Influential Opinion Leaders*

Jakub Steiner[†] Colin Stewart[‡]

Northwestern University University of Toronto

January 31, 2012

Abstract

We present a two-stage coordination game in which early choices of experts with special interests are observed by followers who move in the second stage. We show that the equilibrium outcome is biased toward the experts' interests even though followers know the distribution of expert interests and account for it when evaluating observed experts' actions. Expert influence is fully decentralized in the sense that each individual expert has a negligible impact. The bias in favor of experts results from a social learning effect that is multiplied through a coordination motive. We show that the total effect can be large even if the direct social learning effect is small.

1 Introduction

When large groups of agents seek to coordinate their behavior, it is common for experts to make public recommendations about the best course of action. These experts may have interests that conflict with those of the agents who observe their recommendations. In light of this conflict, do the experts' interests influence mass opinion and behavior? We show that expert endorsements can have a large effect on outcomes, biasing the results toward their own interests. The effect arises even though our model features Bayesian decision-makers who know the distribution of expert biases.

*We thank David Austen-Smith, Mehmet Ekmekci, Yosh Halberstam, Stephen Morris, Alessandro Pavan, Carolyn Pitchik, and Matt Turner 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

One setting to which our model naturally applies is the adoption of network technologies. Consider the decision of whether to adopt a new technology or remain with the status quo when there are positive network externalities (Katz and Shapiro 1986). If potential adopters observe early choices of a few well informed experts, then our results indicate that the equilibrium coordination outcome disproportionately reflects the experts' preferences *after* accounting for the relative prices that experts face. In particular, offering a low price to a relatively small group of early adopters can lead to widespread inefficient adoption even if later buyers know the past prices and all market participants have good information about the quality of the good. In line with the marketing literature (e.g. Watts and Dodds (2007)), opinion leaders become natural marketing targets.

Systematic manipulation of decision-makers' actions 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 contrast, expert influence can arise naturally in models of social learning in which a follower observes choices made by experts whose preferences may differ from her own. We show how a coordination motive may multiply this social learning effect. Moreover, the multiplication can be so large as to create a sizeable total effect even when the direct social learning effect is vanishingly small.

We identify a novel channel through which experts' biases, despite being known, influence the coordination outcome. Agents' optimal actions depend on their beliefs about the majority action. There is a continuum of experts and followers, each of whom possesses private information about the relative popularity of the two available actions. As in the global games literature, private information ensures that there is a unique equilibrium and gives rise to strategic uncertainty. In equilibrium, the coordination outcome is determined by a combination of both the experts' and the followers' biases, enabling us to quantify the influence of experts' preferences.

Each expert is assumed to have an intrinsic bias toward endorsing a particular action, which she trades off against her preference for endorsing the majority action. Since experts' beliefs are based on their private information, endorsements provide information to followers about the future coordination outcome, thereby influencing those followers who also prefer to support the majority action. The model exhibits a subtle interplay between the beliefs and actions of experts and those of followers. Each expert, unable to affect the outcome on her own, treats the outcome as given (albeit uncertain). Each follower also treats the outcome as given, and accounts for the distribution of experts' biases when forming beliefs based on observed endorsements. We show that the interplay between experts and followers skews the outcome toward the average bias of experts. The experts become opinion leaders.

Even though followers are aware of differences between their own biases and those of the experts, their beliefs about the majority action become skewed in the direction of the experts' biases, at least sometimes. While the beliefs of rational followers cannot be manipulated systematically (in the sense that they are correct when averaged across states), they can be affected by experts' biases in some states. The set of states in which the followers fail to filter out the experts' biases turns out to be small, but these happen to be states that are pivotal for the equilibrium outcome.

Consider those experts and followers whose biases are weak enough that, if they were certain about the majority action, they would prefer to support it. In particular, if there is little doubt about the coordination outcome, the optimal actions of agents in these two groups are aligned. Their optimal action depends on their biases only when they are uncertain about the outcome. Since experts are well-informed, followers believe that each expert is likely to be certain of the outcome. Hence followers effectively ignore experts' biases when evaluating their endorsements. Now consider those rare contingencies in which the outcome is so close to a tie that experts are uncertain of the outcome. The distribution of endorsements is biased toward experts' interests in these contingencies. Followers, not knowing that experts are uncertain, ignore the effect of the experts' biases. Consequently, in these contingencies, their choices comply with experts' interests.

Even though contingencies in which well-informed experts cannot predict the majority action are rare when experts are well informed, strategic complementarities can multiply the effect so as to make the action preferred by experts considerably more likely to be adopted. Starting from an equilibrium of the coordination game without experts, introducing experts leads to more followers choosing the experts' preferred action in contingencies where the outcome would otherwise be very close to a tie. This in turn leads to more endorsements of that action in other nearby contingencies, with followers adopting it more often, and so on, multiplying the effect. The size of the effect at each step of this contagion vanishes as experts become very well informed, but the total effect is generally large.

We explicitly characterize the equilibrium of the game with experts and followers. The characterization shows that the presence of experts generally affects the likelihood of coordination on each of the two actions. The influence of experts is monotone in the sense that, if experts' biases shift in favor of one action, the agents are more likely to coordinate on that action. Moreover, the influence of experts is large when followers observe many experts' choices; we show that in the limit as the number of observed choices tends to infinity, the coordination outcome always complies with the experts' preferences.

The global games literature provides several insights into the role of social learning in coordination processes. If social learning is public then the observation of early actions correlates the beliefs 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.[1]

Corsetti, Dasgupta, Morris, and Shin (2004), Edmond (2007), and (Ekmekci 2009) identify effects of expert influence in models with one large player who can act as a coordination device for followers and who internalizes the impact of her action. In contrast, we focus on the case of many experts with negligible individual influence who cannot act as a coordination device, making the mechanism underlying expert influence quite different.

2 Applications

Our model applies to a wide class of settings that combine coordination with social learning. As noted above, this combination arises naturally in the adoption of network technologies, or, more generally, the choice between any goods with positive network externalities (Katz and Shapiro 1985). In this case, the experts in our model can be thought of either as well informed early adopters or as product reviewers who choose to endorse one product and receive some benefit from endorsing a product that ends up being widely adopted. Our results indicate that, even if the pool of experts is negligible relative to the size of the population, the outcome is biased toward their preferences. In particular, lowering the price at which a network good is sold to experts—or

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

giving publicly observed payments to product reviewers—can increase the likelihood of widespread adoption.

Political revolutions provide another natural application of our model. Under a repressive regime, many citizens may prefer to revolt if and only if the revolution is sufficiently likely to succeed. Similarly, opinion leaders who publicly endorse or oppose revolutionary action often have a strong incentive to match their position to the outcome. If the distribution of experts' preferences differs from that of the citizens, the outcome is influenced in favor of the experts' interests. By manipulating the population and preferences of experts, dictators can reduce the likelihood of revolutions even if citizens are perfectly aware of what manipulation is taking place.

Our model can also be applied to democratic elections if (i) well-informed experts convey recommendations to the general electorate, and (ii) some experts and voters prefer, ceteris paribus, to support the winner of the election. A number of empirical papers have documented evidence that experts influence voters' decisions. Page, Shapiro, and Dempsey (1987), Beck, Dalton, Greene, and Huckfeldt (2004), Druckman and Parkin (2008), and DellaVigna and Kaplan (2007) show that voters are affected by media bias, while Gabel and Scheve (2007) show how elites' opinions on European integration influenced popular opinion. Despite this evidence, economic models of elections have typically assumed that voters cannot be systematically manipulated by expert endorsements.

A key assumption of our model is that some agents have an incentive to coordinate with the majority. While this assumption is not standard in the voting literature, there is reason to think that it may hold in at least some elections. In the case of experts, endorsements can 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; given that a particular leader is selected, even party members who supported an opponent may prefer that the leader is elected with a larger vote share. 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)). Left-wing voters choosing between Gore

[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.

and Nader in the 2000 U.S. presidential election provide a clear example of this last motive. As one voter put it:

> "I'd like to vote for Nader, I really would. I like what he stands for. I like the questions he's raising. But I've got to vote for Gore. I'd feel horribly guilty if I woke up the day after the election and Bush won, and if I felt that my particular vote had in any way could have swayed things toward him."[3]

Experts described a similar tradeoff. Three major environmental organizations endorsed the Democratic Party candidate despite their "environmental disappointments with Gore."[4] One can view our model as a reduced form of the decisions of these voters and experts taking the vote for the third candidate (Bush in this case) as given by an exogenous distribution.

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. The core observation in Ekmekci is that a monopolistic expert has both the ability and the incentive to coordinate the expectations of the electorate. When, as in our model, there are many experts with different audiences, experts lack this power, making the mechanism underlying their influence quite different.

[3] Quote from Kevin McGrorty as reported in the New York Times, October 27, 2000, "The 2000 Campaign: The Green Party; Nader's Damage to Gore Most Evident in Oregon."

[4] The New York Times, November 3, 2000, "The 2000 Campaign: The Environment; On a Favorite Issue, Gore Finds Himself on a 2-Front Defense."

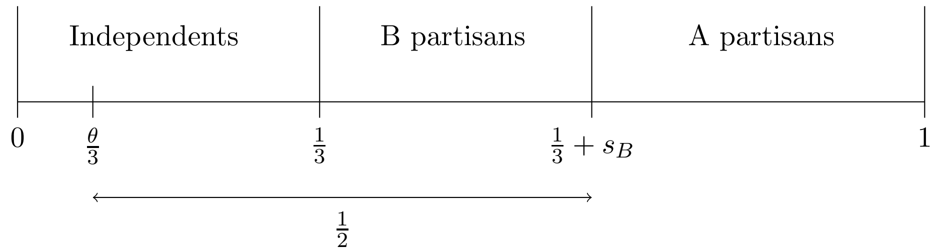

Figure 1: Illustration of the dependence of θ on proportion of partisans for each group.

3 Model

The model features two sets of players: experts and followers. A continuum of followers with unit measure each choose action A or B. Two thirds of followers are "partisans" who have a strong preference for one of the actions and choose it independently of others' actions. The preferred action of the remaining followers depends on their expectations about the outcome. These "independent followers", indexed by $i \in [0, 1/3]$, choose according to two possibly conflicting criteria: they have an intrinsic preference for one of the actions but also prefer to opt for the majority action.Let $a^i \in \{A, B\}$ be the choice of independent follower i, and let $b_f^i \in (-1, 1)$ be a parameter capturing the bias of independent follower i with b_f^i a measurable function of i. Finally, let $w \in \{A, B\}$ (for "winner") be the action chosen by the majority of the followers. The payoff to independent follower i is[5]

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

Independent follower i receives a premium of b_f^i if she chooses action A and a premium of 1 if she chooses the majority action.[6]

Let s_A and s_B be, respectively, the measures of followers for which action A or B is dominant,

[5]Note that only the difference in payoffs between a player's two actions is consequential for equilibrium behavior. In particular, since an individual player has no effect on the coordination outcome, adding an additional term to the payoffs that depends only on which action becomes the majority one would have no effect on equilibrium behavior.

[6]The focus on this payoffs that depend on others' actions only through the majority choice simplifies the analysis considerably but is not necessary for the main effect we identify (namely, the multiplication of the social learning effect through the coordination motive). We conjecture that similar results hold for a broad class of payoffs exhibiting strategic complementarities, including, for example, payoffs in which the second term above is replaced with a linear term equal to the fraction of followers choosing action A.

with $s_A, s_B \in [0, 2/3]$ and $s_A + s_B = 2/3$, and let θ be defined by

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

as illustrated in Figure 1. Action A becomes the majority action if the proportion of independent followers choosing A exceeds θ; otherwise B becomes the majority action.[7] We refer to θ as the state; a larger state is associated with greater partisan support for action B. If $\theta < 0$ then more than half of all followers find action A dominant. Similarly, if $\theta > 1$ then B becomes the majority action. We focus on the coordination outcome for $\theta \in (0, 1)$, where it depends on the independent followers' actions. In this region, the preference for the majority action creates a coordination problem among the independent followers.

In order to capture the idea that agents may be uncertain about the actions of others, we assume that experts and followers possess private information. More specifically, followers receive private information about the state θ consisting of two parts: exogenous signals and observations of experts' actions, who have information about θ. Before describing the details of the information structure, we introduce the experts' payoffs.

There is a continuum of experts with unit measure. Each expert j chooses $a^j \in \{A, B\}$. As described above, depending on the application, an expert's action can be interpreted either as an early choice or as an endorsement of a particular action. As in the case of followers, two thirds of experts have a strong enough preference for one action as to make that action dominant. For simplicity, we assume that these extreme preferences are equally divided between the two actions: one third of experts always choose A, while another third always choose B.[8] The remaining third of experts are independent experts, indexed by $j \in [0, 1/3]$, whose optimal actions depend on their expectations about the coordination outcome. Although we normalize the measure of both groups of agents to 1, the group of experts should be thought of as being small enough as to have a negligible direct impact on the majority action. This is reflected in the assumption that the majority action w is determined purely by the action choices of the followers.

Independent experts' payoffs are similar to those of independent followers. Let $a^j \in \{A, B\}$ be

[7] In case of a tie, the majority action may be chosen arbitrarily.

[8] The assumption that partisan experts are equally divided 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 action of expert j, and let $b_e^j \in (-1, 1)$ capture her bias, 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 with independent followers, independent expert j receives a premium of b_e^j for action A and an additional premium of 1 if she endorses the majority action.

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 follower i receives a private signal $z^i = \theta + \varepsilon^i$. For simplicity, the experts' errors ξ^j are drawn from a continuous distribution F with support $[-1/2, 1/2]$ and density f, and the followers' errors ε^i from a continuous distribution G with support $[-1/4, 1/4]$ and density g.[9] Errors are independent across agents 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. Agents do not know the realized distribution of preferences and derive their beliefs about it from their beliefs about θ.

After the signals have been observed and before followers choose actions, experts simultaneously choose their actions a^j. In addition to her private signal z^i, each follower observes a random sample of n expert actions, where, for simplicity, $n \in \mathbb{N}$ is fixed across followers. The sample is private and taken with uniform probability over all experts, regardless of type. Followers do not observe the biases or signals of the experts in their sample (see Section 6 for a discussion of the role of this assumption). After observing expert actions, the followers simultaneously choose actions a^i.

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

[9]The bounded support of the error terms simplifies exposition but is not necessary for the results.

are common knowledge.

4 Coordination without Opinion Leaders

Before we solve the main model, we consider coordination game in which the followers do not observe the experts' actions (in the notation of Section 3, $n = 0$). In this case, the game reduces to a simultaneous move game among the independent "followers". Using standard global games techniques, we derive the unique monotone Bayesian Nash equilibrium, which has the property that action A becomes the majority action only if partisan support for B does not exceed some critical level. In terms of the state θ, there exists a pivotal state θ_0^* such that A becomes the majority action for $\theta < \theta_0^*$ and B the majority action for $\theta > \theta_0^*$.

Given the threshold θ_0^*, let $\pi_f\left(z^i, \theta_0^*\right)$ be the posterior belief that follower i assigns to A being the majority action after receiving the signal z^i; that is, let

$$\pi_f\left(z^i, \theta_0^*\right) = \Pr\left(\theta < \theta_0^* | z^i\right) = 1 - G\left(z^i - \theta_0^*\right).$$

Independent follower i chooses A if and only if her posterior belief $\pi_f\left(z^i, \theta_0^*\right)$ exceeds a critical probability p_f^i, where p_f^i solves the indifference condition

$$b_f^i + p_f^i = 1 - p_f^i.$$

Note that $p_f^i = \frac{1-b_f^i}{2}$ reflects the bias b_f^i.

By the definition of the pivotal state θ_0^*, the outcome is a tie when $\theta = \theta_0^*$. Since θ is defined to be equal to the share of independent followers leading to a tie, θ_0^* must equal the share of independent followers choosing A in the pivotal state. The *pivotal condition* is given by

$$\theta_0^* = \Pr\left(\pi_f\left(z^i, \theta_0^*\right) > p_f^i \mid \theta_0^*\right), \tag{1}$$

where i is a uniformly drawn independent follower.

When the pivotal state is realized, followers' beliefs p_f^i reflect only the noise in their signals rather than useful information about the coordination outcome. As a result, their beliefs are diffuse in this contingency, as indicated by the following lemma.

Lemma 1. *Posterior beliefs in the pivotal state $\theta_0^* \in [0,1]$ are distributed uniformly on $[0,1]$ regardless of the noise distribution. Thus for any $p \in [0,1]$ and any i, we have*

$$\Pr\left(\pi_f\left(z^i, \theta_0^*\right) > p \mid \theta_0^*\right) = 1 - p. \tag{2}$$

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.

Integrating (2) across the population of independent followers, the pivotal condition (1) implies that

$$\theta_0^* = \frac{1}{1/3}\int_0^{1/3}\left(1 - p_f^i\right) di = \frac{1}{1/3}\int_0^{1/3}\left(1 - \frac{1-b_f^i}{2}\right) di = 1 - \frac{1-b_f}{2} = \frac{1}{2} + \frac{b_f}{2},$$

where b_f denotes the average bias among independent followers. In the absence of experts, the election outcome aggregates the preferences of followers in a natural way. Action A becomes the majority action if it is dominant for sufficiently many followers and/or independent followers are sufficiently biased in its favor. Moreover, the agents are, ex ante, more likely to coordinate on the action preferred by independent followers.

The channel through which the independent followers' bias b_f affects the outcome is best understood through the pivotal condition (1). The pivotal state θ_0^* is determined by the best responses of the independent followers *in the pivotal state*. In this state, the independent followers receive inconclusive signals making them unsure about the coordination outcome, thereby suppressing the significance of the coordination motive. Consequently, their individual choice is affected by their individual biases b_f^i, and the aggregate action is a monotone function of the average bias b_f. The analysis in the next section, where followers observe expert actions, also focuses on behavior in the pivotal state, in which there is considerable strategic uncertainty. In the pivotal state, the behavior of experts who are uncertain about the coordination outcome is affected by their intrinsic biases and it turns out that followers do not filter out the experts' biases. As a result, the experts' biases affect the equilibrium outcome.

5 Coordination with Opinion Leaders

We now return to the model of Section 3 in which each follower observes a random sample of $n \geq 1$ expert actions. We restrict attention to weak perfect Bayesian equilibria in monotone strategies.

As above, any monotone equilibrium gives rise to a pivotal state θ_n^*, such that A becomes the majority action for $\theta < \theta_n^*$ and B the majority action for $\theta > \theta_n^*$. The equilibrium analysis below has the same structure as the analysis of the benchmark game. We take the value of θ_n^* as given, compute the best responses of both the experts and the followers to θ_n^*, and then use the requirement that the outcome in the pivotal state is a tie.

5.1 Experts' Behavior

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

$$b_e^i + p_e^i = 1 - p_e^i. \tag{3}$$

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

Let $l\left(\theta, \theta_n^*, \sigma\right)$ denote the probability that a randomly selected expert chooses A in state θ given the threshold θ_n^*. Taking into account both partisan and independent experts, we have

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

where j is a randomly chosen independent expert.

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

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

The analysis of the experts' behavior is also simple in the pivotal state θ_n^*. By Lemma 1, experts' posterior beliefs are uniformly distributed on $[0,1]$ when $\theta = \theta_n^*$. Therefore, the ex ante probability that independent expert i chooses A in state θ_n^* is $1 - p_e^i$. Combining this observation with (3) gives the following lemma.

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

We have made two observations: (i) in typical states—those outside a σ-neighborhood of θ_n^*—the experts' bias does not influence the distribution of expert actions, and (ii) in the pivotal state the share of experts choosing A increases with their bias. Both observations are important for the analysis of followers' behavior below. Because of (i), when σ is small, followers effectively neglect the experts' biases when evaluating their actions; given followers' information, contingencies in which experts' biases affect expert choice are unlikely. Followers neglect the experts' biases even in the pivotal state in which, because of (ii), experts' biases do shape their actions. Since the equilibrium is determined by the followers' behavior in the pivotal state, the equilibrium outcome reflects the experts' bias even though b_e is commonly known and followers correctly account for it when forming beliefs.

5.2 Followers' Behavior

Next we analyze followers' behavior. Let $p_f(z, \lambda, \theta_n^*, \sigma) = \Pr_\sigma(\theta < \theta_n^* | z, \lambda)$ denote the posterior probability that a follower assigns to A becoming the majority action after observing a signal z and a number λ of experts choosing A (given the threshold θ_n^*). Bayes' rule gives

$$p_f(z, \lambda, \theta_n^*, \sigma) = \frac{\int_{-1/2}^{\theta_n^*} g(z-\theta) \Pr(\lambda|\theta) d\theta}{\int_{-1/2}^{3/2} g(z-\theta) \Pr(\lambda|\theta) d\theta}. \tag{4}$$

The distribution of observed experts' behavior $\Pr(\lambda|\theta)$ depends on the realized state θ and on the experts' strategies. Conditional on θ, λ is binomially distributed with sample size n and success probability $l(\theta, \theta_n^*, \sigma)$.

Let $v(\theta, \theta_n^*, \sigma) \in [0,1]$ denote the share of independent followers choosing A in state θ when all agents play best responses given θ_n^*. We have

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

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

$$\theta_n^* = v(\theta_n^*, \theta_n^*, \sigma), \tag{5}$$

which states that, in the pivotal state θ_n^*, action A is chosen by exactly the right proportion $v(\theta_n^*, \theta_n^*, \sigma)$ of independent followers as to make the outcome a tie.

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

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

The proof is in the appendix.

From this point on we focus on the limit as $\sigma \to 0^+$, in which experts' signals are much more precise than followers' signals. In this limit, followers' posterior beliefs are relatively simple to compute. Let $\pi_f\left(z^i, \theta_n^*\right) = \Pr\left(\theta < \theta_n^* | z^i\right)$ denote the "pre-expert" probability that action A prevails evaluated by follower i conditioning only on her private signal z^i (as opposed to the "post-expert" probability $p_f(z, \lambda, \theta_n^*, \sigma)$). For any $\theta \neq \theta_n^*$, and sufficiently small σ all independent experts choose the majority action. Thus, for $\theta > \theta_n^*$, $\lim_{\sigma\to 0^+} l(\theta, \theta_n^*, \sigma) = \frac{1}{3}$ and for $\theta < \theta_n^*$, $\lim_{\sigma\to 0^+} l(\theta, \theta_n^*, \sigma) = \frac{2}{3}$. Hence, using Bayes rule, $p_f(z, \lambda, \theta_n^*, \sigma)$ converges to

$$\frac{\pi_f(z, \theta_n^*) \binom{n}{\lambda} \left(\frac{2}{3}\right)^{\lambda} \left(\frac{1}{3}\right)^{n-\lambda}}{\pi_f(z, \theta_n^*) \binom{n}{\lambda} \left(\frac{2}{3}\right)^{\lambda} \left(\frac{1}{3}\right)^{n-\lambda} + (1 - \pi_f(z, \theta_n^*)) \binom{n}{\lambda} \left(\frac{1}{3}\right)^{\lambda} \left(\frac{2}{3}\right)^{n-\lambda}}.$$

Straightforward algebraic manipulation leads to the following lemma.

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

$$\lim_{\sigma\to 0^+} p_f(z, \lambda, \theta_n^*, \sigma) = \frac{\pi_f(z, \theta_n^*)}{\pi_f(z, \theta_n^*) + (1 - \pi_f(z, \theta_n^*))\, 2^{n-2\lambda}}. \tag{6}$$

According to the lemma, in the limit, followers treat the experts' choices as informative signals but ignore experts' incentives when evaluating these signals. In particular, the posterior belief increases in the number λ of experts' choices of action A, but is independent of the experts' bias b_e.

Next we characterize the pivotal state in the limit as $\sigma \to 0^+$ using condition (5). Lemma 1 implies that followers' beliefs $\pi_f\left(z^i, \theta_n^*\right)$ before observing experts' choices are uniformly distributed on $[0,1]$, and Lemma 2 determines the distribution of experts' actions in the pivotal state. Finally, Lemma 3 describes, in the limit, followers' beliefs after observing experts' choices. Combining these lemmas yields the following proposition.

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

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

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 $i \sim U[0,1/3]$.[10]

Our main results follow from this proposition. First, the proposition implies that the experts' bias b_e has an unambiguous effect on the pivotal state θ_n^{**}. 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, A becomes more likely to be the majority action if experts' biases shift in its favor.

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

The impact of the experts' bias grows when the number n of observed expert actions becomes large. 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 coordination outcome takes a very simple form in this limit. Unless followers with extreme preferences form a majority, the action favored by experts always prevails.

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.

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

When n is large, the biases of independent followers have no effect on the coordination outcome, which is determined entirely by the preferences of the experts. As the sample of observed experts' actions increases in size, followers view their samples as increasingly reliable indicators of the coordination outcome. Thus even a small expert bias that slightly shifts the distribution of expert actions in the pivotal state forces the outcome toward the experts' bias.

6 Discussion

The influence of experts in our model results from a combination of social learning and coordination. To clarify the roles that these two features play, consider the following variant of the model with no coordination motive: instead of the predominant action being determined by followers' choices, suppose that the pivotal state θ^* is exogenously fixed; independent agents prefer to choose A if and only if $\theta < \theta^*$. As in the model with coordination, the optimal action chosen by an independent expert depends on her bias only in a neighborhood of θ^* where she is uncertain of the optimal action. At all other θ, her preferred action is exactly the ex post optimal choice for every independent follower. When experts have very precise information about θ, contingencies in which they are uncertain are rare, and hence followers effectively neglect the independent experts' biases when evaluating experts' actions. Consequently, when these contingencies arise, followers tend to comply with experts' biases. However, in the absence of a coordination motive, the ex ante probability that expert biases affect followers' behavior vanishes as the precision of the experts' information increases.

When θ^* is determined endogenously by the followers' behavior, the effect of experts' biases is multiplied and does not vanish even if experts have precise information. Consider the effect of a shift in expert bias in favor of action A. Starting from the original equilibrium value of θ^*, this shift generates more expert choice of A in the small neighborhood of θ^* in which experts are uncertain of the outcome. The increase in expert choice of A in turn leads to followers choosing A more often in states close to θ^*, thereby increasing the pivotal threshold. Because of the coordination motive, the increase in the threshold leads to further increases in the number of agents choosing A, repeatedly multiplying the effect. Moreover, no matter how small is the direct social learning effect, the desire to coordinate makes the overall effect non-vanishing.

In our model, followers know only the distribution of preferences of the experts, but do not know the preferences of any particular expert. If followers have perfect knowledge of each expert's bias, then our results do not hold. In this case, followers who observe conflicting choices from independent experts deduce that the state is close to the pivotal one, and are able to correct for experts' biases. If, however, followers observe only a noisy signal of each expert's preference, then results similar to ours continue to hold. As the experts become increasingly informed, followers again effectively neglect those states in which the experts are uncertain about the coordination outcome, believing that conflicting expert actions are more likely to be the result of the presence of partisan experts in the sample. Consequently, the outcome depends on the experts' biases.

The assumption that experts' choices are privately observed by followers is not essential for our results. If instead all followers observe actions of the same n experts (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 the action of a given expert may be observed by many but not all followers (as is natural for marketing or political campaigns). Note that drawing the observed experts 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.[11]

A Proofs

Proof of Lemma 1. First note that $z^i \in [-1/4, 5/4]$ whenever $\theta =\in [0,1]$. Posterior beliefs are given by $\pi_f(z^i, \theta_0^*) = 1 - G(z^i - \theta_0^*)$ for any $z^i \in [-1/4, 5/4]$. If $\theta_0^* \in [0,1]$, then conditional on $\theta = \theta_0^*$ all realized signals are in $[-1/4, 5/4]$. Thus we have

$$\Pr\left(\pi_f(z^i, \theta_0^*) < p \mid \theta_0^*\right) = \Pr\left(G^{-1}(1-p) + \theta_0^* < z^i \mid \theta_0^*\right).$$

[11] 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.

In state θ_0^*, $z^i = \theta_0^* + \varepsilon^i$, and hence

$$\begin{aligned}\Pr\left(G^{-1}(1-p) + \theta_0^* < z^i \mid \theta_0^*\right) &= \Pr\left(G^{-1}(1-p) < \varepsilon^i\right) \\ &= 1 - G\left((G^{-1}(1-p)\right) \\ &= p,\end{aligned}$$

as needed. $\square$

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

Let $q(\lambda^i, \theta)$ denote the probability that follower i observes exactly λ^i endorsements for A when the state is θ (given that the experts' strategies are best responses to θ_n^*). In the pivotal state θ_n^*, $z^i = \theta_n^* + \varepsilon^i$ and hence

$$p_f(z^i, \lambda^i, \theta_n^*, \sigma) = \frac{\int_{-1/2}^{\theta_n^*} g(\theta_n^* + \varepsilon^i - \theta) q(\lambda^i, \theta) d\theta}{\int_{-1/2}^{3/2} g(\theta_n^* + \varepsilon^i - \theta) q(\lambda^i, \theta) d\theta} = \frac{\int_{\theta_n^* + \varepsilon^i - 1/4}^{\theta_n^*} g(\theta_n^* + \varepsilon^i - \theta) q(\lambda^i, \theta) d\theta}{\int_{\theta_n^* + \varepsilon^i - 1/4}^{\theta_n^* + \varepsilon^i + 1/4} g(\theta_n^* + \varepsilon^i - \theta) q(\lambda^i, \theta) d\theta},$$

where the latter equality follows since g has support on $[-1/4, 1/4]$ and the critical state θ_n^* must lie in $[0, 1]$. Using the transformation $\Delta = \theta - \theta_n^*$ gives

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

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

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

where

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

Therefore, we have

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

which does not depend on θ_n^*. □

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

$$
\theta_n^{**} = \Pr\left(\left(\frac{\pi}{1-\pi}\frac{1-p_f^i}{p_f^i}\right)^{1/n} > 2^{1-2\lambda/n}\right), \tag{7}
$$

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 (7) 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 (1985). Network externalities, competition, and compatibility. *The American economic review 75*(3), 424–440.

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.