigma} = \underline{\sigma} = 1/2$.

Now suppose the equilibrium prescribed $\tau_a < \tau_a^*$ upon signal a. For the reasons analogous to those outlined in the proof of Lemma 4 the worst belief for the party playing τ_a^* the voter who gets no additional information can have is $\mu(a,a)$. But for any $\mu \leq \mu(a,a)$, $u(\mu, \tau_a) < u(\mu, \tau_a^*)$ holds. Therefore, the voter if ignorant will prefer τ_a^* to τ_a. Consequently, the payoff to a player with signal a of the deviation τ_a^* will be

$$E_i[\tau_a^* \mid s_i = a] \;=\; \mu(a \mid a)(1-\lambda) + (1-\mu(a \mid a))\lambda\alpha.$$

Since $\sigma \geq \underline{\sigma}$, the expected equilibrium payoff is

$$E_i[\tau_a \mid s_i = a] \leq \mu(a \mid a)/2 + (1-\mu(a \mid a))[(1-\lambda)(1-\underline{\sigma}) + \lambda\alpha].$$

[31] To see the latter, notice that $\alpha\varepsilon + (1-\alpha)(1-\varepsilon) = 1-\varepsilon-\alpha(1-2\varepsilon)$, where the term in parenthesis is positive.

[32] They obviously do not depend on the specific candidate equilibrium under consideration. Rather, they must be satisfied in any separating PBE.

Plugging in and rearranging reveals that

$$\mu(a \mid a)(1-\lambda) + (1-\mu(a \mid a))\lambda\alpha > \mu(a \mid a)/2 + (1-\mu(a \mid a))[(1-\lambda)(1-\underline{\sigma}) + \lambda\alpha]$$

holds if and only if

$$\lambda < \frac{4\alpha^2\varepsilon^2 - 4\varepsilon^3\alpha - 4\alpha^2\varepsilon + 2\varepsilon^3 + 2\alpha\varepsilon + 2\alpha\varepsilon^2 - \alpha + \alpha^2 - 3\varepsilon^2 + \varepsilon}{2[-4\alpha^2\varepsilon + 4\alpha^2\varepsilon^2 - 3\varepsilon^3\alpha + \alpha^2 - \varepsilon^2 - \alpha + 3\alpha\varepsilon + \varepsilon^3]} := \overline{\lambda}_a(\varepsilon, \alpha). \tag{21}$$

It is readily checked that $\overline{\lambda}_a(0, \alpha) = 1/2$ and $\overline{\lambda}_a(\varepsilon, \alpha) > 0$ for all $\alpha > \varepsilon$. Thus, for $\lambda < \overline{\lambda}_a(\varepsilon, \alpha)$ the deviation to τ_a^* upon signal a pays off.

Analogously, if the candidate PBE prescribed $\tau_b > \tau_b^*$, the deviation to τ_b^* upon signal b yields a payoff of

$$E_i[\tau_b^* \mid s_i = b] \quad = \quad \mu(b \mid b)(1-\lambda) + (1-\mu(b \mid b))\lambda(1-\alpha).$$

Since $\sigma \leq \overline{\sigma}$, the expected equilibrium payoff is

$$E_i[\tau_b \mid s_i = b] \leq \mu(b \mid b)/2 + (1-\mu(b \mid b))[(1-\lambda)\overline{\sigma} + \lambda(1-\alpha)].$$

The usual tedious algebra reveals that the deviation pays off if

$$\lambda < \frac{4\alpha^2\varepsilon^2 + 4\varepsilon^3\alpha - 4\alpha^2\varepsilon + \alpha^2 + 6\alpha\varepsilon - 10\alpha\varepsilon^2 - \alpha - \varepsilon + 3\varepsilon^2 - 2\varepsilon^3}{2[-4\alpha^2\varepsilon + 4\alpha^2\varepsilon^2 + 3\varepsilon^3\alpha + \alpha^2 - \varepsilon + 3\varepsilon^2 - \alpha + 5\alpha\varepsilon - 8\alpha\varepsilon^2 - 2\varepsilon^3]} := \overline{\lambda}_b(\varepsilon, \alpha), \tag{22}$$

where $\overline{\lambda}_b(\varepsilon, \alpha) > 0$ for all $\alpha < 1-\varepsilon$ and $\overline{\lambda}_b(0, \alpha) = 1/2$. Letting $\overline{\lambda}(\varepsilon, \alpha) := \min\{\overline{\lambda}_a(\varepsilon, \alpha), \overline{\lambda}_b(\varepsilon, \alpha)\}$ and noting that $\overline{\lambda}_a$ and $\overline{\lambda}_b$ are continuous functions completes the proof. ■

Proof of Lemma 8: Assume the equilibrium prescribed $\tau_b < \tau_b^*$. If player i benefits from the deviation only upon signal $s_i = b$, which remains to be verified shortly, the deviation payoff is

$$E_i[\tau_b^* \mid s_i = b] = \mu(b \mid b)[1-\lambda + \lambda(1-\mu(b, b))] + (1-\mu(b \mid b))\lambda(1-\alpha).$$

Since $\sigma \leq \overline{\sigma}$ the expected equilibrium payoff is no bigger than

$$\mu(b \mid b)/2 + (1-\mu(b \mid b))[(1-\lambda)\overline{\sigma} + \lambda(1-\alpha)].$$

Algebra reveals that this is less than $E_i[\tau_b^* \mid s_i = b]$ for any $\lambda < \overline{\lambda}(\varepsilon, \alpha)$. It thus remains to show that upon signal $s_i = a$ there is no incentive to play τ_b^*. The deviation payoff upon signal $s_i = a$ is

$$E_i[\tau_b^* \mid s_i = a] = (1-\mu(a \mid a))[1-\lambda + \lambda(1-\alpha)] + \mu(a \mid a)\lambda(1-\mu(a, a)).$$

On the other hand, the expected equilibrium payoff is at least

$$\mu(a \mid a)/2 + (1 - \mu(a \mid a)[(1-\lambda)(1-\overline{\sigma}) + \alpha\lambda]$$

since $\overline{\sigma} \geq \sigma$. It can be shown that the expected equilibrium payoff upon $s_i = a$ is strictly larger than $E_i[\tau_b^* \mid s_i = a]$ for any $\lambda < \overline{\lambda}(\varepsilon, \alpha)$. Therefore, under the conditions of the lemma, the deviation to τ_b^* pays off upon signal b and not upon signal a. Consequently, unless $\tau_b = \tau_b^*$, τ_b is not a policy of a separating PSE.

Analogously, upon $s_i = a$ the payoff for a deviator setting τ_a^* is

$$E_i[\tau_a^* \mid s_i = a] = \mu(a \mid a)[1 - \lambda + \lambda\mu(a, a)] + (1 - \mu(a \mid a))\lambda\alpha$$

while the expected equilibrium payoff is no more than

$$\mu(a \mid a)/2 + (1 - \mu(a \mid a))[(1-\lambda)(1-\underline{\sigma}) + \lambda\alpha)]$$

because $\underline{\sigma} \leq \sigma$. Algebra reveals that the former exceeds the latter for any $\lambda < \overline{\lambda}(\varepsilon, \alpha)$. It thus remains to be shown that the deviation to τ_a^* does not pay upon $s_i = b$. The expected deviation payoff is

$$E_i[\tau_a^* \mid s_i = b] = (1 - \mu(b \mid b))[1 - \lambda + \lambda\alpha] + \mu(b \mid b)\lambda\mu(b, b),$$

while the expected equilibrium payoff is, at least,

$$\mu(b \mid b)/2 + (1 - \mu(b \mid b))[(1-\lambda)\underline{\sigma} + \lambda(1-\alpha)]$$

since $\sigma \geq \underline{\sigma}$. For all $\lambda < \overline{\lambda}(\varepsilon, \alpha)$, the latter exceeds the former. Therefore, the deviation to τ_a^* pays off for a party with signal a but not with signal b. Consequently, τ_a is an equilibrium policy in a separating PSE only if $\tau_a = \tau_a^*$. ■

Proof of Proposition 14: Provided it exists, the uniqueness of the pooling PSE with outcome $\tau(\alpha)$ follows along the lines of the model with $\lambda = 0$ studied above.

Consider feedback rule I first. If the candidate equilibrium prescribes pooling at $\tau(\alpha)$, i wins with probability 1/2 upon either signal. If, say, upon $s_i = a$ i plays some smaller policy $\tau \in [\tau(1), \tau(\alpha))$, he wins for sure with probability $\lambda\mu_i(\omega = A \mid s_i = a)$. A sufficient condition for this to be smaller than 1/2 is $\lambda < \frac{1}{2\mu_i(\omega = A \mid s_i = a)}$.

However, the proof is not yet complete as an additional benefit from the deviation to $\tau < \tau(\alpha)$ could be that the voter chooses τ instead of $\tau(\alpha)$ if she remains ignorant. This strategy on behalf of the voter requires that her belief μ be bigger than α. This in turn requires that i upon $s_i = b$ does not have an incentive to play τ. The payoff of this, under the assumption that the voter holds the beliefs just described, is

$$E_i[\tau \mid s_i = b] = (1 - \lambda) + (1 - \mu_i(\omega = B \mid s_i = b))\lambda$$

because τ wins if the voter remains ignorant or if the true state is A and the voter learns this. But $E_i[\tau \mid s_i = b] > 1/2$ for $\lambda < \frac{1}{2\mu_i(\omega=B|s_i=b)}$. Assuming b is weak and a is strong, $\lambda < \frac{1}{2\mu_i(\omega=A|s_i=a)}$ implies $\lambda < \frac{1}{2\mu_i(\omega=A|s_i=a)} < 1 < \frac{1}{2\mu_i(\omega=B|s_i=b)}$. So the deviation to $\tau < \tau(\alpha)$ would pay off to i upon $s_i = b$ for any $\lambda \leq 1$ if the ignorant voter preferred $\tau < \tau(\alpha)$ to $\tau(\alpha)$.[33] Therefore, any deviation would be pooling so that the ignorant voter cannot prefer τ to $\tau(\alpha)$. Consequently, only the informed voter would choose τ and thus the payoff of the deviation was specified correctly as $\lambda\mu_i(\omega = A \mid s_i = a)$. A completely symmetric argument applies for the case when $s_i = a$ is weak. So, if b is a strong signal and a is weak, the pooling PSE exists if and only if $\lambda < \frac{1}{2\mu_i(\omega=B|s_i=b)}$.

Turn now to the case where both signals are strong. Then the pooling PSE exists if and only if $\lambda < \min\left\{\frac{1}{2\mu_i(\omega=A|s_i=a)}, \frac{1}{2\mu_i(\omega=B|s_i=b)}\right\}$. To see this, suppose $\alpha > 1/2$ so that $\mu_i(A \mid a) > \mu_i(B \mid b) > 1/2$. Then for $\lambda \in \left(\frac{1}{2\mu_i(\omega=A|s_i=a)}, \frac{1}{2\mu_i(\omega=B|s_i=b)}\right)$ i deviates upon signal $s_i = a$ to $\tau < \tau(\alpha)$ and attracts the informed voter who learned that $\omega = A$. If $\lambda > \frac{1}{2\mu_i(\omega=B|s_i=b)}$, i deviates upon signal a and attracts both the informed voter who learned the truth is A and the ignorant voter. Analogous arguments apply for $\alpha < 1/2$ and deviations to $\tau > \tau(\alpha)$ upon signal $s_i = b$.

Consider feedback rule II now and assume without loss of generality that signal a is strong. We first show that for $\lambda \leq 1/2$ both types of party i would benefit from a deviation to $\tau < \tau(\alpha)$ if both the voter with signal $s_V = a$ and the voter who remains ignorant choose τ. From this it will follow that for $\lambda \leq 1/2$ the deviation can, at most, pay off for a party with signal $s_i = a$ who will be chosen by the voter if and only if the voter gets the signal $s_V = a$. If both ignorant and signal a types of the voter prefer a τ somewhat smaller than $\tau(\alpha)$ to $\tau(\alpha)$, the expected payoffs of playing τ are

$$\begin{aligned} E_i[\tau \mid s_i = a] &= \mu_i(s_V = a \mid s_i = a)\lambda + (1 - \lambda) \\ E_i[\tau \mid s_i = b] &= \mu_i(s_V = a \mid s_i = b)\lambda + (1 - \lambda). \end{aligned}$$

[33]To see this, notice that $E_i[\tau \mid s_i = b] > 1/2 \Leftrightarrow \lambda > \frac{1}{2\mu_i(\omega=B|s_i=b)}$, where the fraction on the righthand side of the inequality exceeds 1 when signal b is weak (which is equivalent to $\mu_i(\omega = B \mid s_i = b) < 1/2$).

Now $E_i[\tau \mid s_i = a] > 1/2$ for any $\lambda \in [0,1]$ follows because a is a strong signal by assumption. Irrespective of whether b is strong or weak, $E_i[\tau \mid s_i = b] > 1/2$ holds for any $\lambda \leq 1/2$. Thus, both types would benefit from the deviation if the voter chose τ over $\tau(\alpha)$ when remaining ignorant. Therefore, the ignorant voter type cannot prefer τ to $\tau(\alpha)$. Hence, party i can potentially benefit from the deviation only if $s_i = a$, in which case his payoff is $\lambda\mu_i(s_V = a \mid s_i = a)$. But this is strictly less than 1/2, the expected equilibrium payoff, for $\lambda \leq 1/2$.

Now if b is a weak signal, then $E_i[\tau \mid s_i = b] > 1/2$ for any $\lambda \leq 1$. Therefore, the voter will never prefer τ to $\tau(\alpha)$ if she is ignorant. Consequently, the only potentially profitable deviation is one where the voter with signal $s_V = a$ chooses $\tau < \tau(\alpha)$. This requires $\mu_i(s_V = a \mid s_i = a)\lambda > 1/2 \Leftrightarrow \lambda > \frac{1}{2\mu_i(s_V=a|s_i=a)}$. (Notice: For any $\varepsilon < 1/2$, $\mu_i(s_V = a \mid s_i = b) < \mu_i(s_V = a \mid s_i = a)$ holds. So it is never the case that i wants to deviate upon signal $s_i = b$ but not upon $s_i = a$. If $\lambda > \frac{1}{2\mu_i(s_V=a|s_i=b)} > \frac{1}{2\mu_i(s_V=a|s_i=a)}$ both types of i will benefit from the deviation, but the voter, upon $s_V = a$, will still prefer τ to $\tau(\alpha)$ if τ is not too small.)

On the other hand, if both signals are strong, $E_i[\tau \mid s_i = b] = \mu_i(s_V = a \mid s_i = b)\lambda + (1-\lambda) < 1/2$ holds for $\lambda > \frac{1}{2(1-\mu_i(s_V=a|s_i=b)}$, where $\frac{1}{2(1-\mu_i(s_V=a|s_i=b)} < 1$. So there is a deviation that attracts both the voter who is ignorant and the voter with signal $s_V = a$ which is only played by party i upon signal a.

Defining $\hat{\lambda}(\alpha,\varepsilon) := \min\left\{\frac{1}{2(1-\mu_i(s_V=a|s_i=b)}, \frac{1}{2(1-\mu_i(s_V=b|s_i=a)}\right\}$ if $\varepsilon < \alpha < 1-\varepsilon$, $\hat{\lambda}(\alpha,\varepsilon) := \frac{1}{2\mu_i(s_V=a|s_i=a)}$ if $1-\varepsilon < \alpha$ and $\hat{\lambda}(\alpha,\varepsilon) := \frac{1}{2\mu_i(s_V=b|s_i=b)}$ if $\alpha < \varepsilon$ completes the proof. ■

Proof of Proposition 8: Let 1 play τ_k, $\tau_k = \hat{\tau}_k \equiv \arg\max_\tau u(\mu(k,0),\tau)$ upon signal $s_1 = k$ with $k = a, b$. If his signal is k as well, 2 plays $\hat{\tau}_k$ as well. So upon observing 1 and 2 play $\hat{\tau}_k$ the voter holds the belief $\mu(k,k)$ and randomizes between the two policies. If his signal is not k, 2 plays τ'_k defined as the $\tau \neq \hat{\tau}_k$ that solves $u(\alpha,\hat{\tau}_k) = u(\alpha,\tau)$. Upon observing $(\hat{\tau}_k,\tau'_k)$ the voter's belief is α and she is, again, indifferent between the two proposals.

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

Proof of Proposition 9: If both parties pool, the voter's posterior equals the prior both on and off equilibrium. So off and on equilibrium there's no policy she'd prefer to $\tau(\alpha)$. ■

Sketch of Proof of Proposition 10: The fact that τ^O exists is irrelevant in the sense that for any admissible belief $\mu \in [\mu(a,a), \mu(b,b)]$ τ_a^* or τ_b^* (or both) will be preferred to τ^O. ■

Proof of Proposition 11: If all others are playing the pooling policy $\tau^*(\alpha)$, the voter's posterior belief will be α on and off equilibrium because if a party of type k benefits from a deviation then so will he if he is of type h, $h \neq k$. But given this posterior, $\tau^*(\alpha)$ is the voter's bliss point policy. Therefore, no party can possibly gain from a deviation. Uniqueness follows along the usual lines. ■

Sketch of Proof of Lemma 9: Suppose party i of quality O can successfully signal its quality (and benefit from doing so). Then if its quality were C it would benefit from behaving as if it were of $q = O$, and vice versa. ■

B Robustness

B.1 Conditionally Correlated Signals

Heidhues and Lagerlöf (2003) find that signals that are, conditional on the state, positively correlated permit the existence of equilibria where more information is transmitted from the parties than with independent signals. It is therefore important to see what happens in the present set up if one allows for signals that are conditionally correlated.

The main insight from this exercise is that none of the above results are reversed. The intuition for this is quite clear. Since with independent signals the main obstacle to the existence of informative equilibria is that upon receiving its own signal each party is too confident that the other one has received the same signal for $\varepsilon < \alpha < 1 - \varepsilon$, additional correlation will work in the same direction. Therefore, signals that are positively correlated conditional on the state will only reinforce the effects present in the basic model and make the range of the prior larger where the unique generic PSE involves no information transmission.

Correlated signals are modelled as described in Table 1,[34] where $\Pr(s_i = h \mid \omega = K)$ is the

[34] The same technology is used by Bhaskar and van Damme (2002) and Heidhues and Lagerlöf (2003).

	$\Pr(s_2 = k \mid \omega = K)$	$\Pr(s_2 = h \mid \omega = K)$
$\Pr(s_1 = k \mid \omega = K)$	$(1-\varepsilon)^2 + \rho\varepsilon(1-\varepsilon)$	$(1-\rho)\varepsilon(1-\varepsilon)$
$\Pr(s_1 = h \mid \omega = K)$	$(1-\rho)\varepsilon(1-\varepsilon)$	$\varepsilon^2 + \rho\varepsilon(1-\varepsilon)$

Table 1: Signal Technology with Positive Correlation.

probability that i's signal indicates h while the truth is K with $H, K \in \{A, B\}$ and $h, k \in \{a, b\}$ and $\rho \in [0,1]$ is the degree of positive correlation. Observe that $\Pr(s_i = k \mid K) = \Pr(s_j = k \mid \omega = K) + \Pr(s_j = h \mid \omega = K) = 1 - \varepsilon$ as in the previous sections. Notice that for $\rho = 0$ this model coincides with the basic model introduced in Section 2. From the proof of Lemma 2 it becomes also clear that the lemma uses the beliefs $\mu(k \mid k)$ and so on as primitives. Therefore, Lemma 2 will continue to hold with $\rho > 0$. Thus, on equilibrium every party is elected with probability 1/2 independent of its signal.

Next tedious but straightforward derivations reveal that

$$\mu_i(s_j = a \mid s_i = a) = \frac{\alpha(1-2\varepsilon) + (1-\rho)\varepsilon^2 + \rho\varepsilon}{\alpha(1-2\varepsilon) + \varepsilon}, \tag{23}$$

which exceeds 1/2 if and only if $\alpha > \frac{\varepsilon[1-2(\varepsilon(1-\rho)+\rho)]}{1-2\varepsilon} \equiv \underline{\alpha}(\rho)$. Observe that $\underline{\alpha}(\rho)$ decreases in ρ so that is is largest when $\rho = 0$, in which case $\underline{\alpha}(0) = \varepsilon$, which is as it should be. Thus, signal a is strong if and only if $\alpha > \underline{\alpha}(\rho)$. On the other hand, and by a similar argument, signal b is strong if and only $1 - \alpha > \underline{\alpha}(\rho)$ implying $\alpha < \frac{1-3\varepsilon+2\varepsilon[\varepsilon(1-\rho)+\rho]}{1-2\varepsilon} \equiv \overline{\alpha}(\rho)$. Since $\overline{\alpha}(\rho)$ increases in ρ, it is smallest when $\rho = 0$, in which case $\overline{\alpha}(0) = 1 - \varepsilon$, which is as it ought to be. Thus, for any $\rho \in [0,1]$, both signals are strong if $\underline{\alpha}(\rho) < \alpha < \overline{\alpha}(\rho)$. Since for $\rho > 0$, $\underline{\alpha}(\rho) < \varepsilon$ and $\overline{\alpha}(\rho) > 1-\varepsilon$, it is true that the range where both signals are strong is strictly larger with positive correlation. Though obviously positive correlation decreases the information available to the voter in a separating equilibrium, nothing changes conceptionally for the voter. Therefore, in a separating PSE the same two conditions have to be simultaneously satisfied, which will generically not be the case. Therefore:

Proposition 17 *For $\alpha \in (\underline{\alpha}(\rho), \overline{\alpha}(\rho))$, the generically unique PSE outcome is $\tau(\alpha)$. Moreover, $\underline{\alpha}(\rho)$ decreases and $\overline{\alpha}(\rho)$ increases in ρ, so the range for which $\tau(\alpha)$ is the generically unique PSE outcome is larger, the larger the signals' correlation.*

An open issue not yet solved is whether there are separating PSE for $\alpha \notin (\underline{\alpha}(\rho), \overline{\alpha}(\rho))$. A natural conjecture is that the answer is affirmative.

B.2 Asymmetric Signals

Assume now that parties differ with respect to the quality of their signals. Let ε_i be the error probability of party i's signal and assume wlog $\varepsilon_1 < \varepsilon_2 < 1/2$. Assume that the voter still employs a symmetric strategy in the sense that $\gamma(\tau'', \tau') = 1 - \gamma(\tau', \tau'')$ for all τ', τ'', where it may be recalled that $\gamma(\tau'', \tau')$ is the probability that she votes for party when party 1 plays τ'' and party 2 plays τ'. Given the asymmetry in signal technology, this symmetric strategy assumption is certainly less natural than in the model where parties' signal are of identical quality. However, maintaining the symmetric strategy assumption is, obviously, the only way to isolate the effect of departing from the symmetric signals assumption.

Proposition 11 implies immediately that both parties playing $\tau(\alpha)$ independently of their signal is still the unique pooling PSE outcome. Letting τ_i^k be the policy party i plays on equilibrium after signal k, the quality difference of the signals implies $\mu(\tau_1^a, \tau_2^b) > \mu(\tau_1^b, \tau_2^a)$. Upon the signal k both parties must play the same policy τ_k[35] Together with the concavity of u this implies (i) that $u(\mu(\tau_a, \tau_b), \tau_a) \leq u(\mu(\tau_a, \tau_b), \tau_b)$ implies $u(\mu(\tau_b, \tau_a), \tau_b) > u(\mu(\tau_b, \tau_a), \tau_a)$ and (ii) $u(\mu(\tau_a, \tau_b), \tau_a) \geq u(\mu(\tau_a, \tau_b), \tau_b)$ implies $u(\mu(\tau_b, \tau_a), \tau_b) < u(\mu(\tau_b, \tau_a), \tau_a)$. Therefore, it is impossible to satisfy condition (4) in equilibrium. In case (i) policy τ_b is said to be strong while in case (ii) policy τ_a is said to be strong. (This is not related to and not to be confounded with weak and strong signals.) Whenever party 1 plays a separating equilibrium strategy, party 2 will therefore pool on the strong policy. Without loss of generality assume that this is policy τ_b. Voter indifference thus requires that τ_a and τ_b satisfy

$$u(\mu(\tau_a, \tau_b), \tau_a) = u(\mu(\tau_a, \tau_b), \tau_b). \tag{24}$$

On the other hand, for party 1 not to want to deviate from its prescribed equilibrium policy, PSE requires that $\tau_a = \arg\max_\tau u(\mu(a, 0), \tau)$ and $\tau_b = \arg\max_\tau u(\mu(b, 0), \tau)$ if both signals are strong. Thus, equations (24), $\tau_a = \arg\max_\tau u(\mu(a, 0), \tau)$ and $\tau_b = \arg\max_\tau u(\mu(b, 0), \tau)$ impose three independent restriction that will generically not be satisfied. On the other hand, if one signal is weak from the perspective of party 1, then there are only two restrictions: (24) and $\tau_a = \arg\max_\tau u(\mu(a, 0), \tau)$ if b is weak or $\tau_b = \arg\max_\tau u(\mu(b, 0), \tau)$ if a is weak. Thus, if only one signal is strong, then a PSE generically exists where the party with the high quality signal reveals its information. The usual algebra reveals that signal $a(b)$ is still strong for party i iff $\varepsilon_i < \alpha$ ($\alpha < 1 - \varepsilon_i$).

[35] ARGUMENT PERHAPS TO BE EXTENDED: If not, there would be a profitable deviation (since policies would have to be on either side of the hump...).

References

Austen-Smith, D. (1990): "Information Transmission in Debate," *American Journal of Political Science*, 34, 124152.

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

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

——— (2004): "Policy Advice with Imperfectly Informed Experts," *Advances in Theoretical Economics*, Article 1.

Berentsen, A., E. Bruegger, and S. Loertscher (2008): "Learning, public goods provision and the information trap," *Journal of Public Economics*, 92(5-6), 415–436.

Bernhardt, D., J. Duggan, and F. Squintani (2007): "Electoral Competition with privatley-informed candidates," *Games and Economic Behavior*, 58, 1–29.

Bhaskar, V., and E. van Damme (2002): "Moral hazard and private monitoring," *Journal of Economic Theory*, 102, 16–39.

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

——— (2009): "Searching for Good Policies," *Working Paper*.

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

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.

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.

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.

Kartik, N. (2008): "Strategic Communication with Lying Costs," *Working Paper [RES forth].*

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

Krishna, V., and J. Morgan (2001a): "Asymmetric Information and Legislative Rules: Some Amendments," *American Political Science Review*, 95(2), 435–452.

——— (2001b): "A Model of Expertise," *Quarterly Journal of Economics*, pp. 747–775.

Laslier, J.-F., A. Trannoy, and K. Van Der Straeten (2003): "Voting under ignorance of job skills of unemployed: the overtaxation bias," *Journal of Public Economics*, 87(3-4), 595–626.

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

Majumdar, S., and S. Mukand (2004): "Policy Gambles," *American Economic Review*, 94(4), 1207–22.

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

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.

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

Strulovici, B. (2008): "Learning While Voting: Determinants of Collective Experimentation," *Working Paper*.

Vives, X. (2005): "Complementarities and Games: New Developments," *Journal of Economic Literature*, XLIII(2), 437–79.

Wittman, D. (1989): "Why Democracies produce efficient results," *Journal of Political Economy*, pp. 1395–1424.