


WHEN IS A CONTRARIAN ADVISER OPTIMAL?

By

Robert Evans and Sönje Reiche

COWLES FOUNDATION PAPER NO. 1830



COWLES FOUNDATION FOR RESEARCH IN ECONOMICS
YALE UNIVERSITY
Box 208281
New Haven, Connecticut 06520-8281

2023

<http://cowles.yale.edu/>

When Is a Contrarian Adviser Optimal?[†]

By ROBERT EVANS AND SÖNJE REICHE*

We compare contrarian to conformist advice, a contrarian expert being one whose preference bias is against the decision-maker's prior optimal decision. Optimality of an expert depends on characteristics of prior information and learning. If either the expert is fully informed or fine information can be acquired cheaply, then for symmetric distributions F of the state, a conformist (contrarian) is superior if F is single peaked (bimodal). If only coarse information can be acquired, then a contrarian acquires more on average and hence is superior. If information is verifiable, a contrarian has less incentive to hide unfavorable evidence and again is superior. (JEL D72, D82, D83, G34, H71, I12, L94)

A decision-maker (*DM*) has to choose between two decisions, and given the available information, she has a strict preference for one of them. Suppose that she can make use of an expert adviser who has some additional private information but that any such adviser also has an intrinsic preference bias, relative to the *DM*, for one of the decisions. Is it better for the *DM* to use an adviser whose bias is in favor of her initially preferred course of action (a *conformist* adviser) or one whose bias is against her preferred action (a *contrarian* adviser)? Secondly, how is the choice between these two different types of adviser affected by the nature of the information available to that adviser?

It is a standard theme in the psychology literature that there is frequently too much alignment between principals and their expert advisers and that this is detrimental to sound decision-making. The selective exposure hypothesis, which has been extensively studied by social psychologists (see Klapper 1949 and Janis and Mann 1977), holds that people generally seek messages with which they agree and avoid those with which they disagree. Two related phenomena are *bolstering* and *groupthink*. Bolstering,¹ which has also been much studied by social and organizational psychologists, refers to the practice of magnifying the attractions of a chosen alternative and downplaying those of others. For example, Fellner and Marshall (1970) present evidence of potential kidney donors avoiding information that might

*Evans: St. John's College and Department of Economics, University of Cambridge (email robert.evans@econ.cam.ac.uk); Reiche: Department of Economics, Yale University (email soenje.reiche@yale.edu). Leslie Marx was the coeditor for this paper. We are grateful for helpful comments by various seminar audiences and for insightful suggestions by four referees that substantially improved the paper.

[†]Go to <https://doi.org/10.1257/mic.20200204> to visit the article page for additional materials and author disclosure statement(s) or to comment in the online discussion forum.

¹This is related to the phenomenon of *cognitive dissonance*; see Festinger (1957). For applications to economics, see Akerlof and Dickens (1982) and Rabin (1994).

cause them to reconsider their initially preferred decision to donate. Groupthink,² a collective form of defensive avoidance, refers to a situation in which a leader has a group of advisers who share the leader's judgments and provide rationalizations that bolster the leader's preferred course of action. Such situations are characterized by a strong degree of cohesion within the group and conscious or unconscious efforts to exclude expert advisers who might advocate a different course of action. Janis (1972) provides case studies of a number of examples of momentous (and disastrous) decisions that, he argues, exhibited all the symptoms of groupthink, including the Bay of Pigs invasion of Cuba and the decision in December 1941 by the commander in chief of the US Pacific fleet not to prepare for an attack on Pearl Harbor. The concept has subsequently been suggested as a contributing factor in many other debacles, including the 2008 financial crisis.

The economics literature has not directly posed the above questions, but it suggests various relevant conjectures, most of which appear to favor the conformist adviser.³ Better communication or delegation is facilitated by closer alignment of preferences between adviser and *DM* (Crawford and Sobel 1982; Aghion and Tirole 1997). A conformist might be preferred because, unlike a contrarian, he would not overturn the *DM*'s initial choice on the basis of weak evidence (Li and Suen 2004). Advice that the *DM* should change the initial decision may be more credible if it comes from a conformist than if it comes from a contrarian (Cukiermann and Tommasi 1998; Calvert 1985).⁴ If information is hard and endogenous, the conformist may have a greater incentive to acquire information since the evidence is likely to favor the decision he prefers (Che and Kartik 2009).

We consider these issues in the context of a *DM* who can select either an adviser with a contrarian bias or else one with a conformist bias of the same magnitude. We assume symmetry between the biases of the two types of adviser in order to provide a clean test of the effects of the two types of bias. Of the two possible decisions, d_0 and d_1 , *DM* prefers d_0 ex ante.

First, we consider the case in which the *DM* delegates to an adviser who is fully informed. In this case either type of adviser may be superior depending on the distribution of the payoff of action d_1 (we normalize the payoff of d_0 to 0). The distribution determines the expected cost of a type I error (choosing d_1 when d_0 is better) and a type II error (vice versa). If the type I cost is larger, then the conformist is better than the contrarian, and vice versa if the type II cost is larger. This, in turn, implies that what matters is the shape of the distribution of payoffs in the region around zero (when the decision is relatively finely balanced from the *DM*'s point of view). A conflict of interest arises when there is evidence against the adviser's bias, but it is not strong: for example, in the case of a conformist adviser, weak but positive evidence in favor of the contrarian decision d_1 . Plausible payoff distributions are those that are symmetric and either single peaked (e.g., normal) or bimodal (corresponding to a case in which decision d_1 is likely to be either

²See Benabou (2013) for an analysis of groupthink using the tools of economic theory.

³We discuss the related literature in more detail below.

⁴Suen (2004) uses this idea to explain the perpetuation of biased beliefs: people rationally read newspapers with whose political slant they agree.

significantly better than d_0 or else significantly worse). With a single-peaked distribution, weak evidence in favor of d_1 is less likely than weak evidence in favor of d_0 , simply because the ex ante preferred action is d_0 (the mean of the distribution is negative), and the conformist is therefore preferred. With a bimodal distribution, weak evidence in favor of d_0 is less likely than weak evidence in favor of d_1 , and so the contrarian is preferred.

Secondly, we examine the case in which the adviser is initially uninformed but sequentially acquires costly (but not too costly) independent signals and chooses when to stop gathering information and make a decision. It turns out that the optimal choice of adviser depends on the nature of the information structure. Specifically, we consider two polar cases, which we refer to, respectively, as “fine” and “coarse.” A fine information process is one with the property that the posterior expectation of the state converges to the true state. More precisely, for any desired error tolerance ϵ , the ex ante probability that the posterior expectation is within ϵ of the true value can be made arbitrarily close to 1 by committing to collect enough signals. In this case, for a sufficiently small cost of signals, the comparison between advisers is the same as it would be if they were fully informed; that is, the conclusions of the previous section are robust in the sense that optimal learning with sufficiently small learning cost gives the same result.

A coarse information process, on the other hand, only gives *qualitative* information; i.e., it only tells us which action is better for the *DM*. The leading example of a process that is coarse and symmetric (i.e., is not intrinsically biased in favor of either decision) is a binary i.i.d. process such that a high signal is evidence that d_1 is better and a low signal is, symmetrically, evidence that d_0 is better. We show that for such a process, the contrarian (who is biased in favor of d_1) is better for the *DM*. For relatively low cost of information collection, either expert will collect enough information that there is no conflict of interest with the *DM* at the point when the decision is made. However, the *DM* always prefers more information to be gathered, and the contrarian will collect more in expectation than the conformist. Essentially, the reason for this is that the conformist is more likely to stop early and choose d_0 than the contrarian is to stop early and choose d_1 . If, in fact, d_1 is better, then the conformist will search for longer because he is more reluctant to stop and choose d_1 than the contrarian is; similarly, if d_0 is better, then the contrarian will search for longer. But ex ante, d_0 is more likely to be better (d_0 is *DM*'s ex ante preferred action), so in expectation, it is the contrarian who will search longer.

The above results concern the case of soft—i.e., unverifiable—information. We also consider a model in which the adviser may or may not be in possession of hard evidence that he can suppress but not otherwise manipulate; that is, the evidence technology is the one proposed by Dye (1985). We focus on the incentives to reveal information rather than sequentially collect it—either the adviser has some evidence or he does not. In this case, the contrarian adviser is better for *DM* than the conformist. When the adviser is a conformist, the best mechanism for *DM* (and the best equilibrium of the evidence disclosure game) is suboptimal, but when the adviser is a contrarian, there is a mechanism (and an equilibrium of the disclosure game) that delivers the *DM*'s first-best outcome. The optimal mechanism is simple: the decision is d_0 unless the adviser provides evidence that d_1 is better than d_0 for

DM. The contrarian is always willing to provide this evidence if it exists, and this achieves the first best. No such equilibrium exists in the case of the conformist in any mechanism. If there is no evidence, the optimal decision for *DM* is d_0 , so in a sense, d_0 is the “incumbent” action. However, this means that there is an incentive for the conformist to suppress evidence when it weakly favors d_1 since the “incumbent” action is his preferred action. In other words, the contrarian is more willing to submit evidence because he is the one who wants to change the *DM*’s mind.

There is an underlying connection between this result and the result, described above, that when information is coarse and unverifiable, the contrarian is preferred. In that case, the reason is that when an expert’s intrinsically less favored action appears to be optimal, he is more inclined than the other expert would be to carry on searching. Because the contrarian’s less favored action, d_0 , is *ex ante* more likely to be better, he is the one more likely to be in this position and therefore will search longer and make better decisions on average. In a sense, when information is hard, the contrarian is preferred because he is the one who wants to change the *DM*’s mind; when information is coarse and soft, the contrarian is preferred because he is the one who is likelier to want to change his own mind.

Related Literature.—The literature on communication, delegation, and information acquisition is large and growing, but it has not considered the issue that we study—namely, how the *DM*’s choice between experts who are biased in favor of or against the *DM*’s preferred decision depends on the nature of information. The literature on delegation (e.g., Holmstrom 1977, 1984; Melumad and Shibano 1991; Alonso and Matouschek 2005) shows that the optimal delegate is generally one who is aligned with the principal in his response to new information. Our results can be understood as showing what this general principle implies for the relative merits of contrarian and conformist experts in the context of different information structures.

A number of papers have results that are related to ours, some of which were discussed above. Meyer (1991) shows that when a *DM* can choose the bias of a contest, it is optimal for her to bias it in favor of the current leading candidate; this might suggest that the conformist expert is superior since he has a bias against the less favored decision. Our analysis is different in a number of respects; in particular, we have experts who act strategically and have the same prior information as the *DM*.

For the case of fine information, perhaps the closest paper is Li and Suen (2004). Their results are mainly concerned with the optimal combination of biases in settings with multiple experts and, potentially, varied principals, but they also show, in the context of a model with a binary decision and continuous signals, that a conservative *DM* prefers an expert who is more conservative than her. Translated to our model, this implies that the *DM* prefers a conformist expert. We discuss Li and Suen’s (2004) result further in Section III below.

Che and Kartik (2009) study the optimal choice of expert in a setting with hard (verifiable) information and costly information acquisition. They show that it is optimal to choose an expert with a divergence of opinion from the *DM*’s. One crucial difference is that their paper assumes heterogeneous priors, while we assume common priors. One key effect driving their results, which follows from the heterogeneous prior assumption, is the *persuasion effect*: since an expert believes

that new information will move the *DM*'s opinion in the direction of his own (different) opinion, he has an incentive to acquire such information. We discuss the literature on hard evidence further in Section VI below. To the best of our knowledge, there is no previous paper that considers the issue addressed in Section VI.

Li (2001) studies a committee of *DMs* with homogeneous preferences, a binary decision, and costly information acquisition and shows that it is optimal to commit to a threshold that is biased against the ex ante optimal decision. This ex post inefficiency has the effect of overcoming a free-rider problem in acquiring information; essentially, it reduces the attractiveness of the decision that is (constrained) optimal in the absence of new information.⁵ Our result in the case of coarse information has some similarity, but the logic is different. We do not introduce a bias against the agent's ex ante optimal decision (the result applies whatever the contrarian's initially optimal decision); instead, using a contrarian introduces a bias against the principal's initially optimal decision. Furthermore, our result depends on the sequential nature of information collection—it does not apply in general if the agent commits to a certain number of signals.

Other notable papers that consider costly information acquisition are Che and Mierendorff (2019); Dur and Swank (2005); Gerardi and Yariv (2008); and Argenziano, Severinov, and Squintani (2016). Che and Mierendorff (2019) study a model with a binary decision and binary state space in which a *DM* has access to two sources of information, each biased in favor of one of the two possible decisions, and must decide at each moment how much of her attention to devote to each source as well as whether to stop and make a decision. They show that, depending on the prior, the optimal learning strategy may be biased toward the decision that is more likely to be optimal (for relatively extreme beliefs) or less likely to be optimal (for nonextreme beliefs). Our model is different because it concerns delegation to a biased expert, who acts strategically, but a nonbiased information source. Dur and Swank (2005), like us, analyze a model with a binary decision and continuous state space, but they consider the optimal choice of biased expert for a biased principal. The optimal expert is one whose bias is in the same direction as the *DM*'s but less extreme, because of the trade-off between the incentive to collect information, which is highest for an unbiased adviser, and the need for alignment of interest. Gerardi and Yariv (2008) study a model in which two experts engage in costly acquisition of public signals, and they analyze the optimal choice of bias and form of information collection (simultaneous or sequential). They show that for a moderately biased principal, it is optimal for both experts to have identical, extreme biases, opposite to the principal's bias, and to collect information sequentially. Argenziano, Severinov, and Squintani (2016), like us, model information as a sequence of binary signals (though, unlike ours, their expert makes a single decision about how many signals to collect). Their focus too is different from ours: they show that the *DM*, in the context of a Crawford–Sobel (1982) setting, prefers to have the expert collect information and then communicate via cheap talk rather than to delegate the decision or collect the information herself, because she is able to induce overinvestment in information

⁵A related argument is found in Szalay (2005). He shows that it can be optimal to reduce the agent's decision space by eliminating intermediate options in order to increase the incentive to collect information.

acquisition by the expert, even, surprisingly, when she cannot observe how much information has been collected. In our model, there is overinvestment by the contrarian expert but not by the conformist, in the sense that the contrarian searches for longer in expectation than the *DM* would if she had the same information collection technology.

The next section sets up the basics of the model. In Section II we introduce a number of case studies of decision processes that we use subsequently in discussing our results. In Section III we examine the case in which the adviser is fully informed. Section IV contains the analysis of the model of sequential costly acquisition of soft information. In Section V we discuss a number of variations on this model. Section VI considers the case of hard evidence, and Section VII contains some concluding remarks. Those proofs not in the text are in the Appendix.

I. The Setting

We consider a *DM* who must choose between two policies or projects, denoted by d_1 and d_0 . She has the option of making use of an expert who potentially has more information than she has about the merits of the two choices but who may not share her goals: relative to the *DM*, he may be biased in favor of one or the other of the two choices.

We suppose that, ex ante, *DM* strictly prefers one of the decisions (d_0). Her preference, however, could alter as a result of learning the information that the expert has. Our interest is in whether it is better to use an expert who is biased in favor of *DM*'s ex ante preferred choice (a "conformist") or one who is biased against it (a "contrarian"). In order to focus on this question, we consider the choice between two potential experts, denoted E_0 and E_1 , who have symmetric but opposite biases⁶ relative to *DM*. For definiteness, we will interpret d_0 as the status quo policy, and our assumption therefore is that *DM* initially has a preference for the status quo, while the conformist expert is a "conservative" (biased toward the status quo) and the contrarian is an "activist" (biased toward the new policy). However, this assumption is inessential—our interest is in the factors favoring either contrarian or conformist experts, and equivalent results would obtain if *DM*'s initial preference were for d_1 rather than d_0 .

We model the preferences of the three agents as follows. For each agent, the status quo payoff is normalized to zero. *DM*'s payoff from decision d_1 varies, depending on the underlying state of the world, between $-\infty$ and $+\infty$. Without loss of generality, we can redefine the set of states of the world to be $\Pi = (-\infty, +\infty)$; that is, for each $\pi \in \Pi$, we identify state π with the set of underlying states in which the payoff of d_1 is equal to π . Expert E_0 is biased against d_1 in the sense that in every state π , his payoff from d_1 is less than that of *DM* by a fixed amount $b > 0$. Symmetrically, the corresponding payoff for E_1 is always greater by b than it is for *DM*.

⁶The best expert for *DM* is one who shares her preferences. We consider biased experts because we are interested in examining the relative costs of the two types of bias.

More precisely, DM , E_0 , and E_1 have von Neumann–Morgenstern utility functions defined on $\{d_0, d_1\} \times \Pi$ and denoted, respectively, by U_D , U_0 , and U_1 . $U_i(d_j, \pi)$ is the utility of the agent corresponding to i (DM , E_0 , or E_1) for decision d_j ($j = 0, 1$) in state π . For any $\pi \in \Pi$, $U_i(d_0, \pi) = 0$ ($i = D, 0, 1$),

$$U_D(d_1, \pi) = \pi,$$

$$U_1(d_1, \pi) = \pi + b,$$

and

$$U_0(d_1, \pi) = \pi - b.$$

The optimal decision rule for DM , assuming she knows the state π , is therefore to select d_1 if and only if $\pi > 0$ (we assume throughout that each agent would select d_0 if indifferent between the two decisions). However, E_0 would select d_1 if and only if $\pi > b$, while E_1 's rule would be to select d_1 if and only if $\pi > -b$. In other words, there is a conflict of interest between DM and the conformist (E_0) if $\pi \in (0, b]$ —she would want decision d_1 , but the expert, influenced by his bias, would choose d_0 . If $\pi \in (-b, 0]$, there is an opposite conflict in the case that the expert is a contrarian.

We assume that the state of the world π follows a distribution F with a continuous and strictly positive density function f and expectation μ and that $\int_{-\infty}^{\infty} |\pi| dF < \bar{\pi}$ for some finite $\bar{\pi} > 0$. We also assume that, ex ante, DM (hence, a fortiori, E_0) strictly prefers d_0 and that E_1 strictly prefers d_1 ; that is,

$$\mu \equiv \int_{-\infty}^{\infty} \pi dF \in (-b, 0).$$

II. Examples

Here we give some examples of the kind of situation the above is intended to model, with some preliminary discussion. We refer to these examples later in discussing applications of our results.

A. Mergers and Acquisitions

When a firm decides whether to acquire another firm, there is typically an initial phase in which the acquiring firm identifies a set of targets and performs a first valuation analysis. It may then make an offer to the most attractive target. Once the offer is accepted, the acquirer's team of legal and financial experts (investment banks, law firms, financial advisors) carries out in-depth due diligence on the target firm. The decision whether to continue with the acquisition will depend on the results of this due diligence. We can think of the acquiring firm as the DM and its investment bank as the expert adviser. The acquirer's initial preference, at the time it starts the due diligence process, is in favor of acquisition.

It is plausible that any investment bank involved in a merger is biased in favor of the merger taking place; i.e., in our terminology, the adviser in this case is typically a conformist. As part of their compensation, investment banks receive a “success fee,” the value of which is based on the size of the acquisition price. In addition a successful merger often leads to associated transactions that the investment bank will probably be charged with handling, and it raises the likelihood of future business with the acquiring firm.⁷

B. *The Poll Tax*

In the 1980s, Margaret Thatcher’s UK Conservative government proposed to reform the system of local government finance by replacing the existing system, a tax based on assessed property values, with a per capita flat tax on individuals, commonly known as the “poll tax.” One major motivation was to ensure that all voters in local elections were also taxpayers and therefore had a financial stake in the spending decisions of local councils (“no representation without taxation”). The policy was developed by a small and close-knit review group of relatively junior ministers and officials. This group presented its recommendation to a meeting of the prime minister and Cabinet in such a way that it immediately became settled policy and politically unstoppable. The reform was introduced in 1990 and was quickly revealed to be a major political blunder. It was widely regarded as unfair by virtue of its regressiveness; it was very difficult to collect, partly because of administrative difficulties, partly because of its perceived unfairness (many people refused to pay it); and it led to riots in the streets. Thatcher was forced out of office in late 1990, and the poll tax was probably the main single cause of her political demise.

C. *HS2*

High Speed 2 (HS2) is a projected, and highly controversial, high-speed train line in the United Kingdom linking London, Birmingham, and parts of the North of England. Supporters claim that it will increase economic growth and redress regional economic imbalances. Opponents argue that there has been no convincing cost-benefit analysis to show that this is the best way to spend such a large amount of public money (estimated to be below £40 billion when it was first formally taken up by government in 2009 but over £80 billion by the time the decision was made to go ahead (Oakervee Review 2019)) and, furthermore, that the associated environmental damage is out of proportion to any benefit.

⁷“Vodafone’s successful—and complex—acquisition of Mannesmann deepened the firm’s already close relationship with the telecom giant. As a result, Goldman Sachs received mandates for a number of transactions related to the merger. ... In 2013, Goldman Sachs acted as joint financial advisor to Vodafone in its €10.7 billion (US\$14.1 billion) acquisition of Kabel Deutschland Holding AG” (Goldman Sachs 2019).

D. *The Swine Flu Affair*

In January 1976 there was an outbreak of a previously unknown strain of swine flu among army recruits at Fort Dix, New Jersey. Several hundred were infected, some were hospitalized, and one died. Swine flu was thought to have been responsible for the 1918 pandemic, and some epidemiologists believed that a major pandemic would ensue in 1976, killing as many as one million Americans. The head of the Center for Disease Control (CDC) persuaded President Ford to announce an immediate, and unprecedented, program to vaccinate the entire US population within the following few months. It was subsequently widely seen as a fiasco. Some parts of the press regarded it as political and based on inadequate evidence. There were delays in production of the vaccine, partly caused by disputes over indemnification of vaccine manufacturers, and when vaccination finally began, the program had to be halted temporarily when a number of people died immediately after receiving the vaccine (coincidentally, but the press had not been prepared for this eventuality). The program was finally abandoned because of rare side effects. The pandemic itself never materialized. The affair led to reputational damage for the CDC and, it has been argued, to the growth of antivaccine sentiment.

Our analysis assumes that the decision to be made is a binary one (although we discuss a partial generalization to multiple actions in Section V below). We would argue that in many important cases, this is an appropriate assumption. For example, in the case of mergers and acquisitions, once the preferred target firm has been identified and due diligence has begun, the decision is, in effect, binary (although there is negotiation over the price). Often, a decision breaks down into a top-level binary decision (e.g., whether or not to go to war in Iraq) together with multiple subsidiary decisions contingent on the top-level decision (e.g., different methods of prosecuting a war, different diplomatic alternatives to war).

Even if there are, in principle, many different policies that could be followed, political considerations can often force the decision to be effectively binary. For example, economic rationality would require an assessment of how the money allocated for HS2 should best be spent, whether on HS2 or, perhaps, on a variety of other possible transport projects. But once HS2 became politically salient, it was far from clear that this money would be available for these other projects, so the decision became a binary one.⁸

There were, in principle, several possible alternatives to the poll tax, but the review team never seriously considered a local sales tax or local income tax, because it thought they were politically out of the question (King and Crewe 2013). That left only a per capita tax or a reformed property tax (which was adopted after the poll tax was abandoned). In the swine flu affair, doing nothing was regarded as out of the question, so there were, in effect, two realistic options—universal vaccination or a

⁸“HS2 seems to have been considered almost entirely on its own merits or demerits rather than in terms of its merits in comparison with those of other projects, whether in the field of transport or anything else” (King and Crewe 2013, 526).

policy of vaccine manufacture plus stockpiling and waiting to see if the pandemic materialized.

In each of the examples above, the *DM* effectively made use of an expert adviser or homogeneous team of advisers, whether an investment bank in the acquisitions case, the review team for the poll tax, the minister of transport in the case of HS2, or the head of the CDC in the swine flu affair. In each case, the adviser had a bias, whether for structural or ideological reasons.⁹ In general the adviser is not formally delegated the power to make the decision, but we would argue that, de facto, the adviser's decision is pivotal. For example, in the United Kingdom, the prime minister does not have the time or personnel to deal with every issue, even a large one such as HS2, so although the prime minister and Cabinet will vote on it, it is likely that the transport minister will effectively make the decision.

III. The Static Case

We first consider the case in which the expert knows the state of nature and the *DM* delegates the decision to the expert. If *DM* were also fully informed about π , and if $\pi \in (0, b]$, then *DM* would prefer E_1 since he would pick *DM*'s preferred action, d_1 , whereas E_0 would pick d_0 . If $\pi \in (-b, 0]$, then the situation is reversed—she would prefer E_0 . For all other values of π , there is no conflict with either expert, so she is indifferent between them.

Therefore, ex ante, *DM* strictly prefers E_0 to E_1 if and only if

$$\int_0^b U_D(d_1, \pi) - U_D(d_0, \pi) dF(\pi) < \int_{-b}^0 U_D(d_0, \pi) - U_D(d_1, \pi) dF(\pi),$$

that is, if and only if

$$\int_{-b}^b \pi f(\pi) d\pi < 0,$$

and strictly prefers E_1 if and only if the inequality is reversed. This immediately implies the following Proposition.

PROPOSITION 1: *If the decision is delegated to a fully informed expert, then*

(i) *The conformist, E_0 , is strictly preferred by *DM* if*

$$(1) \quad \int_{-b}^0 |\pi| f(\pi) d\pi > \int_0^b \pi f(\pi) d\pi.$$

(ii) *The contrarian, E_1 , is strictly preferred by *DM* if the above inequality is reversed.*

We can interpret the two sides of inequality (1) as, respectively, expected costs of type I and type II errors. If we define the null hypothesis to be the statement that the

⁹Some senior public health officials believed that the swine flu vaccination drive would raise the profile and prestige of public health professionals (Neustadt and Fineberg 1978, 12).

status quo, d_0 , is better for DM , then the LHS (RHS) is the conditional expected cost of making a type I (II) error multiplied by the probability of doing so. If the expected cost of a type I (II) error is higher, then the conformist (contrarian) is better.

Proposition 1 enables us to say more about how the choice between the experts depends on properties of F . For a plausible class of distributions F , the conformist is better if F is unimodal (single peaked), while the contrarian is better if F is bimodal. Note that (1) is equivalent to

$$\int_0^b \pi [f(\pi) - f(-\pi)] d\pi < 0,$$

so that a sufficient condition for E_0 to be preferred is that $f(\pi) < f(-\pi)$ for all $\pi \in (0, b)$.

We refer to a distribution F with continuous, differentiable density f and mean μ as *unimodal* if f has exactly one peak (local maximum). F is *bimodal* if f has exactly two peaks and one trough (local minimum). F is *symmetric* if $f(\mu - z) = f(\mu + z)$ for all $z \in \mathbb{R}$.

COROLLARY 1: *If the decision is delegated to a fully informed expert, then*

- (i) *The conformist, E_0 , is strictly preferred by DM if F is unimodal and either (a) F is symmetric or (b) f has peak $\pi' \leq -b$.*
- (ii) *The contrarian, E_1 , is strictly preferred by DM if F is bimodal, the right-hand peak π_2 satisfies $\pi_2 \geq b$ and either (a) F is symmetric or (b) the trough π_3 of f satisfies $\pi_3 \leq -b$.*

PROOF:

(i) (a) See Figure 1. (b) In this case, $f'(\pi) < 0$ on $[-b, b]$, so $f(\pi) < f(-\pi)$ for all $\pi \in (0, b)$. (ii) (a) f is increasing on (μ, b) , so for $\pi \in [0, -\mu]$, $f(-\pi) < f(\pi)$. For $\pi \in (-\mu, b)$, let $\alpha(\pi) = \max\{\pi' \mid \pi' < \pi, f(\pi') = f(\pi)\} = 2\mu - \pi$. Then f is decreasing on $(\alpha(\pi), \mu)$ and $-\pi \in (\alpha(\pi), \mu)$, so $f(-\pi) < f(\pi)$. See Figure 2. (b) $f'(\pi) > 0$ on $[-b, b]$. ■

For example, if the value of the project d_1 is distributed normally, then (i) implies that the conformist is better. A symmetric normal distribution F with negative mean might arise, for example, in the context of the following standard model. At the outset, π is normally distributed with zero mean and variance σ_π^2 . DM observes a signal of the form $\pi + \epsilon$, where ϵ is normally distributed with zero mean and variance σ_ϵ^2 . The realization of the signal is negative. The posterior distribution of π , which is the distribution we describe as F , is then normal with negative mean.

For an example of a symmetric bimodal distribution, suppose that $\pi = \mu + \eta$, where η is symmetric with zero mean and $|\eta|$ follows a gamma distribution $\Gamma(\alpha, \lambda)$ with $\alpha > 2$ and $\alpha - 1 > \lambda(b - \mu)$. The distribution of η , and hence of π , is bimodal since $|\eta|$ has density

$$h(x) = \frac{\lambda^\alpha x^{\alpha-1} e^{-\lambda x}}{\Gamma(\alpha)}$$

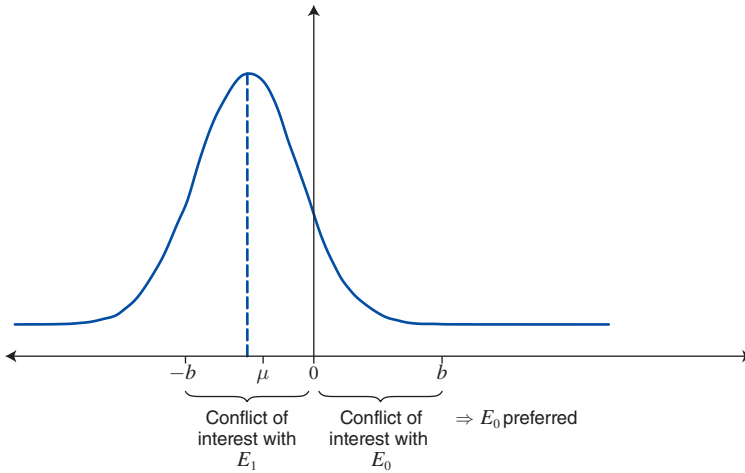


FIGURE 1

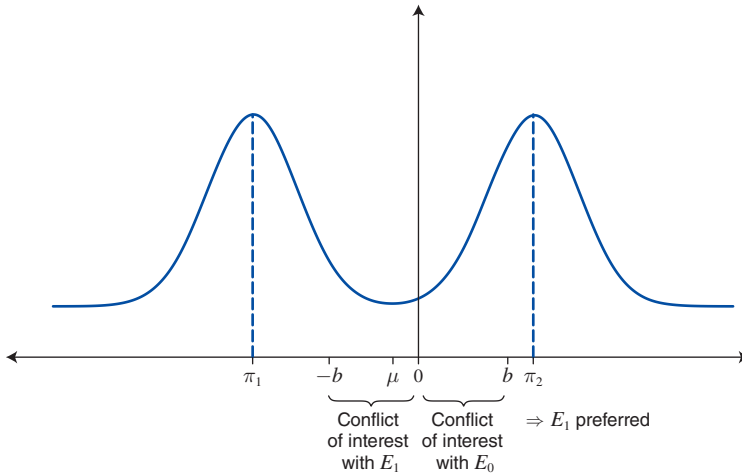


FIGURE 2

with support $[0, \infty]$, where Γ is the gamma function. Therefore, $h(0) = 0$, the slope of h is also 0 at $x = 0$, and h increases to a peak at $x = (\alpha - 1)/\lambda$. Therefore, both peaks of the density of π are outside $[-b, b]$ if $\alpha - 1 > \lambda(b - \mu)$. Part (ii) of the Corollary implies that the contrarian is better in this example.

In the Corollary, we focus on the properties of symmetry, single peakedness, and bimodality because they are plausible properties of payoff distributions. It should be clear, though, that they are relevant only via their implications for the distribution in the interval $[-b, b]$. For example, the contrarian is superior for any distribution, whether bimodal or not, in which the density is low around the (negative) mean but increasing for a sufficient interval to the right of the mean.¹⁰

¹⁰For a paper that studies comparative statics across type distributions in the context of communication, see Szalay (2012).

A bimodal distribution corresponds to a case in which d_1 is likely to be either significantly better than the status quo or else significantly worse. Since the status quo is better ex ante, the downside of the project d_1 is greater than its upside in expectation (for a symmetric distribution), and this implies, somewhat counterintuitively, that the contrarian is better. The reason for this is that the disadvantage for the *DM* of the conformist is that he will choose the status quo when there is weak but positive evidence in favor of the project; the disadvantage of the contrarian is that he will choose the project when there is weak but positive evidence in favor of the status quo. Essentially, in the unimodal case, there is more likely to be weak positive evidence for the status quo than for the project, and vice versa in the bimodal case.

The result for the unimodal case is reminiscent of a result of Li and Suen (2004), but the intuition is somewhat different to the one that they give, namely that the conformist is better because he will only change the decision (relative to the decision *DM* would have made unaided) if there is strong evidence for doing so. But *DM* would also want to change her decision if there were weak (but positive) evidence for doing so, and the contrarian will do that. The key point is whether or not the advantage thereby gained by the *DM* is outweighed by the disadvantage that the contrarian will also change the decision when there is weak (but positive) evidence against doing so; this depends on the shape of the distribution F .

The result in part (ii) of the Corollary requires, in addition to a bimodal distribution, that the bias is not too large. For a sufficiently high bias, the conformist must be preferred because the limit of $\int_{-b}^b \pi f(\pi) d\pi$ as b increases is $\mu < 0$. Essentially, for a high enough bias, either expert would, in effect, ignore his information, so *DM* prefers the conformist.

The Corollary suggests that it may be useful to distinguish between “normal” projects—those whose return distributions have, roughly speaking, the shape of a normal distribution—and “high-gain/high-risk” projects, corresponding to the bimodal case. Decisions about mergers and acquisitions may generally fall into the normal category. It is plausible that the true valuation of a target firm should be distributed approximately normally, conditional on the first valuation analysis, and that due diligence should ensure that the adviser is well informed about the state. The Corollary therefore provides some explanation for the common practice of using an investment bank as adviser since, as we argued in Section II, the investment bank typically has a bias in favor of acquisition and hence is a conformist. HS2, on the other hand, may be an example of a “high-gain/high-risk” project, suggesting that a contrarian adviser may be advisable. Although it may turn out that there are high economic gains from the project, it is equally plausible that the economic returns could be modest and the political costs very high—for example, because of the environmental consequences in marginal political constituencies.¹¹

¹¹ One relatively likely outcome ex ante was that large amounts would be spent on developing and starting the project before its cancellation because of unacceptable cost overruns or for political reasons, giving a large overall loss.

IV. Delegation with Sequential Costly Information Acquisition

We now suppose that the expert is initially uninformed but has access to a sequence of costly signals that are informative about the state π and can decide how much of this information to acquire. In the previous section, the utility to the *DM* of the expert depended only on the extent to which the expert's delegated decision aligned with *DM*'s preferred decision. Now it depends on two factors: (i) the degree of alignment with the decision eventually taken by the expert (measured with respect to the conditional belief at the time of decision) and (ii) the quality of the information collected, which will depend on the length of search that the expert undertakes.

The principal delegates the decision to expert E_i . The expert chooses to either take a decision from $\{d_0, d_1\}$ or pay $c > 0$ for a piece of information—that is, a signal drawn from a distribution that varies with the true state π . Having observed the signal, he decides to either take a decision or pay $c > 0$ for a second signal, and he continues in this way until the decision is made. There is no upper bound on the number of signals that he can observe. If the process ends after t signals with decision d , his overall payoff is $U_i(d, \pi) - ct$, where $U_i(d, \pi)$ is the utility function defined in Section I. *DM*'s utility in that event is $U_D(d, \pi)$. If the process never ends (the expert observes infinitely many signals), *DM*'s payoff is 0 and we take E_i 's to be $-\infty$. In the analysis, we assume for convenience that there is a discrete, infinite time structure and only one signal can be acquired in a period, but this is not necessary—an alternative interpretation is that the expert can acquire signals arbitrarily rapidly, with the proviso that they are collected in sequence. There is no discounting.

Let H_t be the set of histories of length $t \geq 0$, with a typical element of H_t being $h_t = (s_0, \dots, s_{t-1})$ for $t \geq 1$, where s_t is the signal realization in period t . h_0 is the null history, and $H_0 = \{h_0\}$. We denote by H the set of infinite sequences of signal realizations. A policy σ for E_i consists of an initial choice $\sigma^0 \in \{d_0, d_1, C\}$, where C stands for “Continue,” and for each $t \geq 1$, a function σ^t that maps H_t to the set $\{d_0, d_1, C\}$. An optimal policy for E_i is one that, for any time $t \geq 0$ and any history at time t , maximizes his continuation expected utility.¹² We assume that the expert chooses such a policy.¹³ The question is: which type of expert is better for the principal?

We focus on two different models of the information process, which can be regarded as polar cases: *fine information* and *coarse information*. If information is fine, then an expert with any arbitrary bias would, if he collects enough signals, learn the state accurately enough to be able to take his optimal decision with probability arbitrarily close to 1. If it is coarse, an expert can only ever learn which of the decisions is better for the *DM*. We will see that the optimal choice of expert will depend on the nature of the information process.

First, we consider the fine information case.

¹²For simplicity, we assume that if indifferent between stopping and continuing, the expert stops, and if indifferent between d_0 and d_1 , he chooses d_0 .

¹³An optimal stopping rule will exist; see DeGroot (1970).

A. Fine Information

We define a learning process by a probability distribution Ψ over $\Pi \times H$, the set of (state, signal-sequence) pairs, with marginal over Π given by F . Ψ_t is the probability distribution over $\Pi \times H_t$ implied by Ψ . For $h_t \in H_t$, let Φ_{h_t} be the expert's conditional probability distribution over Π after he observes the history h_t . We assume that for all $t \geq 0$ and all $h_t \in H_t$, Φ_{h_t} exists and has a finite expectation, denoted by $E_{\Phi_{h_t}}(\pi)$.

DEFINITION: *The information process is fine if, for any $\epsilon_1 > 0$ and $\epsilon_2 > 0$, there exists $t(\epsilon_1, \epsilon_2)$ such that if $t > t(\epsilon_1, \epsilon_2)$,*

$$(2) \quad \Pr_{\Psi} \left[\left| E_{\Phi_{h_t}}(\pi) - \pi \right| \leq \epsilon_1 \right] \geq 1 - \epsilon_2.$$

In (2) the LHS is shorthand for $\Psi_t \left[(\pi, h_t) \in \Pi \times H_t : \left| E_{\Phi_{h_t}}(\pi) - \pi \right| \leq \epsilon_1 \right]$. In other words, by collecting enough signals, the ex ante probability can be made arbitrarily close to 1 that the (state, signal-sequence) pair is such that the expert's posterior expectation of π is arbitrarily close to the true value of π .

LEMMA 1: *Suppose that information is fine. Then DM's expected payoff, in the sequential information acquisition model with information cost c and a contrarian (resp. conformist) expert, approaches, as c tends to zero, DM's payoff when the expert is a fully informed contrarian (resp. conformist).*

The main step in the proof of Lemma 1 is to show that in the sequential information collection model, E_i 's ex ante expected utility shortfall, relative to his full information optimum, is close to zero when the search cost is close to zero. An equivalent statement then holds for *DM*. The implication of Lemma 1 is that if *DM* would strictly prefer one expert to the other if both experts were fully informed about the state, then the same strict preference obtains in our sequential information collection model for low enough c . This immediately implies:

COROLLARY 2: *Suppose that the information process is fine. Then statements (i) and (ii) of Proposition 1 and statements (i) and (ii) of Corollary 1 apply for the sequential information acquisition model if $c < \bar{c}$, where $\bar{c} > 0$ depends on the learning process Ψ and hence on F .*

This can be thought of as a kind of robustness result—Proposition 1 and its Corollary are robust to the uncertainty about the state that arises endogenously from optimal learning, as long as it is fine.

B. Coarse Information

In this case, we assume that in each period $t \geq 0$, as long as the expert has not stopped, he observes the realization of a binary signal, s_t . Either there is a “positive” signal ($s_t = 1$), which strictly increases the expectation of π , or a “negative”

one ($s_t = 0$), which strictly lowers it. Because we want to isolate the effect of the expert's bias, we are mainly interested in information structures that are symmetric with respect to the two decisions. In our leading example of a coarse symmetric structure, the signals $\{s_t\}$ follow an i.i.d. process given by $\Pr(s_t = 1 | \pi) = \alpha$ if $\pi > 0$ and $\Pr(s_t = 1 | \pi) = 1 - \alpha$ if $\pi \leq 0$, where $\alpha \in (0.5, 1)$. The learning here is coarse in the sense that although the process will eventually reveal which of d_1 or d_0 is better for DM , it will never reveal any finer information about the value of π . The fact that the probability of the low signal conditional on the state being low equals the probability of the high signal conditional on the state being high implies a certain symmetry of the signal process.

More generally, by a coarse, symmetric learning process, we mean one that satisfies the four properties in Assumption 1 below.

We define the stochastic process W by $W(0) = E(\pi | h_0) \equiv \mu$ and $W(t) = E(\pi | h_t)$ for $t \geq 1$; that is, $W(t)$ is the conditional expectation of π at the start of period t . We assume that for any t , any $h_t \in H_t$, and any s_t , $W(t+1) \neq W(t)$. For $h_t \in H_t$ ($t \geq 0$) and $j \in \{0, 1\}$, let (h_t, j) be the history consisting of h_t followed by the signal realization j . We denote by \hat{W} the set of all values that the process W can take; that is, $\hat{W} \equiv \bigcup_{h_t \in H} \{E(\pi | h_t)\}$. We assume for convenience that $0 \in \hat{W}$.

ASSUMPTION 1 (*Coarse, Symmetric Learning*):

- (i) (*Markov*) For any $t \geq 0$, $W(t+1)$ depends only on $W(t)$ and s_t .
- (ii) Given any history $h_t \in H$, $E(\pi | (h_t, 0, 1)) = E(\pi | (h_t, 1, 0)) = E(\pi | h_t)$.
- (iii) \hat{W} is symmetric about 0; that is, if $\tilde{\pi} \in \hat{W}$, then $-\tilde{\pi} \in \hat{W}$.
- (iv) $\bar{w} \equiv \sup(\hat{W}) > b$.

The Markov assumption, (i), implies that the learning process is coarse in the sense that the effect of a signal on the expectation of π depends only on the expectation prior to the signal, not on the particular path of realizations that led to that expectation. (i) and (ii) together imply that this expectation moves through the discrete set \hat{W} one step at a time: it moves up one after a positive signal and down one after a negative one. (ii) and (iii) imply symmetry of the learning process in two senses. By (ii), a positive and negative signal cancel each other out at any stage, and by (iii), starting from a position of indifference between the two decisions (on the part of DM), any sequence in which the number of positive signals less the number of negative signals is $n > 0$ moves the expectation of π up by the same amount as a sequence with n net negative signals would move it down.¹⁴ The symmetry assumption rules out the possibility that, for example, it is more costly to gather positive information about d_1 than about d_0 . (iv)

¹⁴Note that this is not the case if the starting point is a nonzero value of $W(t)$; in general, a positive signal will move the conditional expectation up by a different amount than a negative one would move it down.

ensures that learning is worthwhile for *DM*—with enough positive evidence, d_1 is optimal for either expert, and also, by symmetry (i.e., by Assumption 1(iii)), enough negative evidence makes d_0 optimal since $\underline{w} < -b$, where $\underline{w} \equiv \inf(\hat{W})$.

We maintain Assumption 1 for the remainder of this subsection. Our leading example¹⁵ mentioned above satisfies Assumption 1 if the initial distribution F at the time of delegation, when $t = 0$, derives from an earlier prior distribution (F_0) over Π that is symmetric about 0 and satisfies $E_{F_0}(\pi | \pi > 0) > b$, and, before time zero, there have been some public signals of the form described, the realizations of which have mostly been negative. Then $\bar{w} = E_{F_0}(\pi | \pi > 0)$ and $\underline{w} = -\bar{w}$: as positive signals accumulate, the conditional expectation converges to $E_{F_0}(\pi | \pi > 0)$, and as negative ones accumulate, it converges to $E_{F_0}(\pi | \pi \leq 0)$.¹⁶

For one situation that may plausibly be modeled by a coarse symmetric process, consider a *DM* who is a politician who is only interested in how much the proposed project, or policy, will increase or decrease her party’s chances of winning the next election. The only available advisers care about this too but also have an ideological position, one way or the other. Since the detailed effects of the policy are difficult for the electorate to predict or understand, the available evidence that can be collected can only give diffuse information on its electoral effect—it cannot give useful information on exactly how many votes will be won or lost, but it can give a fairly clear idea, if enough is collected, whether it will be positive or negative.

The One-Signal Case.—Before analyzing the infinite horizon model, we briefly consider the case in which the expert has at most one signal to observe before taking the decision. Denote the signal realization by s . By Assumption 1 and $0 \in \hat{W}, E(\pi | s = 1) \leq 0 < b$, so E_0 will not pay for the signal, because d_0 would be optimal for him regardless of the signal realization. If $E(\pi | s = 0) \geq -b$, then E_1 would not pay for the signal either, for the same reason. In that case the *DM* has no use for an expert—she should simply take decision d_0 . On the other hand, if $E(\pi | s = 0) < -b$, then E_1 will, if the cost of the signal is low enough, pay for the signal and choose d_0 if $s = 0$, d_1 if $s = 1$. However, *DM* prefers d_0 regardless of the signal (strictly if $s = 0$), so the conformist (or, equivalently, nondelegation) is better. The conclusion is that the contrarian has more incentive to collect information than the conformist but the contrarian’s misalignment of interest with *DM* makes him inferior to the conformist. We will see that the finding that the contrarian searches more for information applies also in the infinite horizon case and that

¹⁵The leading example is the most natural one, but other processes satisfy Assumption 1. For example, take any set $\Pi_1 \subset \Pi$ such that Π_1 and Π/Π_1 are symmetric, i.e., $\pi \in \Pi_1$ if and only if $-\pi \in \Pi/\Pi_1$ and $E(\pi | \pi \in \Pi_1) > b$. Then, replacing $(-\infty, 0]$ by Π/Π_1 and $(0, \infty)$ by Π_1 in the definition of the leading example (and making the symmetry assumption described in the current paragraph) gives such a process. It is worth mentioning that coarse learning is also equivalent to the following behavioral (not necessarily Bayesian) rule. In a given sample of signals, let n_+ and n_- denote, respectively, the numbers of positive and negative signals, and let $n = n_+ - n_-$. Given this sample, *DM* selects d_1 if and only if $n > 0$, E_0 selects d_1 if and only if $n > \rho$ for some $\rho > 0$, and E_1 does so if and only if $n > -\rho$.

¹⁶For a continuous, nondegenerate, symmetric, mean-zero distribution F_0 , denote $E_{F_0}(\pi | \pi > 0)$ by w^+ , and let $p_t(h_t)$ denote $\Pr(\pi > 0 | h_t)$. Then, $W(t) = (2p_t(h_t) - 1)w^+$. Letting $\nu_t(h_t) = p_t(h_t)/(1 - p_t(h_t))$, Bayes’ Rule gives $\nu_t(h_t) = \alpha(1 - \alpha)^{-1}\nu_{t-1}(h_{t-1})$ if $h_t = (h_{t-1}, 1)$, i.e., if there is a positive signal at $t - 1$. Therefore, ν_t increases without bound if all signals are positive; hence, $p_t \rightarrow 1$ and $W(t) \rightarrow w^+$. Standard results (e.g., Blackwell and Dubins 1962) also imply that if $\pi > 0$, p_t converges almost surely to 1.

unlike in the one-signal case, it implies, for low search costs, that the contrarian is then strictly better than both the conformist and nondelegation.

The Infinite Horizon Case.—Our infinite horizon model is a version of the classical sequential decision problem (see Wald 1947).

It is well known that there is a unique¹⁷ optimal policy for E_i and this policy takes a threshold form.

PROPOSITION 2: *For any search cost $c > 0$, there exist $\underline{W}_i(c) \in \hat{W}$ and $\overline{W}_i(c) \in \hat{W}$, ($i = 0, 1$), where*

$$\underline{W}_0(c) \leq b \leq \overline{W}_0(c), \underline{W}_1(c) \leq -b \leq \overline{W}_1(c),$$

and, omitting the dependence on c , $\underline{W}_1 = -\overline{W}_0$, $\underline{W}_0 = -\overline{W}_1$ and $\overline{W}_1 < \overline{W}_0$ (hence, $\underline{W}_1 < \underline{W}_0$), such that

- (i) E_i 's optimal policy is: at any time t , (a) stop and choose d_0 if $W(t) \leq \underline{W}_i$, (b) stop and choose d_1 if $W(t) \geq \overline{W}_i$, and (c) otherwise continue;
- (ii) (a) $\lim_{c \rightarrow 0} \underline{W}_i(c) = \underline{w}$, $\lim_{c \rightarrow 0} \overline{W}_i(c) = \overline{w}$. (b) $\underline{W}_i(c)$ (resp. $\overline{W}_i(c)$) is increasing (resp. decreasing) in c . (c) $\lim_{c \rightarrow \infty} \underline{W}_0(c) = \lim_{c \rightarrow \infty} \overline{W}_0(c) = b$ and $\lim_{c \rightarrow \infty} \underline{W}_1(c) = \lim_{c \rightarrow \infty} \overline{W}_1(c) = -b$.

Some of the results in this Proposition are standard, but for completeness, we provide a full proof in the Appendix.

Assuming that the initial expectation $W(0)$ lies in the continuation interval of the chosen expert, the process stops when the expectation reaches one or the other endpoint of this interval. Because of their opposite biases, E_1 's interval is shifted to the left, compared with E_0 's interval. The fact that each of the two intervals is symmetric to the other, about zero, follows from the fact that E_0 and E_1 are symmetric with respect to the two decisions, because of their utility functions and the symmetry of the information process (Assumption 1). (ii) implies that if c is low enough, the expert will stop only when the conditional expectation is close to its minimum or maximum. If c is high enough, then the expert does not search but instead takes an action in period 0. Note that the Proposition implies that for low enough c , either expert will be aligned with DM , in the sense that when they take a decision, the posterior belief will be such that the decision is optimal for DM .

PROPOSITION 3: *If the learning process is coarse and symmetric, $W(0) \in (-b, 0)$, and c is low, the contrarian, E_1 , is preferred to the conformist.*

¹⁷Given our assumption in footnote 12.

PROOF:

By our assumptions, $\underline{w} < -b < W(0) = \mu < 0$. By Proposition 2(ii), $\lim_{c \rightarrow 0} \underline{W}_0(c) = \underline{w}$, so there exists \hat{c} such that if $c < \hat{c}$, $\underline{W}_0(c) < W(0) < 0 < \overline{W}_1(c)$. Assume that $c < \hat{c}$. Then either expert would search (and, moreover, would choose *DM*'s optimal choice on stopping). Let $\tilde{W}_t(z)$ be the random process followed by the expectation of π beginning at time zero with $E(\pi) = z$. $W(t)$ is then equal to $\tilde{W}_t(\mu)$. Given $(x, y, z) \in \hat{W}^3$ such that $x < z < y$, let $\theta((x, y); z)$ be the probability that $\tilde{W}_t(z)$ hits x before it hits y .

Conditional on the process $W(t)$ hitting \overline{W}_1 before it hits \underline{W}_0 , *DM* prefers E_0 to E_1 since E_1 would at this point stop and choose d_1 , whereas E_0 would continue searching. The conditional expected payoff increment that *DM* obtains in that event from E_0 (compared with E_1) is $\theta((\underline{W}_0, \overline{W}_0); \overline{W}_1)(-\underline{W}_0)$. This is because¹⁸ the decision would only differ from d_1 if \underline{W}_0 is hit before \overline{W}_0 , in which case the benefit of d_0 over d_1 is $-\underline{W}_0$.

Symmetrically, conditional on hitting \underline{W}_0 before \overline{W}_1 , *DM* prefers E_1 to E_0 by

$$(1 - \theta((\underline{W}_1, \overline{W}_1); \underline{W}_0))(\overline{W}_1),$$

and by symmetry, this is equal to $\theta((\underline{W}_0, \overline{W}_0); \overline{W}_1)(-\underline{W}_0)$. Since the payoff advantage of E_0 conditional on hitting \overline{W}_1 before \underline{W}_0 equals the payoff advantage of E_1 conditional on hitting \underline{W}_0 before \overline{W}_1 , the question is which of these thresholds is more likely to be hit first. That is, E_1 is strictly better for *DM* than E_0 if

$$\theta((\underline{W}_0, \overline{W}_1); W(0)) > 1 - \theta((\underline{W}_0, \overline{W}_1); W(0)).$$

By the optional stopping theorem, *DM*'s expected payoff from stopping when either \underline{W}_0 or \overline{W}_1 is reached, and choosing d_1 in either case, equals the current expected payoff from d_1 , so

$$\theta((\underline{W}_0, \overline{W}_1); W(0))\underline{W}_0 + [1 - \theta((\underline{W}_0, \overline{W}_1); W(0))]\overline{W}_1 = W(0).$$

Therefore, since $\underline{W}_0 = -\overline{W}_1$,

$$[1 - \theta((\underline{W}_0, \overline{W}_1); W(0))] - \theta((\underline{W}_0, \overline{W}_1); W(0)) = \frac{W(0)}{\overline{W}_1} < 0$$

since $W(0) = \mu < 0$. ■

As the proof of Proposition 3 shows, if c is low enough that each expert's continuation interval strictly contains 0, either expert would, on making a decision, be aligned with *DM*, so the reason that the contrarian is preferred is that he will search for longer, which benefits *DM*. Essentially, the contrarian searches more because

¹⁸This argument relies on the optional stopping theorem; see the proof of Proposition 2(i).




FIGURE 3

the initial expectation is further from the endogenously determined decision threshold in the case of the contrarian than in the case of the conformist. As can be seen from Figure 3, $W(0)$ is closer to the middle of the search interval for the contrarian than for the conformist, by virtue of symmetry and the fact that $W(0) < 0$. That is, the conformist, E_0 , is comparatively likely to stop early and choose his preferred option; the chance of E_1 doing so is lower because E_0 's preferred option, d_0 , is ex ante likely to be better than d_1 . If, in fact, d_1 is better, then E_0 will probably search for longer because he is more reluctant to stop and choose d_1 than E_1 is; similarly, if d_0 is better, then E_1 will search for longer. But ex ante d_0 is more likely to be better, so in expectation it is the contrarian who will search longer. The contrarian is more likely to find evidence that he does not “like” and so to search for longer, implying a higher likelihood of making the correct decision.

A related observation is that the contrarian is more likely to take account of evidence that goes against his bias: $W(0)$ is closer to \underline{W}_1 than it is to \overline{W}_0 , so the contrarian, E_1 , is more likely to end up choosing d_0 (at \underline{W}_1) than E_0 is to choose d_1 (at \overline{W}_0). In other words, because d_0 is ex ante more likely to be better, E_1 is more inclined to act on evidence in favor of d_0 than E_0 is to act on evidence in favor of d_1 .

The fact that the choice between experts hinges on which will search for longer distinguishes the coarse from the fine information case. When information is fine, the potential conflict of interest between the expert and DM does not vanish when the expert collects enough information. In this case DM still benefits in expectation from the collection of more information, but when c is low—and, hence, a large amount of information is collected by either expert—the optimal choice between the two experts depends overwhelmingly on which has a greater conflict of interest.

There is a parallel between the findings that the contrarian is superior for coarse information structures and for fine information structures when the initial distribution F over π is bimodal. In the coarse case, if the expert searches for long enough, his posterior expectation will converge to either \underline{w} or \overline{w} . That is, the distribution of the limit belief is bimodal. However, the underlying reason for the superiority of the contrarian differs across the two cases. In the coarse case, the reason is that the contrarian will search more; in the fine bimodal case, it is that the contrarian has a lower expected conflict of interest when the two decisions are relatively finely balanced.

Proposition 3 may shed light on the poll tax and swine flu episodes. In each case, the likely effects of the policies were not clearly amenable to precise prediction, but it could have been reasonably well predicted, with enough research, whether or not they were the better policy. The poll tax review team were conformists in the sense that they were biased toward Thatcher's initially preferred policy. Proposition 3 would then suggest that they would not look hard enough at the available choices before making the decision. Similarly, it would suggest that Ford, rather than effectively allowing the head of the CDC, David Sencer, to make the decision about the immunization program, should have delegated it to a more sceptical official or body.

According to King and Crewe, "numerous other interested parties ... might have been expected to have views about—and possibly even to harbour doubts about—the wisdom and practicality of introducing a flat-rate tax. But the new policy had already been adopted before any of them were informed and consulted." One civil servant told King and Crewe that the review team never seriously addressed the issue of compliance.¹⁹

Two political scientists who studied the swine flu decision process concluded that Sencer did not search hard enough for contrary evidence: "Sencer concentrated on the worst case in the short run. So did his superiors. Had they thought equally hard about the likely case in the long run—side effects and suits but no pandemic—the issue of diminishing credibility for CDC would have loomed large. ... Had [Sencer] not sought control of operations it would be still better." One of the *DMs* told them, "Hell, the thing that was needed in planning the swine flu program was a day around the table brainstorming Murphy's Law: 'If anything can go wrong, it will.'"²⁰

V. Variations on the Model

A. Other Mechanisms

Instead of delegating the decision to the expert, *DM* could, in principle, use some other mechanism. Suppose that this mechanism takes the form of a message space M and a function ρ that maps M to $\Delta(\{d_0, d_1\})$, the set of probability distributions over $\{d_0, d_1\}$. The interpretation is that after collecting all the signals he wants to collect, the expert sends a message to *DM*, and the mechanism then outputs a decision. Delegation is equivalent to a mechanism in which, for some $m_0 \in M$, $\rho_0(m_0) = 1$ and, for some $m_1 \in M$, $\rho_1(m_1) = 1$, where $\rho(m) = (\rho_0(m), \rho_1(m))$. In this setting, *DM* can do no better, if the information process is coarse, than to delegate the decision if she uses an adviser:

LEMMA 2: *In the coarse information case, if the expert is E_0 , then delegation is an optimal mechanism for *DM*. If the expert is E_1 , then there is a threshold \hat{c} such that if $c \leq \hat{c}$, delegation is an optimal mechanism for *DM*, and otherwise it is optimal for *DM* to take the decision herself.*

¹⁹King and Crewe (2013, 74, 306).

²⁰Neustadt and Fineberg (1978, 98, 93).

Essentially, the reason for this is that any communication mechanism with commitment, if it responds to the expert's information, is equivalent to delegation because the decision is binary. In framing the mechanism design problem in this way, we make several implicit assumptions. Firstly, the information that the expert collects is soft, i.e., nonverifiable. Secondly, the *DM* is not able to make decision-contingent or message-contingent money payments to the expert. Thirdly, the *DM* cannot design a dynamic mechanism that the expert must play through time, as information collection takes place. The latter may, for example, be due to the fact that she cannot observe how many signals the expert collects, perhaps because the expert, as mentioned above, may be able to collect arbitrary numbers of signals, sequentially, within a period.

Alternatively, *DM* could use a communication game without commitment—that is, the expert decides when to stop searching and then sends a cheap talk message to *DM*, who then decides what action to take. Clearly, this cannot be better for *DM* than the optimal mechanism with commitment.²¹

When the information process is fine, the issue of which mechanism is optimal is substantially more complicated. Consider, for example, a stochastic mechanism with two messages, m_0 and m_1 ; after m_0 , the probability of decision d_0 is α , and after m_1 , it is $\beta \neq \alpha$. Because of the garbling, the expert has less incentive to collect information than in the delegation mechanism. Depending on the distribution of π , this could improve *DM*'s expected payoff relative to delegation. For example, if there is a high probability that π is in the conflict-of-interest region, *DM* may want the expert to make the decision with imperfect information. Characterizing the optimal mechanism in this case seems challenging.

B. More than Two Actions

Although, as we argued in Section II, many important decisions in business, politics and other arenas boil down in practice to binary choices, many other decision problems have a higher dimension than two. Here, we show that many of our results extend to at least some decision frameworks with more than two potential actions.

Suppose that there is a set $D = \{d^1, \dots, d^n\}$ of actions available and that D is partitioned into two disjoint nonempty sets D_0 and D_1 , where $D = D_0 \cup D_1$ and $D_0 \cap D_1 = \emptyset$, such that E_i ($i = 0, 1$) is biased toward decisions in D_i . More specifically, there is a payoff vector (π^1, \dots, π^n) that is distributed according to a joint distribution function H . *DM*'s payoff from action d^j is $U_D(d^j, \pi^j) = \pi^j$, whereas E_i 's payoff is $U_i(d^j, \pi^j) = \pi^j + b$ if $d^j \in D_i$ and otherwise $U_i(d^j, \pi^j) = \pi^j$.

For example, it may be that E_0 is politically right leaning and policies in D_0 are business friendly, whereas E_1 is left leaning and D_1 consists of actions that favor the protection of the environment. Note that there is no conflict of interest between *DM* and either expert *within* an action category, but there is a conflict of interest *across* actions in different categories.

²¹ It can be shown that for small c , there is an equilibrium of this communication game that is equivalent to the equilibrium of the delegation game.

Assume that the experts are ex ante better informed than DM and know which action within each category, D_0 or D_1 , is the best action; call these best actions d_0 and d_1 , respectively. However, they do not know the exact payoff of each action and, consequently, do not know which of the two potentially optimal actions, d_0 or d_1 , is indeed better for DM . Denote $\max\{\pi^j | d^j \in D_i\}$ by π_i ($i = 0, 1$) and $\pi_1 - \pi_0$ by π . Let F be the distribution over π induced by H and assume that $E_F[\pi] < 0$. This framework is then equivalent to the binary-action model studied above. E_0 is a conformist because he is biased in favor of the group of actions that has the higher ex ante expected payoff.

If the expert does not initially know the ranking of payoffs within each set D_i but can acquire costly signals about each action and the information process is fine (i.e., each π^i can be learned with arbitrary precision if enough signals are collected), then our results above for the fine information setting will apply.

In contrast to the binary-action setting, DM will always do better in this multiple-action framework by delegating the decision to E_0 than by taking the decision herself. Note, however, that an information revelation mechanism with commitment could dominate delegation in this case.²²

C. Discounting

If information is coarse and DM discounts the future, using a relatively low discount factor δ , then there is potentially a conflict for DM between speed of decision and accuracy of decision. As we have seen, E_1 provides (for low enough c) the more accurate decision (because he searches longer in expectation), but E_0 is faster (in the case in which there are fixed discrete times at each of which at most one signal can be observed). If δ is low enough, then E_0 will be preferred. It is clear, on the other hand, that Proposition 3 will still apply if DM is sufficiently patient. Furthermore, it is straightforward to show that if the expert also discounts the future, in addition to paying the per-period information-gathering cost, the form of his optimal strategy will be as given in Proposition 2. As c goes to 0 and the discount factor simultaneously goes to 1, the continuation interval expands to fill the interval (\underline{w}, \bar{w}) . Therefore, a small amount of discounting does not affect our result that E_1 is preferred to E_0 . Similarly, it is easy to see that a small amount of discounting would not affect our conclusions for the fine information case.

VI. Hard Evidence

In the previous sections, we assumed that the information collected by the expert is soft (i.e., unverifiable). In this section we consider a setting in which the expert may have hard evidence, such as documents and other kinds of data, that are informative about the state of nature π . We focus on the incentives to disclose

²²Che, Dessein, and Kartik (2013) consider a similar problem in which a biased advisor recommends an action (out of n) to an uninformed DM . There is no problem of information collection; the authors concentrate on the issue of truthful information revelation. They show that a revelation mechanism, in which the expert panders his recommendation toward conditionally better-looking actions, is better for DM than delegation.

evidence.²³ That is, there is a single period, and the expert may, or may not, already be in possession of hard evidence. The evidence (or lack of it) that a given expert E_i has defines his type. The structure of evidence that we assume is the one introduced by Dye (1985), which has been widely used in the economics and accounting literatures to analyze disclosure. In a Dye evidence model, the agent either has no evidence or else has evidence that proves his type. We denote by $\phi \in (0, 1)$ the probability that he has evidence. If he has evidence, he can choose between disclosing this evidence and disclosing nothing. Otherwise, he can disclose nothing. We assume that there is a finite set of possible types. Most of the literature on evidence disclosure assumes type-independent preferences (e.g., Grossman 1981; Milgrom 1981; Glazer and Rubinstein 2004, 2006; and Hart, Kremer, and Perry 2017), an assumption that does not apply in our setting. However, since there are only two possible actions, our model satisfies *simple type dependence* as defined by Ben-Porath, Dekel, and Lipman (2019)—henceforth, BDL—and so we make use of their results.

We can consider either a disclosure game (without commitment by DM) or a mechanism (with commitment). The disclosure game has two stages. In stage 1, E_i makes a cheap talk announcement of a type and either submits evidence that proves his type or else submits no evidence (which, formally, we refer to below as submitting t_0). If he does not have any evidence, then he can only submit t_0 in addition to sending a cheap talk message. At stage 2, the DM , having observed the message and the disclosed evidence, chooses either d_0 or d_1 . The mechanism is the same except that the mechanism itself defines the stage 2 choice, and DM is committed to it.

We denote the finite set of types by T . Each type in T has strictly positive probability, given by the prior distribution over π and the evidence-generating process. The type without evidence is denoted by $t_0 \in T$. For any $t \in T$, let $\pi(t)$ be the expert's posterior expectation of π , derived from the prior on π and the evidence corresponding to t (so $\pi(t_0) = \mu$). We assume that T contains types t such that $\pi(t) \in (0, b)$ and also types t such that $\pi(t) \in (-b, 0)$; that is, there is strictly positive probability that the evidence creates a conflict of interest between the expert, of either type, and DM .²⁴

A deterministic mechanism is a function $\rho: T \times T \rightarrow \{d_0, d_1\}$. Here, $\rho(t_j, t_k) = d_i$ means that after the expert has claimed to be type $t_j \in T$ and submitted evidence $t_k \in T$, the decision is d_i . It is *incentive compatible* for E_i if and only if, for all $t, t_i \in T$, $E(U_i(\rho(t, t), \pi) | t) \geq E(U_i(\rho(t_i, t_0), \pi) | t)$ (it is optimal for him to prove his type if he can).²⁵ By Theorem 1 of BDL, there is an optimal incentive-compatible mechanism for DM that is deterministic, and there is a perfect Bayesian equilibrium (PBE) of the disclosure game that has the same outcome

²³But see the remark at the end of this section about a model in which the expert must search for evidence.

²⁴To see that simple type dependence is satisfied take, e.g., E_0 . Using the notation of BDL (2019, 533), let $u_i(d_0) = 0, u_i(d_1) = 1$ and let $T_i^+ = \{t \in T | \pi(t) > b\}$. By our assumption that E_0 prefers d_0 when $\pi = b$, this represents his preferences. Let $v(d_0, t) = 0$, and define $\bar{v}_i(t)$ in such a way that $v(d_1, t) = \pi(t)$.

²⁵By the appropriate Revelation Principle (Bull and Watson 2007; Deneckere and Severinov 2008), there is no loss of generality in restricting to incentive-compatible mechanisms in which the cheap talk messages are type declarations.

function (mapping types to decisions) as this optimal mechanism. The following Proposition shows that the contrarian is the superior expert in this setting.

PROPOSITION 4: *In the Dye evidence model:*

- (i) *If the expert is a conformist, then in any mechanism—and, hence, in any equilibrium of the disclosure game—the DM’s expected payoff is strictly less than her first-best optimal expected payoff.*
- (ii) *If the expert is a contrarian, then there is an incentive-compatible mechanism, and a PBE of the disclosure game, that achieves the DM’s optimal payoff for any π .*

PROOF:

- (i) We show that any pure strategy PBE of the disclosure game is suboptimal for *DM*. This implies, by Theorem 1 of BDL, that any mechanism, whether deterministic or stochastic, is also suboptimal. In any pure strategy equilibrium, the decision taken after evidence is submitted is pinned down by the posterior belief corresponding to the evidence, so there are potentially three types of equilibrium, distinguished by what happens after no evidence is submitted, whether (a) the decision varies with the cheap talk message, (b) the decision is d_0 , or (c) the decision is d_1 . Case (a) is equivalent to delegation since E_0 can simply withhold evidence, if there is any, and select his optimal decision via the cheap talk message. This is suboptimal for *DM* since, for some types of E_0 , E_0 ’s optimal action differs from *DM*’s. Case (b) too is equivalent to delegation—if E_0 prefers d_1 , then he must have evidence that persuades him that d_1 is better, and this evidence will also persuade *DM* since she is less biased against d_1 than E_0 is. If E_0 prefers d_0 , he can withhold evidence. Equilibria of type (c) are suboptimal for *DM* since d_1 is suboptimal in the event that no evidence exists.²⁶
- (ii) In the case of the contrarian, E_1 , an incentive-compatible first-best optimal mechanism is as follows. $\rho(t, t') = d_0$ if $t' = t_0$ or if $\pi(t') \leq 0$. Otherwise, $\rho(t, t') = d_1$. It is optimal for E_1 to submit evidence if he has it since he can induce d_1 if and only if he has evidence that gives rise to a posterior expectation of π greater than zero, in which case he prefers d_1 to d_0 . It is also optimal for him to tell the truth in the cheap talk message. It is straightforward to construct a corresponding equilibrium of the disclosure game. ■

Essentially, the reason that the contrarian is better than the conformist is that if there is no evidence at all, then d_0 is optimal for *DM*—in a sense, d_0 is the “incumbent” decision. But that implies that the conformist has an incentive to

²⁶Equilibria of types (a) and (b) always exist; those of type (c) exist if ϕ , the probability that evidence exists, is not too low. Which equilibrium is best for *DM* depends on the value of ϕ . There is a threshold for ϕ below which the outcome function of (a) and (b) is better and above which that of (c) is better.

deviate by suppressing evidence in favor of d_1 . There is no corresponding incentive for the contrarian since suppressing evidence can only induce his less favored decision. To put the point another way, the contrarian has more incentive to submit evidence because he is the one who wants to change the *DM*'s mind. The reasoning behind this result clearly relies on the structure of evidence being as assumed by Dye (1985)—in effect, the evidence, if any, is either indivisible or else revelation of part of it exposes the existence of other parts. While this assumption is appropriate for many contexts,²⁷ clearly there are other forms of evidence that may obtain in other contexts.

We have assumed that the two experts have the same absolute size of bias b . Note, however, that Proposition 4 does not require this. For any size of bias $b \in (0, 1)$, the first best is achievable, and the equilibrium is as described above, when the expert is a contrarian with this bias. Furthermore, for any strictly positive bias, there is no first-best equilibrium when the expert is a conformist.

Here, unlike in Section IV, we have assumed that the expert does not engage in costly search for information. Consider an alternative model in which the expert initially has no evidence but can search for it and, with some probability, such evidence exists. Each period of search costs $c > 0$ and implies a fixed strictly positive probability of finding the evidence, if it exists. As soon as he discovers evidence, he plays the disclosure game above; as long as he has not found any evidence, he decides in each period to either search for another period or stop and play the disclosure game. The *DM* cannot observe the number of periods of search. In equilibrium the expert searches until either he discovers evidence or the posterior probability that none exists reaches a threshold, which depends on c . It is straightforward to show that for small c , the equilibrium outcomes of this game approximate those in the static game analyzed above.

VII. Concluding Remarks

The advantages of allowing, indeed encouraging, dissenting voices in the context of policy debates have long been known. (See Dewatripont and Tirole 1999 for an analysis of the merits of adversarial advocacy.) Our results suggest some advantages that contrarians may have that extend beyond their value in the context of debate. A contrarian is less likely to suppress negative evidence, more likely to act on information that is contrary to his initial bias, and more willing to exert effort to collect more information. On the other hand, the optimal choice of expert depends on the nature of the informational environment. In particular, if the decision payoff is normally distributed and information is soft and fine, then a conformist expert has the advantage that a conflict of interest with the principal is less likely than it would be with a contrarian.

At a fundamental level, the choice between a contrarian and a conformist hinges on which has the greater conflict of interest with the *DM*. In the fine nonverifiable information case, this is directly evident. In the cases of verifiable information and

²⁷For example, Che and Severinov (2017) use an evidence disclosure model of this kind to study the effect of legal advice about a client's disclosure decision.

coarse nonverifiable information, the contrarian has the lesser conflict because in each case, he is the one who is more inclined, directly or indirectly, to choose the option he “dislikes,” which in turn is due to the fact that this option is the one favored by the initial beliefs. In the verifiable information case, the conformist is (compared to the contrarian) unwilling to reveal evidence against his “own” action because that action is the default. In the coarse information case, the contrarian is more likely to end up choosing the option he “dislikes,” when it is in fact optimal, because, since that option is ex ante favored, he is likely to search for longer than the conformist.

APPENDIX

PROOF OF LEMMA 1:

Given $t \geq 1$ and $(\pi, h_t) \in \Pi \times H_t$, let $\mu_i(\pi, h_t)$ be E_i 's utility shortfall, relative to E_i 's full information optimum, if he stops at t . In the case of E_1 (the definition for E_0 is similar), this is defined by (i) $\mu_1(\pi, h_t) \equiv \pi + b$ if $\pi > -b$ and $E_{\Phi_{h_t}}(\pi) \leq -b$, (ii) $\mu_1(\pi, h_t) \equiv -\pi - b$ if $\pi \leq -b$ and $E_{\Phi_{h_t}}(\pi) > -b$, and (iii) $\mu_1(\pi, h_t) \equiv 0$ otherwise (e.g., in case (i), the optimum decision is d_1 , giving payoff $\pi + b$, but E_1 's choice is d_0 , giving payoff zero). Then a policy with stopping time T (where T is a random variable) is ex ante optimal for E_i if it minimizes

$$E_{\Psi}[\mu_i(\pi, h_T) + cT],$$

that is, the ex ante expectation, with respect to the joint distribution over π and (infinite) signal sequence, of the total search cost plus the utility shortfall when the decision is taken. Note that an optimal policy must be ex ante optimal.

For each cost $c > 0$, fix a particular optimal policy for E_i ($i = 0, 1$) and denote it by $\sigma_i(c)$. Without loss of generality, suppose the expert is E_1 .

CLAIM 1: *Given any $\xi > 0$, there exists a time $\hat{\tau}(\xi)$ such that if $t > \hat{\tau}(\xi)$,*

$$E_{\Psi}\mu_1(\pi, h_t) < \xi.$$

That is, if E_1 adopts the policy of stopping at t , regardless of the history, and taking the optimal decision given the updated belief at the stopping time, the ex ante expectation of the utility shortfall is bounded by ξ .

PROOF OF CLAIM 1:

Given $t > 0$,

$$E_{\Psi}(\mu_1(\pi, h_t)) = \int_{\Pi} E_{\Psi}(\mu_1(\pi, h_t) | \pi) dF(\pi).$$

Since $\int_{-\infty}^{\infty} |\pi| dF < \bar{\pi}$ for some finite $\bar{\pi}$, there exists $K > 0$ such that

$$\int_{\{\pi: |\pi| > K\}} |\pi| + b dF(\pi) < \frac{\xi}{2}.$$

Therefore, since $\mu_1(\pi, h_t) \leq |\pi| + b$ for all (π, h_t) ,

$$E_{\Psi}(\mu_1(\pi, h_t)) \leq \int_{\{\pi: |\pi| \leq K\}} E_{\Psi}(\mu_1(\pi, h_t) | \pi) dF(\pi) + \frac{\xi}{2}.$$

If $\pi \in (-b - \xi/2, -b + \xi/2)$, then $\mu_1(\pi, h_t) < \xi/2$ for all h_t (for example, if $-b - \xi/2 < \pi < -b$, the utility shortfall is either 0 or $-\pi - b < \xi/2$). Hence, it is sufficient to show that $E_{\Psi}(\mu_1(\pi, h_t) | \pi \in D) < \xi/2$, where

$$D \equiv \{\pi \in \Pi: |\pi| \leq K, \pi \notin (-b - \xi/2, -b + \xi/2)\}.$$

If $\pi \notin (-b - \xi/2, -b + \xi/2)$, then $\mu_1(\pi, h_t) > 0$ only if $|E_{\Phi_{h_t}}(\pi) - \pi| > \xi/2$ (otherwise π and the posterior expectation of π lie on the same side of $-b$). By the definition of a fine information process, for any $\epsilon_2 > 0$, there exists $t(\xi/2, \epsilon_2)$ such that if $t > t(\xi/2, \epsilon_2)$,

$$\begin{aligned} & \Pr_{\Psi} \left[|E_{\Phi_{h_t}}(\pi) - \pi| > \xi/2 | \pi \in D \right] \Pr_{\Psi}(\pi \in D) \\ & + \Pr_{\Psi} \left[|E_{\Phi_{h_t}}(\pi) - \pi| > \xi/2 | \pi \notin D \right] \Pr_{\Psi}(\pi \notin D) < \epsilon_2, \end{aligned}$$

so

$$\Pr_{\Psi} \left[|E_{\Phi_{h_t}}(\pi) - \pi| > \xi/2 | \pi \in D \right] < \frac{\epsilon_2}{\Pr_{\Psi}(\pi \in D)};$$

hence,

$$E_{\Psi}(\mu_1(\pi, h_t) | \pi \in D) < \frac{(K + b)\epsilon_2}{\Pr_{\Psi}(\pi \in D)}$$

since $\mu_1(\pi, h_t)$ has an upper bound of $K + b$ on D . To prove the Claim, take $\epsilon_2 = \left[\xi \Pr_{\Psi}(\pi \in D) \right] \left[2(K + b) \right]^{-1}$.

CLAIM 2: *The ex ante expected utility shortfall to E_1 from the optimal policy $\sigma_1(c)$ converges to zero as c goes to zero.*

PROOF OF CLAIM 2:

Take a decreasing sequence of costs $c_1, c_2, \dots, c_i, \dots$, converging to 0, and $\eta > 0$ (where η is small). Take c_i in the above sequence such that $\hat{\tau}(\eta/2)c_i < \eta/2$. Then, by Claim 1, the sum of the expected shortfall and the expected cost from the policy of stopping at $t = \hat{\tau}(\eta/2)$ is less than $(\eta/2) + (\eta/2) = \eta$. A fortiori, the same is true of $\sigma_1(c_i)$, since the latter is optimal, and so the expected shortfall from $\sigma_1(c_i)$ is less than η for high enough i . This proves Claim 2.

For $h \in H$, let $d_1^c(h)$ be the final decision made if the strategy is $\sigma_1(c)$ and the signal sequence is h ; also, let $H_1^c \equiv \{h \in H: d_1^c(h) = d_1\}$. By Claim 2, E_1 's

expected payoff from his optimal policy converges to his full information optimum, i.e.,

$$\lim_{c \rightarrow 0} \int_{\{(\pi, h): h \in H_1^c\}} (\pi + b) d\Psi = \int_{-b}^{\infty} (\pi + b) dF(\pi).$$

The full information optimal policy for E_1 is to choose d_1 if and only if $\pi > -b$, which implies choosing d_1 with probability $1 - F(-b)$. Therefore, Claim 2 implies that in the limit as $c \rightarrow 0$, E_1 chooses d_1 with probability $1 - F(-b)$ in the sequential model, i.e., $\lim_{c \rightarrow 0} \Psi(\{(\pi, h): h \in H_1^c\}) = 1 - F(-b)$, so

$$\lim_{c \rightarrow 0} \int_{\{(\pi, h): h \in H_1^c\}} \pi d\Psi = \int_{-b}^{\infty} \pi dF(\pi).$$

This proves the Lemma. ■

PROOF OF PROPOSITION 2:

- (i) $W(t)$ is a time-homogeneous Markov process, and therefore the optimal decision at t is determined by $W(t)$, independently of t .

Suppose first that the expert is E_0 . If it is optimal for him at time t to stop and choose d_0 , then it must be that $W(t) \leq b$; conversely, if it is optimal for him to stop and choose d_1 , then $W(t) > b$. For some $W(t)$, it must be optimal to stop and choose d_0 ; otherwise, for any $W(t)$, it would be optimal to choose d_1 immediately; however, if $W(t) < b$, then choosing d_0 is strictly better than this.

Define \underline{W}_0 as the maximum of values of W (in the discrete set \hat{W}) such that it is optimal for E_0 to stop and choose d_0 at conditional expectation W . Then $\underline{W}_0 \leq b$. Consider any $w' < \underline{W}_0$. It must be strictly optimal to stop when the expectation is w' . This is because if the expectation is to reach b from w' , it must pass through \underline{W}_0 , at which point the process will end with decision d_0 ; since d_0 is optimal whether E_0 stops now or continues, it is sub-optimal to continue.

A symmetrical argument applies to high values of $W(t)$, so we conclude that E_0 's optimal policy takes a threshold form, with thresholds \underline{W}_0 and $\overline{W}_0 > \underline{W}_0$.

The two experts have preferences that are symmetric with respect to the two decisions, and by Assumption 1, the stochastic processes determining $W(t)$ and $-W(t)$ are also symmetric. It follows that if d_0 and d_1 are interchanged and $-W(t)$ replaces $W(t)$, the problem faced by E_1 is identical to the one faced by E_0 . Hence, there exist thresholds \underline{W}_1 and $\overline{W}_1 > \underline{W}_1$ such that E_1 stops and chooses d_0 if $W(t) \leq \underline{W}_1$, stops and chooses d_1 if $W(t) > \overline{W}_1$, and otherwise continues. Symmetry implies that $\underline{W}_0 = -\overline{W}_1$ and $\underline{W}_1 = -\overline{W}_0$. Suppose that $\overline{W}_1 \geq \overline{W}_0$. If $W(t) = \overline{W}_0$, then E_0 stops and chooses d_1 , and it is weakly optimal for E_1 to continue. Let θ be the probability that E_1

eventually chooses d_0 , conditional on $W(t) = \overline{W}_0$ (in the case $\overline{W}_0 = \overline{W}_1$, this is the probability of choosing d_0 if he deviates at t by continuing and thereafter plays his optimal strategy). Then the additional expected payoff to E_0 , compared with stopping now, from continuing and following E_1 's optimal strategy, is $\theta(-\underline{W}_1 + b) - c\tilde{t}$, where \tilde{t} is the expected time to stopping, conditional on $W(t) = \overline{W}_0$ and E_1 's strategy being followed. To see this, note that since W is a bounded martingale and the stopping time is almost surely finite, the optional stopping theorem (see Williams 1991) implies that the expected payoff from stopping at \underline{W}_1 or \overline{W}_1 , whichever is hit first, but choosing d_1 in either event, is equal to the current expected payoff from d_1 . Therefore, the expected payoff, ignoring search costs, to E_0 if he follows E_1 's strategy, less his current expected payoff from d_1 , is the probability of choosing d_0 times his expected payoff increment in that event, which is $-\underline{W}_1 + b$. The corresponding expected payoff increment for E_1 , compared with choosing d_1 now, is $\theta(-\underline{W}_1 - b) - c\tilde{t}$, so if E_1 weakly prefers to continue, E_0 strictly prefers to continue. This shows that $\overline{W}_0 > \overline{W}_1$.

- (ii) (a) Suppose that the expectation at time t , $W(t)$, satisfies $\underline{w} < W(t) < b$ and E_0 is the chosen expert. For a given cost c , compare E_0 's expected payoff from (a) stopping now and choosing d_0 (the optimal choice conditional on stopping) and (b) an alternative strategy of continuing at t and thereafter stopping only when either $\{W \leq W(t)\}$ is next hit or $b + \eta$ ($\eta > 0$) is hit, whichever comes first, in the first case choosing d_0 and in the second case choosing d_1 . The additional expected payoff of (b) over that of (a) is $\hat{\theta}(\eta) - c\hat{t}$, where $\hat{\theta}$ is the probability of hitting $b + \eta$ first and \hat{t} is the expected stopping time. This is because in the event of hitting $b + \eta$ first, DM 's expected gain from d_1 rather than d_0 is $b + \eta$, and so E_0 gains this, less his bias parameter b . Since \hat{t} is finite,²⁸ this quantity is strictly positive for low enough c , so it cannot be optimal to stop when the expectation is $W(t)$. A symmetric argument for the case $W(t) \geq b$ then establishes the result. (b) To prove monotonicity of $\underline{W}_0(c)$ (the proof for $\overline{W}_0(c)$ is similar), consider \underline{c} and \bar{c} such that $0 < \underline{c} < \bar{c}$, and suppose that $\underline{W}_0(\bar{c}) < \underline{W}_0(\underline{c})$. Denote by $\sigma_0(c)$ E_0 's optimal strategy when the cost is c . Conditional on $W(t) = \underline{W}_0(\underline{c})$, the expected gain from playing $\sigma_0(\bar{c})$ rather than $\sigma_0(\underline{c})$ in the continuation, when the cost is c , is given by $\hat{\theta}(\overline{W}_0(\bar{c}) - b) - c\hat{t}$, where $\hat{\theta}$ is the probability, conditional on $W(t) = \underline{W}_0(\underline{c})$, of hitting $\overline{W}_0(\bar{c})$ before $\underline{W}_0(\bar{c})$ and \hat{t} is the expected stopping time of $\sigma_0(\bar{c})$, conditional on $W(t) = \underline{W}_0(\underline{c})$. By definition, $\hat{\theta}(\overline{W}_0(\bar{c}) - b) - \bar{c}\hat{t} \geq 0$. Therefore, $\hat{\theta}(\overline{W}_0(\bar{c}) - b) - \underline{c}\hat{t} > 0$, which contradicts the fact that it is optimal for E_0 to stop at $\underline{W}_0(\underline{c})$ when the cost is \underline{c} . (c) For high enough c , it is clearly sub-optimal to search. ■

²⁸By the analysis of the gambler's ruin problem; see Cox and Miller (1965).

PROOF OF LEMMA 2:

Suppose that E_1 is the expert, and take an optimal mechanism, ρ^* , for DM , together with an optimal strategy for E_1 , given the mechanism. This is a search strategy together with a message strategy that maps the posterior expectation of π , after search stops, to a probability distribution over messages. Suppose that the mechanism is responsive to the expert, in the sense that there exist messages m and m' in M such that $\rho(m) \neq \rho(m')$.

We can take the expert's type to be his posterior expectation of π . Given type $\pi \in \Pi$ and decision $d \in \{d_0, d_1\}$, let $\Pr(d|\pi)$ be the probability of decision d when the expert has type π . Suppose that there exist π, π' , and π'' in Π such that $\Pr(d_0|\pi) < \Pr(d_0|\pi') < \Pr(d_0|\pi'')$. Then $\Pr(d_0|\pi') \in (0, 1)$, and so, by our tie-breaking rule (the expert chooses d_0 when indifferent), type π' has a profitable deviation, either to the message strategy of type π'' , if he weakly prefers d_0 to d_1 , or to that of type π , if he strictly prefers d_1 . This contradiction shows that there can only be two values of $\Pr(d_0|\pi)$ as π varies. Therefore, without loss of generality, we can assume that the optimal mechanism, ρ^* , has only two messages: m_0 , sent by types who weakly prefer d_0 , and m_1 , sent by types who strictly prefer d_1 .

Let $\rho_1^*(m_1) \equiv \beta$ and $\rho_1^*(m_0) \equiv \alpha < \beta$. If ρ^* is not a delegation mechanism then $0 < \beta - \alpha < 1$ (for a delegation mechanism $\beta - \alpha = 1$). Suppose the optimal mechanism is not a delegation mechanism. If, when he stops, E_1 strictly prefers d_1 (and hence sends message m_1), his expected payoff, gross of search cost, is $\beta(\pi + b)$; if he prefers d_0 , it is $\alpha(\pi + b)$. That is, the model is equivalent to a sequential search model as analyzed in the text except that the payoffs are given by $\tilde{U}_1(d_1, \pi) = \beta(\pi + b)$ and $\tilde{U}_1(d_0, \pi) = \alpha(\pi + b)$. Adapting the proof of Proposition 2, it follows that E_1 's search strategy is defined by thresholds $\underline{W}_1(\alpha, \beta, c) \leq -b$ and $\overline{W}_1(\alpha, \beta, c) \geq -b$, where we can take it that $\underline{W}_1(\alpha, \beta, c) < 0 < \overline{W}_1(\alpha, \beta, c)$; otherwise, DM would do better not to use the expert and to choose d_0 instead, so the optimal mechanism would not be responsive to the expert.

Let DM 's expected payoff, given ρ^* , when E_1 uses this strategy be $V_D(\alpha, \beta, c)$. Suppose that DM considers instead a mechanism ρ' characterized by α' and β' , where $\alpha' \leq \alpha$ and $\beta' \geq \beta$, with at least one of these inequalities strict.

Denote $\underline{W}_1(\alpha', \beta', c)$ by \underline{W}'_1 , $\overline{W}_1(\alpha', \beta', c)$ by \overline{W}'_1 , $\underline{W}_1(\alpha, \beta, c)$ by \underline{W}_1 , and $\overline{W}_1(\alpha, \beta, c)$ by \overline{W}_1 .

Then $\underline{W}'_1 \leq \underline{W}_1$ and $\overline{W}'_1 \geq \overline{W}_1$. To see this, suppose that $\underline{W}'_1 > \underline{W}_1$. Given parameters (α, β) , E_1 strictly prefers to continue at \underline{W}'_1 (recall our assumption that E_1 stops if indifferent). Given (α', β') , the expected cost, at \underline{W}'_1 , of continuing and using the thresholds \underline{W}_1 and \overline{W}_1 is the same as for (α, β) , but the expected benefit is higher. This is because this expected benefit in the case of (α, β) is $(\beta - \alpha)(\overline{W}_1 + b)$ multiplied by the probability of reaching \overline{W}_1 before \underline{W}_1 , whereas in the case of (α', β') , it is $(\beta' - \alpha')(\overline{W}_1 + b)$ multiplied by the same probability. Therefore, it must be optimal to continue at \underline{W}'_1 given parameters (α', β') . Contradiction. The case $\overline{W}'_1 < \overline{W}_1$ is ruled out similarly.

$V_D(\alpha', \beta', c) > V_D(\alpha, \beta, c)$, where $V_D(\alpha', \beta', c)$ is DM 's expected payoff given mechanism ρ' , when E_1 uses his corresponding optimal strategy. To see this, note

that, conditional on reaching \underline{W}_1 before \overline{W}_1 , $V_D(\alpha, \beta, c) = \alpha \underline{W}_1$, whereas, using the notation and arguments of the proof of Proposition 3, conditional on reaching \underline{W}_1 before \overline{W}_1 ,

$$V_D(\alpha', \beta', c) = \alpha' \underline{W}_1 + \left(1 - \theta\left(\frac{\underline{W}'_1, \overline{W}'_1}{\underline{W}_1}\right)\right)(\beta' - \alpha') \overline{W}'_1 \geq \alpha' \underline{W}_1 \geq \alpha \underline{W}_1.$$

A similar argument shows that $V_D(\alpha', \beta', c) \geq V_D(\alpha, \beta, c)$, conditional on reaching \overline{W}_1 before \underline{W}_1 ; furthermore, at least one of these two inequalities is strict. This shows that a nondelegation mechanism that is responsive to the expert can be improved upon. A similar argument gives the same conclusion for E_0 .

If the expert is E_0 , then delegation is weakly better for DM than the optimal nonresponsive mechanism (i.e., taking action d_0 for sure). If the expert is E_1 , then delegation is better than taking action d_0 for sure if and only if c is sufficiently low, so that $\underline{W}_1(c) < 0 < \overline{W}_1(c)$. ■

REFERENCES

- Aghion, Philippe, and Jean Tirole. 1997. "Formal and Real Authority in Organizations." *Journal of Political Economy* 105 (1): 1–29.
- Alonso, Ricardo, and Niko Matouschek. 2005. "Optimal Delegation." *Review of Economic Studies* 75 (1): 259–93.
- Argenziano, Rossella, Sergei Severinov, and Francesco Squintani. 2016. "Strategic Information Acquisition and Transmission." *American Economic Journal: Microeconomics* 8 (3): 119–55.
- Akerlof, George A., and William T. Dickens. 1982. "The Economic Consequences of Cognitive Dissonance." *American Economic Review* 72 (3): 307–19.
- Benabou, Roland. 2013. "Groupthink: Collective Delusions in Organizations and Markets." *Review of Economic Studies* 80: 429–62.
- Ben-Porath, Elchanan, Eddie Dekel, and Barton L. Lipman. 2019. "Mechanisms with Evidence: Commitment and Robustness." *Econometrica* 87 (2): 529–66.
- Blackwell, David, and Lester Dubins. 1962. "Merging of Opinions with Increasing Information." *Annals of Mathematical Statistics* 33 (3): 882–86.
- Bull, Jesse, and Joel Watson. 2007. "Hard Evidence and Mechanism Design." *Games and Economic Behavior* 58 (1): 75–93.
- Calvert, Randall L. 1985. "The Value of Biased Information: A Rational Choice Model of Political Advice." *Journal of Politics* 47 (2): 530–55.
- Che, Yeon-Koo, Wouter Dessein, and Navin Kartik. 2013. "Pandering to Persuade." *American Economic Review* 103 (1): 47–79.
- Che, Yeon-Koo, and Navin Kartik. 2009. "Opinions as Incentives." *Journal of Political Economy* 117 (5): 815–60.
- Che, Yeon-Koo, and Konrad Mierendorff. 2019. "Optimal Dynamic Allocation of Attention." *American Economic Review* 109 (8): 2993–3029.
- Che, Yeon-Koo, and Sergei Severinov. 2017. "Disclosure and Legal Advice." *American Economic Journal: Microeconomics* 9 (2): 188–225.
- Cox, D.R., and H.D. Miller. 1965. *The Theory of Stochastic Processes*. London: Chapman and Hall.
- Crawford, Vincent P., and Joel Sobel. 1982. "Strategic Information Transmission." *Econometrica* 50 (6): 1431–51.
- Cukierman, Alex, and Mariano Tommasi. 1998. "When Does It Take a Nixon to Go to China?" *American Economic Review* 88 (1): 180–97.
- DeGroot, Morris H. 1970. *Optimal Statistical Decisions*. New York: McGraw-Hill.
- Deneckere, Raymond, and Sergei Severinov. 2008. "Mechanism Design with Partial State Verifiability." *Games and Economic Behavior* 64 (2): 487–513.
- Dewatripont, Mathias, and Jean Tirole. 1999. "Advocates." *Journal of Political Economy* 107 (1): 1–39.

- Dur, Robert, and Otto H. Swank.** 2005. "Producing and Manipulating Information." *Economic Journal* 115 (500): 185–99.
- Dye, Ronald A.** 1985. "Disclosure of Nonproprietary Information." *Journal of Accounting Research* 23 (1): 123–45.
- Fellner, C.H., and J.R. Marshall.** 1970. "Kidney Donors—The Myth of Informed Consent." *American Journal of Psychiatry* 126 (9): 1245–51.
- Festinger, Leon.** 1957. *A Theory of Cognitive Dissonance*. Stanford, CA: Stanford University Press.
- Gerardi, Dino, and Leeat Yariv.** 2008. "Costly Expertise." *American Economic Review* 98 (2): 187–93.
- Goldman Sachs.** 2019. "Vodafone Acquires Mannesmann in the Largest Acquisition in History." <https://www.goldmansachs.com/our-firm/history/moments/2000-vodafone-mannesmann-merger.html> (accessed February 3, 2021).
- Glazer, Jacob, and Ariel Rubinstein.** 2004. "On Optimal Rules of Persuasion." *Econometrica* 72 (6): 1715–36.
- Glazer, Jacob, and Ariel Rubinstein.** 2006. "A Study in the Pragmatics of Persuasion: A Game Theoretical Approach." *Theoretical Economics* 1: 395–410.
- Grossman, Sanford J.** 1981. "The Informational Role of Warranties and Private Disclosure about Product Quality." *Journal of Law and Economics* 24 (3): 461–83.
- Hart, Sergiu, Ilan Kremer, and Motty Perry.** 2017. "Evidence Games: Truth and Commitment." *American Economic Review* 107 (3): 690–713.
- Holmstrom, Bengt.** 1977. "On Incentives and Control in Organizations." PhD diss. Stanford University.
- Holmstrom, Bengt.** 1984. "On the Theory of Delegation." In *Bayesian Models in Economic Theory*, edited by M. Boyer and R. Kihlstrom, 115–41. New York: North-Holland.
- Janis, Irving L.** 1972. *Victims of Groupthink: A Psychological Study of Foreign-Policy Decisions and Fiascos*. Boston: Houghton Mifflin.
- Janis, Irving L., and Leon Mann.** 1977. *Decision Making*. New York: The Free Press.
- King, Anthony, and Ivor Crewe.** 2013. *The Blunders of Our Governments*. London: Oneworld.
- Klapper, Joseph T.** 1949. "The Effects of Mass Media." Unpublished.
- Li, Hao.** 2001. "A Theory of Conservatism." *Journal of Political Economy* 109 (3): 617–36.
- Li, Hao, and Wing Suen.** 2004. "Delegating Decisions to Experts." *Journal of Political Economy* 112 (S1): 311–35.
- Melumad, Nahum D., and Toshiyuki Shibano.** 1991. "Communication in Settings with No Transfers." *RAND Journal of Economics* 22 (2): 173–98.
- Meyer, Margaret A.** 1991. "Learning from Coarse Information: Biased Contests and Career Profiles." *Review of Economic Studies* 58 (1): 15–41.
- Milgrom, Paul R.** 1981. "Good News and Bad News: Representation Theorems and Applications." *Bell Journal of Economics* 12 (2): 380–91.
- Neustadt, Richard E., and Harvey V. Fineberg.** 1978. *The Swine Flu Affair*. Washington, DC: National Academies Press.
- Oakervee Review.** 2019. *Oakervee Review of HS2*. London: Oakervee Review. <https://www.gov.uk/government/publications/oakervee-review-of-hs2>.
- Rabin, Matthew.** 1994. "Cognitive Dissonance and Social Change." *Journal of Economic Behavior and Organization* 23 (2): 177–94.
- Suen, Wing.** 2004. "The Self-Perpetuation of Biased Beliefs." *Economic Journal* 114 (495): 377–96.
- Szalay, Deszö.** 2005. "The Economics of Clear Advice and Extreme Options." *Review of Economic Studies* 72 (4): 1173–98.
- Szalay, Deszö.** 2012. "Strategic Information Transmission and Stochastic Orders." Unpublished.
- Wald, Abraham.** 1947. *Sequential Analysis*. New York: Wiley.
- Williams, David.** 1991. *Probability with Martingales*. Cambridge, UK: Cambridge University Press.