University of Toronto
Department of Economics

Working Paper 403

Influential Opinion Leaders

By Jakub Steiner and Colin Stewart

April 16, 2010

Influential Opinion Leaders*

Jakub Steiner† — Northwestern University

Colin Stewart‡ — University of Toronto

April 9, 2010

Abstract

We present a simple model of elections in which experts with special interests endorse candidates and endorsements are observed by the voters. We show that the equilibrium election outcome is biased towards the experts' interests even though voters know the distribution of expert interests and account for it when evaluating endorsements. Expert influence is fully decentralized in the sense that individual experts have no incentive to exert influence. The effect arises when some agents prefer, ceteris paribus, to support the winning candidate and when experts are much better informed about the state of the world than are voters.

1 Introduction

In the lead up to elections, many experts make public recommendations about which candidate to vote for. Do the experts' interests influence mass opinion and behavior? We show that expert endorsements can have a large effect on election outcomes, biasing the results toward their own interests. Our model features Bayesian voters who know the distribution of expert biases and a large number of experts who have negligible individual influence. The effect arises as a result of herding multiplied through a coordination motive. We show that the total effect can be large even if the direct herding effect small.

Well-informed agents convey recommendations to the general electorate through at least two institutionalized channels. First, the media customarily publish editorials endorsing particular

*We thank David Austen-Smith, Mehmet Ekmekci, Stephen Morris, and Alessandro Pavan for helpful comments. The second author is grateful to SSHRC for financial support of this research.

†email: j-steiner@kellogg.northwestern.edu

‡email: colin.stewart@utoronto.ca

political candidates. Second, campaign donors endorse candidates indirectly through political advertisements financed from donations. We examine elites' influence in a model in which first each expert endorses one candidate and then voters, each observing a random sample of endorsements, elect a candidate by majority rule. The election outcome is biased toward experts' interests under two assumptions: (i) experts are much more informed than are voters about the distribution of preferences, and (ii) some experts and voters prefer, ceteris paribus, to support the winner of the election.

Despite empirical evidence to the contrary,[1] economic models of elections have typically assumed that voters cannot be systematically manipulated by expert endorsements except insofar as their beliefs about candidate type are affected. Indeed, systematic manipulation of decision makers by experts' interests may appear to be at odds with rational choice. A Bayesian decision-maker accounts for experts' biases when evaluating their advice, potentially offsetting the experts' influence. For instance, in the cheap talk literature the bias of the informed agent typically results in a limitation on credible communication rather than consistent manipulation of the principal.

In this paper we identify a novel channel through which experts' biases, despite being known, influence the election outcome. In our model, because some agents prefer to support the winner, the election is a coordination game in which optimal actions depend on agents' beliefs about the election outcome. There is a continuum of experts and voters, each of whom possesses private information about the relative strength of the candidates (based on the number of partisan voters on each side). As in the global games literature, private information ensures that there is a unique equilibrium. In equilibrium, the election outcome is determined by the candidates' relative strength and the agents' biases, enabling us to quantify the influence of experts' preferences.

Each expert is assumed to have an intrinsic bias toward endorsing a particular candidate, which she trades off against her preference for endorsing the winner. Since experts' beliefs are based on their private information, endorsements provide information to voters about the future election outcome, thereby influencing those voters who also prefer to support the winner. The experts become opinion leaders. Even though voters are aware of differences between their own biases and those of the experts, their beliefs about the winner are affected by the experts' biases, at least

[1] See Page, Shapiro, and Dempsey (1987), Beck, Dalton, Greene, and Huckfeldt (2004), Druckman and Parkin (2008), DellaVigna and Kaplan (2007), and Gabel and Scheve (2007).

sometimes. While the beliefs of rational voters cannot be manipulated systematically, they can be "fooled" in some contingencies. The set of contingencies in which the voters fail to filter out the experts' biases turns out to be small, but these happen to be contingencies that are pivotal for the equilibrium outcome.

Consider independent experts and independent voters whose biases are weak enough that, if they were certain of who would win the election, they would prefer to support the winner. In particular, if there is little doubt about who will win, the optimal actions of agents in these two groups are aligned. An agent's optimal action depends on her bias only when she is uncertain about the outcome. Since experts are well-informed, voters believe that each expert is likely to be certain of the outcome. Hence voters effectively ignore experts' biases when evaluating their endorsements. Consequently, in contingencies so close to a tie that experts are unsure of the outcome, the distribution of endorsements is biased toward experts' interests, which in turn biases the vote in the same direction.

Even though contingencies in which experts are uncertain are rare when experts are well informed, strategic complementarities can multiply the effect so as to make the candidate preferred by experts considerably more likely to win. Starting from an equilibrium of the simultaneous voting game without experts, introducing experts leads to more votes for their preferred candidate in contingencies where the election would otherwise be very close. This in turn leads to more endorsements of that candidate in other nearby contingencies, which generates more votes, and so on, multiplying the effect. The size of the effect at each step vanishes as experts become very well informed, but the total effect is generally non-vanishing.

We explicitly characterize the equilibrium of the game with experts and voters. The characterization shows that the presence of experts generally affects the likelihood that each candidate wins. The influence of experts is monotone in the sense that, if experts' biases shift in favor of one candidate, that candidate becomes more likely to win the election. Moreover, the influence of experts is large when voters observe many endorsements; we show that in the limit as the number of observed endorsements tends to infinity, the candidate preferred by the experts always wins whenever partisan voters for the opposing candidate do not form a majority.

The key assumption that agents favor the winner of an election may hold for various reasons. In the case of experts, endorsements are public, and can therefore affect relationships with the elected politician (for example through access). For voters, a preference for the winner may be due to conformism as in Callander (2007) and Callander (2008).[2] In elections to select a party leader, conformism arises naturally from a desire to keep the party united for the general election. Alternatively, a coordination motive may arise among voters in elections with multiple candidates if a majority with preferences split between two similar candidates need to coordinate to defeat a Condorcet loser (see Cox (1997), Forsythe, Myerson, Rietz, and Weber (1993), and Myatt (2007)).

Aside from elections, our framework can be applied to many other economic settings that naturally exhibit an incentive to "vote" for the winner. For example, in the adoption of technologies with network externalities (Katz and Shapiro 1986), if potential adopters observe early choices of a few well informed experts, then the model suggests that the equilibrium coordination outcome disproportionately reflects the experts' preferences. In particular, the price at which the good is offered to early adopters can have a large impact on eventual adoption even if later buyers know the past prices and market participants have good information about the quality of the good. In addition, in line with the marketing literature (e.g. Watts and Dodds (2007)), opinion leaders become natural marketing targets.

The view that rationality limits manipulation by experts has been common in the political economy literature. Coate (2004) and Prat (2006) emphasize that rational voters can account for the interests behind expensive political campaigns. In a similar vein, DellaVigna and Kaplan (2007) interpret their empirical evidence of media impact as either a temporary phenomenon that will disappear as biases are learned, or as a consequence of irrationality among voters. To the extent that opinion manipulation has appeared in political economy modelling, it has generally been assumed in an ad hoc form lacking explicit foundations (see, e.g., Shachar and Nalebuff (1999), Grossman and Helpman (2002), and Murphy and Shleifer (2004)). One exception is Ekmekci (2009), who shows how a single expert with a known bias can manipulate an election by acting as a coordination device among voters. In our model, experts cannot act as coordination devices because each expert's endorsement is observed by a negligible fraction of voters.

The global games literature provides several insights on the role of social learning in coordination processes. If social learning is public then the observation of early actions correlates the beliefs

[2] See Callander (2008) for a review of the literature. Callander traces the use of conformism in voting theory back to Hinich (1981) and cites psychological literature beginning with Asch (1951) that documents conformism empirically.

of late movers. This correlation can lead to equilibrium multiplicity as in Angeletos, Hellwig, and Pavan (2007), or to a disproportionate influence of one large opinion leader as in Corsetti, Dasgupta, Morris, and Shin (2004). Social learning in our model is private, thereby preserving the informational heterogeneity needed to ensure equilibrium uniqueness.[3] Edmond (2007) studies strategic information transmission between citizens and government with conflicting preferences and finds that the government can manipulate the actions of rational citizens. Both Corsetti, Dasgupta, Morris, and Shin (2004) and Edmond (2007) study the influence of one large player who internalizes the impact of her action. In our model, the equilibrium reflects the experts' interests even though experts have negligible individual influence.

2 Model

A continuum of voters elect candidate A or B by majority rule. Two thirds of voters are partisans who always vote for their preferred candidate independent of others' actions. The remaining third of voters are independent voters whose preferred candidate depends on their expectations about the election outcome. Independent voters, indexed by $i \in [0, 1/3]$, vote according to two possibly conflicting criteria: they have an intrinsic preference for one of the candidates but also prefer to vote for the winner. Let $a^i \in \{A, B\}$ be the vote of independent voter i, let $w \in \{A, B\}$ be the winner of the election, and let b_v^i be a parameter in $(-1, 1)$ capturing the bias of independent voter i with b_v^i a measurable function of i. The payoff to independent voter i is[4]

$$u_v(a^i, b_v^i, w) = b_v^i \mathbb{1}_{a^i=A} + \mathbb{1}_{a^i=w}.$$

Independent voter i receives a premium of b_v^i if she votes for candidate A and a premium of 1 if she votes for the winner.

Let s_A and s_B be the measures of partisan voters supporting candidates A and B respectively, with $s_A, s_B \in [0, 2/3]$ and $s_A + s_B = 2/3$, and let θ be the proportion of independent votes for

[3]See Dasgupta (2007) for the pioneering study of private social learning in global games.

[4]Note that only the difference in payoffs between a voter's two actions is consequential for equilibrium actions. In particular, since an individual voter has no effect on the electoral outcome, adding an additional term to the payoffs that depends only on who wins the election would have no effect on equilibrium behavior.

candidate A that would lead to a tie; that is, θ is defined by

$$s_A + \frac{\theta}{3} = s_B + \frac{1-\theta}{3}.$$

Candidate A wins if the proportion of independent votes for A exceeds θ; otherwise B wins.[5] We refer to θ as the state. If $\theta < 0$ then more than half of all voters are partisans supporting candidate A, and hence A wins regardless of the independent votes. Similarly, if $\theta > 1$ then candidate B wins. We focus on the election outcome for $\theta \in (0, 1)$, where it depends on the independent voters' actions. In this region, the preference to vote for the winner creates a coordination problem among the independent voters. If, for $\theta \in (0, 1)$, θ was common knowledge among the independent voters, the game would have multiple equilibria with each candidate winning in some equilibrium.

Instead of complete information, voters receive private information about the state θ consisting of two parts: exogenous signals and endorsements from experts who have information about θ. Before describing the details of the information structure, we introduce the experts' payoffs.

There is a continuum of experts. Each expert j casts an endorsement $a^j \in \{A, B\}$ in favor of one of the two candidates. As in the case of the voters, one third of the experts are independents and two thirds are partisans. For simplicity, we assume that partisan experts are equally divided between the candidates: half of the partisan experts always endorse candidate A, while the other half always endorse candidate B.[6] The remaining third of experts are independent experts, indexed by $j \in [0, 1/3]$, whose optimal actions depend on their expectations about the election outcome.

Independent experts' payoffs are similar to those of independent voters. Let $a^j \in \{A, B\}$ be the endorsement by independent expert j, and let b_e^j be a constant in $(-1, 1)$ capturing the bias of independent expert j with b_e^j a measurable function of j. The payoff to independent expert j is

$$u_e(a^j, b_e^j, w) = b_e^j \mathbb{1}_{a^j=A} + \mathbb{1}_{a^j=w}.$$

As for independent voters, independent expert j receives a premium of b_e^j if she endorses candidate A and an additional premium of 1 if she endorses the winner.

[5] The tie-breaking rule is irrelevant since ties occur with zero probability.

[6] The assumption that partisan experts are equally divided between the two candidates can be relaxed. With a different distribution, the signs of the comparative statics we study remain the same, but the magnitude of the effects may differ.

The information structure and timing are as follows. First, the state θ is drawn from a uniform distribution on $[-1/2, 3/2]$. Then each independent expert j receives a private signal $x^j = \theta + \sigma\xi^j$, and each independent voter i receives a private signal $z^i = \theta + \varepsilon^i$. The experts' errors ξ^j are drawn from a continuous distribution F with support on $[-1/2, 1/2]$ and density f, and the voters' errors ε^i from a continuous distribution G with support on $[-1/4, 1/4]$ and density g. Errors are independent across players and independent of θ. The parameter σ, which is assumed to lie in the interval $(0, 1/2]$, scales the noise in experts' signals. Our results focus on the limit as σ tends to 0.

After the signals have been observed and before the election, experts simultaneously choose endorsements a^j. In addition to her private signal z^i, each voter observes a random sample of n endorsements, where, for simplicity, $n \in \mathbb{N}$ is fixed across voters. The sample is private and taken with uniform probability over all experts, regardless of type. Voters do not observe the biases or signals of the experts in their sample. After observing endorsements, the voters simultaneously choose votes a^i.

Let $\lambda^i \in \{0, \ldots, n\}$ denote the number of endorsements of candidate A in voter i's sample. A strategy for an independent expert maps each signal x^j to an endorsement $a^j \in \{A, B\}$. A strategy for an independent voter maps each pair (z^i, λ^i) to a vote $a^i \in \{A, B\}$. A strategy for an expert is monotone if there is some threshold signal above which she endorses candidate B and below which she endorses candidate A. A strategy s^i for a voter is monotone if (i) $s^i(z, \lambda) = B$ implies that $s^i(z', \lambda') = B$ whenever $z' \geq z$ and $\lambda' \geq \lambda$, and (ii) $s^i(z, \lambda) = A$ implies that $s^i(z', \lambda') = A$ whenever $z' \leq z$ and $\lambda' \leq \lambda$. We restrict attention to monotone strategies. All parameters of the model, including biases, and all distributions are common knowledge.

3 Elections without Opinion Leaders

Before we solve the main model, we consider elections in which the voters do not observe the experts' endorsements (in the notation of Section 2, $n = 0$). In this case, the election reduces to a simultaneous move game among the independent voters. We derive a monotone Bayesian Nash equilibrium with a pivotal state θ^* such that candidate A wins for $\theta < \theta^*$ and B wins for $\theta > \theta^*$.

Given the threshold θ^*, let $\pi_v(z^i, \theta^*)$ be the posterior belief that voter i assigns to candidate A

winning after receiving the signal z^i; that is, let

$$\pi_v(z^i, \theta^*) = \Pr\left(\theta < \theta^* | z^i\right) = 1 - G\left(z^i - \theta^*\right).$$

Independent voter i votes for candidate A if and only if her posterior belief $\pi_v(z^i, \theta^*)$ exceeds a critical probability p_v^i, where p_v^i solves the indifference condition

$$b_v^i + p_v^i = 1 - p_v^i.$$

Note that $p_v^i = \frac{1-b_v^i}{2}$ reflects the independent voters' bias b_v^i.

By the definition of the pivotal state θ^*, the election results in tie when $\theta = \theta^*$. Since θ is defined to be equal to the share of independent votes leading to a tie, θ^* must equal the share of independent voters supporting candidate A in the pivotal state. The *pivotal condition* is given by

$$\theta^* = \Pr\left(\pi_v(z^i, \theta^*) > p_v^i \mid \theta^*\right), \tag{1}$$

where i is a uniformly drawn independent voter. The following lemma allows us to use the preceding condition to compute the pivotal state regardless of the distribution of noise in voters' signals.

Lemma 1. *Posterior beliefs in the pivotal state $\theta^* \in [0,1]$ are distributed uniformly on $[0,1]$ regardless of the noise distribution. Thus for any $p \in [0,1]$ and any i, we have*

$$\Pr\left(\pi_v(z^i, \theta^*) > p \mid \theta^*\right) = 1 - p.$$

The uniform property of posterior beliefs in the lemma has been used in Guimaraes and Morris (2007) and Steiner (2006). For convenience, we include the proof in the appendix.

Using the lemma, the pivotal condition (1) implies that

$$\theta^* = \frac{1}{1/3}\int_0^{1/3} \left(1 - p_v^i\right) di = \frac{1}{1/3}\int_0^{1/3} \left(1 - \frac{1-b_v^i}{2}\right) di = 1 - \frac{1-b_v}{2} = \frac{1}{2} + \frac{b_v}{2},$$

where b_v denotes the average bias among independent voters. The election outcome without expert endorsements aggregates the preferences of voters in a natural way. Candidate A wins if she has sufficient support among partisan voters and/or independent voters are sufficiently biased in her favor. Moreover, the candidate preferred by independent voters is, ex ante, more likely to win the election.

The channel through which the independent voters' bias b_v affects the outcome is best understood through the pivotal condition (1). The pivotal state θ^* is determined by the best responses of the independent voters *in the pivotal state*. In this state, the independent voters receive inconclusive signals making them unsure about the election outcome. Consequently, their individual voting behavior is affected by their individual biases b_v^i, and the aggregate vote is a function of the average bias b_v. The analysis in the next section, where voters observe expert endorsements, also focuses on behavior in the pivotal state, where there is large strategic uncertainty. In the pivotal state, the behavior of experts who are uncertain about the outcome of the election is affected by their intrinsic biases and it turns out that voters do not filter out the experts' biases. As a result, the experts' biases affect the equilibrium outcome.

4 Elections with Opinion Leaders

We now return to the model of Section 2 in which each voter observes a random sample of $n > 0$ expert endorsements. Since we do not use any other equilibrium concept, we write simply "equilibrium" to mean weak perfect Bayesian equilibrium. As in the benchmark game, we restrict attention to equilibria in monotone strategies.

As above, any monotone equilibrium gives rise to a pivotal state θ^*, such that candidate A wins for $\theta < \theta^*$ and candidate B wins for $\theta > \theta^*$. The equilibrium analysis below has the same structure as the analysis of the benchmark game. We take the value of θ^* as given, compute the best responses of both the experts and the voters to θ^*, and then use the requirement that in the pivotal state the election results in a tie.

4.1 Experts' Behavior

We begin by considering the best responses of experts. Given the threshold θ^*, independent expert j endorses candidate A if and only if her posterior belief $\pi_e(x^j, \theta^*, \sigma)$ that $\theta < \theta^*$ exceeds a critical probability p_e^i. The critical probability again satisfies the indifference condition

$$b_e^i + p_e^i = 1 - p_e^i. \tag{2}$$

Note that $p_e^i = \frac{1-b_e^i}{2}$ reflects the experts' bias b_e^i.

Let $l(\theta, \theta^*, \sigma)$ denote the probability that a randomly chosen expert endorses candidate A in state θ given the threshold θ^*. Taking into account both partisan and independent experts, we have

$$l(\theta, \theta^*, \sigma) = \frac{1}{3} + \frac{1}{3} \Pr\left(\pi_e(x^j, \theta^*, \sigma) > p_e^j \mid \theta\right),$$

where j is a random independent expert.

The analysis of experts' behavior is particularly simple if the realized state θ is sufficiently far from the pivotal state θ^* relative to the noise in the experts' signals. In that case, every independent expert correctly forecasts and endorses the winner. Thus we have

$$l(\theta, \theta^*, \sigma) = \begin{cases} \frac{1}{3} & \text{if } \theta > \theta^* + \sigma, \\ \frac{2}{3} & \text{if } \theta < \theta^* - \sigma. \end{cases}$$

The analysis of the experts' behavior is also simple when the realized state θ is exactly equal to the pivotal state θ^*. By Lemma 1, experts' posterior beliefs are uniformly distributed on $[0,1]$ in the pivotal state θ^*. Therefore, the ex ante probability that independent expert i endorses candidate A in state θ^* is $1 - p_e^i$. Combining this observation with (2) gives the following lemma.

Lemma 2. *Let b_e denote the average bias of independent experts. For any $\theta^* \in [0,1]$ and any $\sigma > 0$, we have $l(\theta^*, \theta^*, \sigma) = \frac{1}{2} + \frac{b_e}{6}$. In particular, in the pivotal state, the share of experts endorsing candidate A is strictly increasing in b_e and is independent of θ^* and σ.*

We have made two observations (i) in typical states—those outside a σ-neighborhood of θ^*—the experts' bias does not influence the distribution of endorsements, and (ii) in the pivotal state the experts' support of candidate A increases with their bias. Both observations are important for the analysis of voters' behavior. Because of (i), at least when σ is small, voters effectively neglect the experts' biases when updating their beliefs based on endorsements. Because of (ii), the experts' biases affect voters' decisions in the pivotal state. Since the equilibrium is determined by the voters' behavior in the pivotal state, the equilibrium outcome depends on the experts' bias even though b_e is commonly known and voters correctly account for it when forming beliefs.

4.2 Voters' Behavior

Next we analyze voters' behavior. Let $p_v(z, \lambda, \theta^*, \sigma)$ denote the posterior probability that a voter assigns to candidate A winning the election after observing a signal z and a number λ of endorsements for A (given the threshold θ^*). We have

$$p_v(z, \lambda, \theta^*, \sigma) = \frac{\int_{-1/2}^{\theta^*} g(z-\theta) \Pr(\lambda|\theta) d\theta}{\int_{-1/2}^{3/2} g(z-\theta) \Pr(\lambda|\theta) d\theta}. \tag{3}$$

The distribution of endorsements $\Pr(\lambda|\theta)$ depends on the realized state θ and on the experts' behavior. Conditional on θ, λ is binomially distributed with parameters n and $l(\theta, \theta^*, \sigma)$.

Let $v(\theta, \theta^*, \sigma) \in [0, 1]$ denote the share of independent votes for candidate A in state θ when all players play best responses given θ^*. We have

$$v(\theta, \theta^*, \sigma) = \Pr\left(p_v(z^i, \lambda^i, \theta^*, \sigma) > p_v^i \mid \theta\right).$$

As in the benchmark game, θ^* must satisfy the condition

$$\theta^* = v(\theta^*, \theta^*, \sigma), \tag{4}$$

which states that, in the pivotal state θ^*, candidate A receives exactly the right proportion $v(\theta^*, \theta^*, \sigma)$ of independent votes as to make the election result in a tie.

Due to the symmetry of the model with respect to θ, $v(\theta^*, \theta^*, \sigma)$ is independent of θ^*. It follows that the pivotal state is uniquely determined.

Proposition 1. *The voting game has a unique monotone equilibrium.*

The proof is in the appendix.

From now on we focus on the limit as $\sigma \to 0$, in which experts' signals $x^i = \theta + \sigma\xi^i$ are much more precise than voters' signals. In this limit, voters' posterior beliefs are relatively simple to compute. As in the benchmark game, let $\pi_v(z^i, \theta^*) = \Pr(\theta < \theta^*|z^i)$ denote the probability that candidate A wins evaluated by voter i using only her private signal z^i.

Lemma 3. *For every z, λ, and θ^*, we have*

$$\lim_{\sigma \to 0^+} p_v(z, \lambda, \theta^*, \sigma) = \frac{\pi_v(z, \theta^*)}{\pi_v(z, \theta^*) + (1 - \pi_v(z, \theta^*))\, 2^{n-2\lambda}}. \tag{5}$$

According to the lemma, in the limit, voters treat the experts' endorsements as informative signals but ignore experts' incentives when evaluating these signals. In particular, the posterior belief increases in the number λ of endorsements for A, but is independent of the experts' bias b_e. In the limit, two thirds of experts endorse A whenever $\theta < \theta^*$ and two thirds of experts endorse B whenever $\theta > \theta^*$. Equation (5) follows from the usual Bayesian updating formula. The proof of Lemma 3 is in the appendix.

Next we characterize the pivotal state in the limit as $\sigma \to 0^+$ using condition (4). Lemma 1 implies that $\pi_v(z^i, \theta^*)$ is uniformly distributed on $[0, 1]$, and Lemma 2 determines the distribution of endorsements in the pivotal state. Finally, Lemma 3 describes voters' beliefs in the limit. Combining these lemmas yields the following proposition.

Proposition 2. *As $\sigma \to 0^+$, the pivotal threshold θ^* converges to*

$$\theta^{**} = \Pr\left(\frac{\pi}{\pi + (1-\pi)2^{n-2\lambda}} > p_v^i\right),$$

where π, λ and i are independent random variables with $\pi \sim U[0, 1]$, $\lambda \sim B\left(n, \frac{1}{2} + \frac{b_e}{6}\right)$ and $i \sim U[0, 1/3]$.[7]

The main results of this note follow from this proposition. First, the proposition implies that the experts' bias b_e has an unambiguous effect on the pivotal state θ^{**}. Since the distribution of λ is increasing in b_e (in the sense of first-order stochastic dominance), and the posterior belief $\frac{\pi}{\pi+(1-\pi)2^{n-2\lambda}}$ is increasing in λ, the characterization of Proposition 2 leads to the following corollary, which indicates that candidate A becomes more likely to win if experts' biases shift in her favor.

Corollary 1. *The pivotal threshold θ^{**} is strictly increasing in the experts' average bias b_e.*

The impact of the experts' bias grows when the number n of observed endorsements becomes large. Let $\theta^{**}(n)$ denote the pivotal threshold θ^{**} in the limit as $\sigma \to 0^+$ when each voter observes

[7] Here $B(n, p)$ denotes the binomial distribution for n draws with probability p.

n endorsements. Consider $\lim_{n\to\infty} \theta^{**}(n)$, corresponding to the equilibrium outcome in the ordered limit in which first $\sigma \to 0^+$ and then $n \to \infty$. As the following corollary indicates, the outcome of the election takes a very simple form in this limit. Unless partisan voters in favor of one candidate form a majority, the candidate favored by experts always wins.

Corollary 2. *If the experts' average bias b_e is positive then* $\lim_{n\to\infty} \theta^{**}(n) = 1$*. If the experts' average bias b_e is negative then* $\lim_{n\to\infty} \theta^{**}(n) = 0$*.*

The proof is in the appendix.

When n is large, the biases of independent voters have no effect on the election outcome, which is determined entirely by the preferences of the independent experts. As n increases, the observed endorsements become increasingly reliable indicators of the election outcome. Thus even a small expert bias that slightly shifts the distribution of endorsements in the pivotal state forces the outcome toward the experts' bias.

5 Discussion

The influence of experts in our model results from a combination of herding and coordination. To clarify the roles that these two features play, consider a variant of the model with no coordination motive. Instead of the winner being determined by followers' votes, suppose that the pivotal state θ^* is exogenously fixed; candidate A wins whenever $\theta < \theta^*$ and otherwise B wins. As in the model with coordination, the optimal endorsement chosen by an independent expert depends on her bias only in contingencies in which she is uncertain of the outcome. Otherwise, the candidate she endorses is exactly the ex post optimal choice for every independent voter. When experts have very precise information about θ, contingencies in which they are uncertain are rare, and hence voters effectively neglect the independent experts' biases when evaluating endorsements. Consequently, when these contingencies arise, voting behavior depends on experts' biases. However, in the absence of a coordination motive, the ex ante probability that expert biases affect voting behavior vanishes as the precision of the experts' information increases.

When θ^* is determined endogenously by voting behavior, the effect of experts' biases is multiplied and does not vanish if experts have precise information. Consider the effect of a shift in expert bias in favor of candidate A. Starting from the original equilibrium value of θ^*, this shift generates more endorsements of A in the small neighborhood of θ^* in which experts may be uncertain of the election outcome. The increase in endorsements in turn leads to more votes for A in states close to θ^*, thereby increasing the pivotal threshold. Because of the coordination motive, the increase in the threshold leads to further endorsements and votes for A, multiplying the effect. Moreover, no matter how small is the direct herding effect, the desire to coordinate makes the overall effect non-vanishing.

In our model, voters know only the distribution of preferences of the experts, but do not know the preferences of any particular expert. This assumption is natural in settings such as campaign finance, where experts may be interpreted as donors to political campaigns, provided that the voters observe only political advertisements without a link to individual donors. In other settings, such as editorial bias, it may be more natural to assume that the voters have some knowledge of the experts' political preferences. If voters have perfect knowledge of each expert's bias, then our results do not hold. In this case, voters who observe conflicting endorsements from independent experts deduce that the state is close to the pivotal one, and are able to correct for experts' biases. If, however, voters observe only a noisy signal of each expert's preference, then results similar to ours continue to hold. As the experts become increasingly informed, voters again effectively neglect those states in which the experts are uncertain about the election outcome, believing that conflicting endorsements are more likely to be the result of partisan experts. Consequently, the outcome depends on the experts' biases.

Voters in our model have intrinsic incentives to vote for a particular candidate. This formulation differs from the purely instrumental voting found in pivotal voter models, where voters have preferences only over the outcome of the election. If voters have an incentive to vote for the winner and the number of voters is large, instrumental voting reduces to a special case of our model. When there are many voters, each voter assigns only negligible probability to being pivotal. Thus the incentive to vote according to one's personal bias is completely crowded out by the incentive to vote for the winner (see Callander (2007)). Thus if we replaced the intrinsic bias term for voters in our model with an instrumental payoff term, the model would reduce (since the population is infinite) to our model with $b_v = 0$.

The assumption that endorsements are privately observed by voters is not essential for our results. If instead all voters observe the same n endorsements (drawn at random from the continuum of experts), then the equilibrium is again unique and exhibits the same features as in the private case. Moreover, although the equilibria in the two cases involve different thresholds, they converge to the same limit as n grows large. This strongly suggests that experts can also exert influence over the outcome in intermediate cases where a given endorsement may be observed by many but not all voters (as may be natural for campaign advertising or editorials in the media). Note that drawing the public endorsements at random from a continuum precludes any signaling motive on the part of the experts. We conjecture that incorporating such a motive would only strengthen the influence of experts.[8]

A Proofs

Proof of Lemma 1. Posterior beliefs are given by $\pi_v(z^i, \theta^*) = 1 - G(z^i - \theta^*)$ for any $z^i \in [-1/4, 5/4]$. If $\theta^* \in [0, 1]$, then conditional on $\theta = \theta^*$ all realized signals are in $[-1/4, 5/4]$. Thus we have

$$\Pr\left(\pi_v(z^i, \theta^*) < p \mid \theta^*\right) = \Pr\left(G^{-1}(1-p) + \theta^* < z^i \mid \theta^*\right).$$

In the state θ^*, $z^i = \theta^* + \varepsilon^i$, and hence

$$\begin{aligned} \Pr\left(G^{-1}(1-p) + \theta^* < z^i \mid \theta^*\right) &= \Pr\left(G^{-1}(1-p) < \varepsilon^i\right) \\ &= 1 - G\left((G^{-1}(1-p)\right) \\ &= p, \end{aligned}$$

as needed. $\square$

Proof of Proposition 1. We prove that the distribution of $p_v(z^i, \lambda^i, \theta^*, \sigma)$ conditional on θ^* does not depend on θ^*. It follows that $v(\theta^*, \theta^*, \sigma) = \Pr\left(p_v > p_v^i \mid \theta^*\right)$ does not depend on θ^* and hence the pivotal state condition $\theta^* = v(\theta^*, \theta^*, \sigma)$ has a unique solution.

Let $q(\lambda^i, \theta)$ denote the probability that voter i observes exactly λ^i endorsements for A when

[8] Proposition 7 of Corsetti, Dasgupta, Morris, and Shin (2004) pertains to a model closely related to a variant of our model with one expert who has a signaling motive. In their setting, the expert exerts a large influence over the outcome.

the state is θ (given the experts' strategies). In the pivotal state θ^*, $z^i = \theta^* + \varepsilon^i$ and hence

$$p_v(z^i, \lambda^i, \theta^*, \sigma) = \frac{\int_{-1/2}^{\theta^*} g(\theta^* + \varepsilon^i - \theta) q(\lambda^i, \theta) d\theta}{\int_{-1/2}^{3/2} g(\theta^* + \varepsilon^i - \theta) q(\lambda^i, \theta) d\theta}.$$

Using the transformation $\Delta = \theta - \theta^*$ and recalling that G has support on $[-1/4, 1/4]$ gives

$$p_v(z^i, \lambda^i, \theta^*, \sigma) = \frac{\int_{\varepsilon^i - 1/4}^{0} g(\varepsilon^i - \Delta) q(\lambda^i, \theta^* + \Delta) d\Delta}{\int_{\varepsilon^i - 1/4}^{\varepsilon^i + 1/4} g(\varepsilon^i - \Delta) q(\lambda^i, \theta^* + \Delta) d\Delta}.$$

To prove that the last expression does not depend on θ^*, we show that, for each Δ, the distribution of λ^i conditional on the state being $\theta^* + \Delta$ does not depend on θ^*. The random variable λ^i is distributed according to the Binomial distribution $B\left(n, l\left(\theta^* + \Delta, \theta^*, \sigma\right)\right)$, and thus it suffices to prove that $l(\theta^* + \Delta, \theta^*, \sigma)$ does not depend on θ^*. Accordingly, note that

$$l(\theta^* + \Delta, \theta^*, \sigma) = \frac{1}{3} + \frac{1}{3} \Pr\left(\pi_e(x^j, \theta^*, \sigma) > p_e^j \mid (\theta = \theta^* + \Delta)\right),$$

where

$$\pi_e(x^j, \theta^*, \sigma) = \Pr(\theta < \theta^* \mid x^j) = 1 - F\left(\frac{x^j - \theta^*}{\sigma}\right).$$

Therefore, we have

$$\begin{aligned} l(\theta^* + \Delta, \theta^*, \sigma) &= \frac{1}{3} + \frac{1}{3} \Pr\left(x^j < \sigma F^{-1}\left(1 - p_e^j\right) + \theta^* \mid (\theta = \theta^* + \Delta)\right) \\ &= \frac{1}{3} + \frac{1}{3} \Pr\left(\sigma \xi^j < \sigma F^{-1}\left(1 - p_e^j\right) - \Delta\right), \end{aligned}$$

which does not depend on θ^*. □

Proof of Lemma 3. For any $\theta \neq \theta^*$, there exists sufficiently small σ such that in the state θ, all independent experts endorse the winner of the election. Thus, for $\theta > \theta^*$, $\lim_{\sigma \to 0^+} l(\theta, \theta^*, \sigma) = \frac{1}{3}$ and for $\theta < \theta^*$, $\lim_{\sigma \to 0^+} l(\theta, \theta^*, \sigma) = \frac{2}{3}$.

The integrand in both the numerator and the denominator of (3) is bounded and hence

$p_v(z, \lambda, \theta^*, \sigma)$ converges to

$$\frac{\int_{-1/2}^{\theta^*} g(z-\theta)\binom{n}{\lambda}\left(\frac{2}{3}\right)^{\lambda}\left(\frac{1}{3}\right)^{n-\lambda} d\theta}{\int_{-1/2}^{\theta^*} g(z-\theta)\binom{n}{\lambda}\left(\frac{2}{3}\right)^{\lambda}\left(\frac{1}{3}\right)^{n-\lambda} d\theta + \int_{\theta^*}^{3/2} g(z-\theta)\binom{n}{\lambda}\left(\frac{1}{3}\right)^{\lambda}\left(\frac{2}{3}\right)^{n-\lambda} d\theta}.$$

Straightforward algebraic manipulation shows that the last expression is equal to the right hand side of (5). □

Proof of Corollary 2. Rearranging the expression for the threshold in Proposition 2 yields

$$\theta^{**}(n) = \Pr\left(\left(\frac{\pi}{1-\pi}\frac{1-p_v^i}{p_v^i}\right)^{1/n} > 2^{1-2\lambda/n}\right), \tag{6}$$

where π, λ and i are independent random variables with $\pi \sim U[0,1]$, $\lambda \sim B\left(n, \frac{1}{2}+\frac{b_e}{6}\right)$ and $i \sim U[0, 1/3]$. Note that, on the one hand, the left-hand side of the inequality in (6) converges in probability to 1. On the other hand, since λ/n converges in probability to $1/2+b_e/6$, the right-hand side converges in probability to $2^{-b_e/3}$. The result follows since $2^{-b_e/3} < 1$ if $b_e > 0$ and $2^{-b_e/3} > 1$ if $b_e < 0$. □

References

Angeletos, G., C. Hellwig, and A. Pavan (2007). Dynamic global games of regime change: Learning, multiplicity, and the timing of attacks. *Econometrica 75*(3), 711–756.

Asch, S. (1951). Effects of group pressure upon the modification and distortion of judgments. In H. Guetzkow (Ed.), *Groups, Leadership and Men*, pp. 177–190. Carnegie Press.

Beck, P., R. Dalton, S. Greene, and R. Huckfeldt (2004). The social calculus of voting: Interpersonal, media, and organizational influences on presidential choices. *American Political Science Review 96*(1), 57–73.

Callander, S. (2007). Bandwagons and momentum in sequential voting. *Review of Economic Studies 74*(3), 653–684.

Callander, S. (2008). Majority rule when voters like to win. *Games and Economic Behavior 64*(2), 393–420.

Coate, S. (2004). Pareto-improving campaign finance policy. *American Economic Review 94*(3), 628–655.

Corsetti, G., A. Dasgupta, S. Morris, and H. Shin (2004). Does one Soros make a difference? A theory of currency crises with large and small traders. *Review of economic Studies 71*(1), 87–113.

Cox, G. (1997). *Making votes count: strategic coordination in the world's electoral systems.* Cambridge Univ Pr.

Dasgupta, A. (2007). Coordination and delay in global games. *Journal of Economic Theory 134*(1), 195–225.

DellaVigna, S. and E. Kaplan (2007). The Fox News Effect: Media Bias and Voting*. *The Quarterly Journal of Economics 122*(3), 1187–1234.

Druckman, J. and M. Parkin (2008). The impact of media bias: How editorial slant affects voters. *The Journal of Politics 67*(04), 1030–1049.

Edmond, C. (2007). Information manipulation, coordination and regime change. *NYU typescript*.

Ekmekci, M. (2009). Manipulation through political endorsements. *Journal of Economic Theory 144*(3), 1227–1248.

Forsythe, R., R. Myerson, T. Rietz, and R. Weber (1993). An experiment on coordination in multi-candidate elections: The importance of polls and election histories. *Social Choice and Welfare 10*(3), 223–247.

Gabel, M. and K. Scheve (2007). Estimating the effect of elite communications on public opinion using instrumental variables. *American Journal of Political Science*, 1013–1028.

Grossman, G. and E. Helpman (2002). *Special interest politics.* The MIT Press.

Guimaraes, B. and S. Morris (2007). Risk and wealth in a model of self-fulfilling currency attacks. *Journal of Monetary Economics 54*(8), 2205–2230.

Hinich, M. (1981). Voting as an act of contribution. *Public choice 36*(1), 135–140.

Katz, M. and C. Shapiro (1986). Technology adoption in the presence of network externalities. *The Journal of Political Economy 94*(4), 822–841.

Murphy, K. and A. Shleifer (2004). Persuasion in politics. *American Economic Review 94*(2), 435–439.

Myatt, D. (2007). On the theory of Strategic Voting. *Review of Economic Studies 74*(1), 255–281.

Page, B., R. Shapiro, and G. Dempsey (1987). What moves public opinion? *The American Political Science Review 81*(1), 23–43.

Prat, A. (2006). Rational voters and political advertising. *The Oxford handbook of political economy*, 50–63.

Shachar, R. and B. Nalebuff (1999). Follow the leader: Theory and evidence on political participation. *American Economic Review 89*(3), 525–547.

Steiner, J. (2006). Coordination in a mobile world. *CERGE-EI Working Paper 295.*

Watts, D. and P. Dodds (2007). Influentials, networks, and public opinion formation. *Journal of Consumer Research 34*(4), 441–458.