Sequential Expert Advice: Superiority of Closed-Door Meetings*

Parimal K. Bag† Tridib Sharma ‡

March 2, 2018

Abstract

Sequential advice to a decision maker by experts with career concerns has previously been studied under transparency. In a two-expert model and a Bayesian decision maker (D), our finding is that secrecy (weakly) dominates transparency, yielding superior decisions for D. The subtle insight is that secrecy empowers experts moving late to be pivotal more often. We show two other results: (i) only secrecy enables the second expert to partially communicate her information and its high precision level to D and swing the decision away from the first expert's recommendation; (ii) if the experts on average have high precision, then the second expert is effective only under secrecy.

We also show that the efficacy of secrecy is further enhanced if either (1) the experts are allowed to revise their advice following each other's initial recommendation (i.e., deliberate), or (2) the experts are allowed to make detailed recommendations (i.e., give advice and also report its quality). Expanding the message space under deliberation or under detailed recommendation, both generate the possibility of fully revealing equilibria, leading to informationally efficient choice by D.

JEL Classification: D82, D83, D23, C72. *Key Words:* Career concern, cheap talk, sequential advice, deliberations, transparency, secrecy, signal revelation, partial type revelation.

***Acknowledgements:** We thank Masaki Aoyagi and three anonymous referees for detailed comments and suggestions on an earlier draft. Also, for criticisms and related observations at various stages of this work, we thank Murali Agastya, Sourav Bhattacharya, Martin Bodenstein, Lawrence Christiano, Sudipto Dasgupta, Alessandro Lizzeri, Andy McLellan, Nicolas Melissas, Arijit Sen, Levent Ulku, and seminar/conference participants at University of Queensland, University of Durham, Singapore Management University, Indian Institute of Management (Calcutta and Bangalore), Delhi School of Economics, ISI-Delhi Growth and Development Conference, Canadian Economic Theory Conference, and the Royal Economic Society Conference at Bristol. Neng Qian provided research assistance, and a supporting research grant from the Singapore Ministry of Education is gratefully acknowledged. For errors or omissions, we remain responsible.

†Department of Economics, National University of Singapore, Faculty of Arts and Social Sciences, AS2 Level 6, 1 Arts Link, Singapore 117570; E-mail: ecsbpk@nus.edu.sg

‡ITAM, Mexico; E-mail: sharma@itam.mx

1 Introduction

Many decisions are based on the advice given sequentially by multiple experts. Parliamentary sessions debate over issues of national importance where members of different political parties give their views and respond to others' views on the right course of action. Examples abound – whether to go to war, what health reforms to implement, whether to tighten immigration policy, agree to a trade treaty, or break away from a club such as the European Union on economic merit. Each such action taken will later on prove to be either a right or wrong action. Similarly, in modern corporations CEOs may consult their deputies for opinions on crucial decisions such as business expansion, joint venture, r&d, etc.

In a majority of the cases, however, a decision maker due to time constraint may not have the luxury of getting the opinion of each expert scrutinized by others. Also, the advice provided may not be *able* to adequately transmit all the relevant information. For instance, part of the information is about the quality of recommendation which reflects on the expert's expertise. Experts rarely admit that the quality of their advice might not be high. Decision maker therefore has to manage with the limited information that the expert voluntarily transmits instead of thinking that the expert would self-certify the quality of advice. Another reason could be imperfections in language or comprehension. For example, a project appraisal report may not include the intricate properties of its statistical estimators. Our primary focus in this paper will be on a scenario where two experts sequentially provide cheap talk advice in two stages. The advice is constrained by the messages available.

Our secondary focus will be on two alternative procedures of advice – deliberation and detailed recommendation. Under deliberation experts have the option of revising their recommendations in light of the opinion of others.[1] Deliberations are also sequential in nature with commentators giving their views, endorsing or countering prevailing opinions. Prolonged deliberations may allow experts to potentially transfer all their information, overcoming language and comprehension imperfections. In monetary policy committees, select members debate over the inflation target or interest rate. The head of states such as the Prime Minister, or the leader of a major political party, may rely on their trusted advisors, who give their conflicting views about the suitability of important decisions in closed-door meetings in a more discussion and deliberative style. A judge may conduct cases *in-camera* (or *in-chambers*), i.e., hear evidence and arguments in private as opposed to open court

[1] This is of secondary focus because the main result here is in accordance with the intuition derived from the analysis of the sequential advice model.

trial.[2] Detailed recommendation captures the case where experts are not constrained by limitations in language or comprehension.

We ask whether a decision maker, in order to select a decision corresponding with an unknown *state*, should make the experts' recommendations public, by disclosing who made what recommendations when (transparency), or maintain secrecy by informing the public of only the summary decision. The true state will eventually become known. One principal assumption will be that the experts have "career concerns," an idea in organizational economics introduced by Holmstrom (1999): they care mainly about the perception of outsider(s), i.e. the public, about their ability in predicting the true state. In the examples of parliamentary and political debates as well as advice by corporate managers there is no explicit payment involved per each advice. Instead, giving opinions are seen as part of the job.[3]

When the decision maker has the final authority to make a yes/no decision, he may base his decision not just on the number of each type of recommendation but also on the order of recommendations. This is so especially when the experts can see and, hence, learn from each other's recommendation. In choosing to take the same decision following the same recommendations, but in different orders, the decision maker may be wasting valuable information. In this regard, sequentiality of advice introduces a new dynamic not prevalent in simultaneous decision making with an exogenous rule, e.g., in one-shot voting.

In the two-stage sequential advice game, with constrained message sets, our main result is that a Bayesian decision maker would always prefer the most informative equilibrium under secrecy over any of the equilibrium under transparency (Proposition 4). One key reason for this is that secrecy enables late recommendations to often be decisive even if early recommendations come from fairly high (average) quality experts (Corollary 3); under transparency an early opinion is more likely to stifle valuable late views (Proposition 1). In addition, only secrecy creates a positive chance for the second expert's intrinsic high

[2]See Jan. 4, 2010 report, "Democratic leaders plan secret health reform deliberations"; source: http://www.usnews.com/opinion/blogs/peter-roff/2010/01/04/democratic-leaders-plan-secret-health-reform-deliberations. A recently edited book titled "*Secrecy and Publicity in Votes and Debates*" (2015; edited by Jon Elster) discuss a range of issues involving secret vs. open debates in politics. In Chapter 12 ("Secret Votes and Secret Talk"), one of the authors, John Ferejohn, discusses publicness of votes (leading to decisive public choices) by elected representatives but also talks about background political deliberations that "are either completely or partly veiled from outside scrutiny." See also some of the other chapters, especially, "Secret-Public Voting in FDA Advisory Committees" by P. Urfalino and P. Costa.

[3]Politicians (Congressmen in the USA) appeal to a higher level audience – heads of important committees in the House or the Senate, electorate who may win him or her a Senate seat in future election, etc. 'Skills' and 'achievements' (e.g., "sponsoring major pieces of legislation, delivering famous speeches, casting decisive votes on important issues" etc.) are reckoned to be important considerations in influencing a politician's career prospects (see Diermer, Keane and Merlo, 2005).

ability to be communicated to the decision maker, thus helping him make better decisions (Proposition 2). When the restriction on message sets is removed we have an additional result: under secret deliberation as well as detailed recommendation, we show that at times all experts are able to reveal all their information (Proposition 6). Surprisingly, the scope of information revelation is the same under both these formats of communication.

Interactions among three principal forces shape the experts' information revelation incentives: the prior bias, the lead opinion bias, and conformity bias. The first and second biases are exogenous in our setup and impact on the experts' beliefs about the likely state. If one state is favored over the other in terms of the prior beliefs, then we say that there is a *prior bias*. For any given prior bias favoring a specific state, the *lead opinion bias* suggests how seriously the second expert should view the first expert's recommendation. As the experts' average precision levels improve, so does the lead opinion bias. The third bias, if and when it exists (and it can arise only endogenously), is induced by the beliefs of the outsider. Typically, when an expert's recommendation (actual or perceived) does not match the realized state, the outsider updates his beliefs towards the expert having low predictive skills. If such downward revision were to be greater when the alternative to status quo is found to have been wrongly predicted, one would expect experts to sometimes recommend in favor of the status quo even if they expect it to occur with probability less than one-half. This we call the *conformity bias*. It inclines experts toward recommending the decision favored by the prior.

With transparency, at times the second expert is able to to overturn a recommendation made against the favored state. But under secrecy, the second expert may also be able to overturn a recommendation made in support of the favored state. For small prior biases, the lead opinion bias comes into play. Under transparency, when the lead opinion bias is large the second expert herds and thus becomes redundant. Under secrecy, however, an endogenous conformity bias acts against the lead opinion bias, making the second expert's recommendation relevant over a larger range of parameters (Propositions 3 and 1, and Lemma 7). Even in the absence of conformity bias, under secrecy, the second expert may be able to sometimes communicate, through *partial herding*, all relevant information – her signal about the state as well as its quality when the quality is high (Proposition 2). Surprisingly, this is so even when advice is constrained by a limited message set. For these reasons, the decision maker favors secrecy over transparency.

■ **Related literature.** Our paper follows the sequential cheap-talk advice literature started by Scharfstein and Stein (1990), and studied extensively by Ottaviani and Sorensen (2001;

2006a,b,c). These papers focused on how financial experts, who care about their reputation in predicting assets' returns or an unknown *state*, tend to herd in their recommendations, or conform to some prior expectation of the unknown state.[4] While these papers address an important class of problems, what has not been considered before is whether the transparent sequential advice considered by these authors is indeed an ideal protocol for good decision making. Studying within the same sequential advice framework, this paper argues that conducting the sequential advice in closed-door meetings, which is very plausible in many political and organizational decision contexts, yields clearly superior decisions.

Also related are the papers of Levy (2007a,b), and Visser and Swank (2007) who have studied information transmission by experts but in the context of voting. There are important differences between these two papers and our model: first of all, sequencing of expert advice in our analysis (as opposed to simultaneous voting) allows the second expert to learn from the first expert's action. Second, our decision maker does not commit to a decision rule (as in voting) and instead updates his prior and optimally selects a decision.

Levy analyzes a committee decision model using voting. Three experts, motivated by career concerns, simultaneously and independently vote on an action each, and the decision is determined by a given voting rule (unanimity or majority rule). The main argument is that with secretive voting experts are more likely to conform to pre-existing biases either in the voting rule or in the prior, while transparency often leads to contrarian voting. One of Levy's main findings is that under the unanimity rule, secretive voting may *sometimes* induce better decisions than a transparent procedure.[5]

Visser and Swank study a somewhat different model where career-concerned experts with private signals about the suitability of a public project engage in simultaneous information exchange, followed by voting. Smart experts observe the accurate information whereas dumb experts observe completely uninformative signal, and the exerts do not know whether they are dumb or smart. The authors find that transparency aligns experts' interests better with the first-best (or public) objective. Under secrecy, experts' behaviors are characterized by conformity and strategic suppression of individual information. Our results on the superiority of secret advice and/or deliberation contrast with their result on the superiority of transparent deliberation. Besides simultaneity of information exchange and voting (as opposed to

[4] Austen-Smith (1993) studied a sequential referral mechanism involving biased experts on the question of suitability of a legislative bill. One of the concerns was whether sequential referral is better than simultaneous referral. Our setting is different in that the experts are disinterested in what decision gets implemented.

[5] The question of transparency has been analyzed in other applications also by Sibert (2003), Gersbach and Hahn (2008), Seidmann (2011), among others.

our sequential advice/deliberation), Visser and Swank made an important assumption, that the project's implementation by itself would indicate to the outside world that the majority of experts got the correct signal and hence could agree to implement the project; as such the market does not get to observe the ex-post accuracy of the decision taken. Thus there is a group bias towards conformity. Transparency, on the other hand, kills off the incentive to wrongly implement the project because any disagreement in the information exchange will already have hurt the group's collective reputation. In contrast, in our setup (and also in Levy) the market gets to see the actual state and hence conformism, per say, does not enhance reputation.

Issues of deliberations in decision making have been analyzed in both voting models and mechanism design setting but where experts do not necessarily have career concerns. Jackson and Tan (2013) analyze a simultaneous move binary voting model with a prior deliberations phase added in, where multiple experts, simultaneously, can either truthfully reveal or conceal their privately observed signals with no possibility of cheap-talk style misrepresentation. Wolinsky (2002) considered how allowing partial communication between two experts with different pieces of information relevant for a decision, before the experts communicate their information to a decision maker, helps making a better decision, compared to full or no communication. The communication between the experts in Wolinsky are hidden from the decision maker. Meade and Stasavage (2008) have shown that after meetings of the Federal Reserve Board, and in these meetings experts (with career concerns) report their private signals sequentially and then vote sequentially, conformity increased markedly if the individual reports, that the authors call deliberations, are made public.[6] Recently Iaryczower, Shi and Shum (2016) structurally estimate the effect of pre-vote deliberation by judges on decisions of US appellate courts and find that it can improve or worsen the accuracy of decisions. The works of Jackson and Tan, and Wolinsky are primarily about information aggregation by specific voting rules/mechanism for an exogenously given deliberation protocol. In particular, their deliberations, including the one in Iaryczower et al., take the form of *simultaneous* exchange of information, as opposed to our sequential advice or back-and-forth exchanges and gradual revelation.

Finally, our paper relates to the communication games literature. In a cheap-talk, two-person bimatrix game where one player is better informed about the game to be played than

[6]This herding result is along similar lines as in Ottaviani and Sorensen (2001) and our Proposition 1. Our back-and-forth advice/reporting in Section 7 captures the spirit of deliberations better than the sequential reporting protocol of Meade and Stasavage (2008).

the other, Aumann and Hart (2003) have shown the merit of protracted exchange of messages for mutually beneficial information revelation.[7] They show that trying to achieve revelation "too soon" will unravel, in contrast to our result in Proposition 6 on the equivalence of the equilibrium revelation sets.[8]

In the next two sections we present the decision maker's problem, the advice protocols and a technical result on partition of expert types. The core analysis is developed in Sections 4–7. In Section 8, we discuss how simultaneous advice alters some of the results. Final remarks appear in Section 9. The proofs are relegated to Appendix A and a Supplementary file.

2 Decision maker's problem

A decision maker, D, has to solicit recommendations (advice) from two experts. There is an outside observer O, to be referred as the public or the "market", whose evaluation of the experts' abilities confers the only benefits (payoffs) on the experts.

Formally, two experts make their recommendations to D sequentially about a payoff relevant state $\omega \in \{a, b\}$. Throughout e is a generic label for an expert, with the first mover referred as i and second mover as j. The two experts, D and O share a common prior on the states that favor state a: $\Pr(a) = q$, where $q \in (\frac{1}{2}, 1)$ will be referred as the ***prior bias***.

Each expert privately observes a signal $s_e \in \{\alpha, \beta\}$. Let

$$\Pr(s_e = \alpha \mid \omega = a) = \Pr(s_e = \beta \mid \omega = b) = t_e$$

be the quality of expert e's signal, that we call e's precision level (or ability) $t_e \in \{\xi, \lambda\}$, with $\frac{1}{2} < \xi < \lambda < 1$. Experts are privately informed about their abilities that are i.i.d., with $\Pr(t_e = \lambda) = \theta$, $0 < \theta < 1$ for $e = i, j$. Let

$$k \equiv \theta\lambda + (1 - \theta)\xi$$

[7]Forges and Koessler (2008) analyze multi-round but unmediated message games for certifiable types and show how delayed information certification and multi-round communication is required to achieve some equilibrium payoffs. See also Krishna and Morgan (2004).

[8]Of course the two models are not directly comparable: ours involve two experts and a decision maker, whereas Aumann and Hart (2003) has only one informed party; our experts do not take any outcome relevant action, only the decision maker does; our imperfectly informed experts take turns in talking, whereas in their model both the informed and the uninformed talk simultaneously in a sequence of (possibly infinite number) rounds before choosing actions; etc.

be the prior that any expert will observe the correct signal. The index k will also measure the influence of expert i, i.e., the first mover, on expert j's recommendation and will be referred as the ***lead opinion bias***.

Define

$$T_i = T_j = \{(\alpha,\xi),(\alpha,\lambda),(\beta,\xi),(\beta,\lambda)\},$$

where the elements of T_i and T_j represent the private information (types) of experts i and j, and are denoted by τ_i and τ_j.

An expert is randomly drawn by D with probability $\frac{1}{2}$ to move first. After observing i's recommendation to D, expert j makes her recommendation. Let $V_i : T_i \to \{A, B\}$ and $V_j : T_j \times \{A, B\} \to \{A, B\}$ be the pure strategies by experts i and j, with a typical recommendation profile $(\nu_i, \nu_j(\nu_i))$ denoted by ν. Thus, $\nu \in V \equiv \{A, B\} \times \{A, B\}$. Denote the set of strategies of expert e as:

$$\mathbf{V}_e = \{(\nu_e(\alpha,\xi),\nu_e(\alpha,\lambda),\nu_e(\beta,\xi),\nu_e(\beta,\lambda)) \mid \nu_e(s_e,t_e) \in \{A,B\}\}. \tag{1}$$

Though we do not explicitly write the first expert's recommendation in $\mathbf{V}_j$, j's strategy will depend on i's recommendation as in the specification V_j.

Note that each expert has four two-dimensional types. Furthermore, the message set $\{A, B\}$ is of cardinality two.

D uses a **Bayesian decision rule** $d : V \to \{A, B\}$ to choose between actions A and B. After the decision, the *true* state is revealed and D receives a payoff $\pi_D(d, \omega)$, where

$$\begin{aligned} \pi_D(A,a) = \pi_D(B,b) &= 1, \\ \pi_D(B,a) = \pi_D(A,b) &= 0. \end{aligned} \tag{2}$$

Thus, action A (resp. B) is D's ideal decision in state a (resp. b).

All of the above, except realizations of types, states and signals, are common knowledge among experts, D and O.

We now state the two protocols and the payoffs of the experts.

[***Transparency* or $\wp = t$**] O *observes* d*, the state of the world* ω *and the sequence of moves (which expert moves first and which second) as well as the recommendations made by the experts. In particular,* O *observes a realization of the outcome,* $(\nu_i, \nu_j, d, \omega)$*, and Bayes-updates his beliefs regarding the experts' abilities denoted by* $\Pr(t_i \mid \nu_i, \nu_j, d, \omega)$*. The* *expected abilities of* i *and* j*, as well as their payoffs, are*

$$\begin{aligned} E^{\wp=t}(t_i|\nu_i,\nu_j,d,\omega) &= \Pr(t_i=\lambda \mid \nu_i,\nu_j,d,\omega)\lambda + \Pr(t_i=\xi \mid \nu_i,\nu_j,d,\omega)\xi, \\ E^{\wp=t}(t_j|\nu_i,\nu_j,d,\omega) &= \Pr(t_j=\lambda \mid \nu_i,\nu_j,d,\omega)\lambda + \Pr(t_j=\xi \mid \nu_i,\nu_j,d,\omega)\xi. \end{aligned} \tag{3}$$

[***Secrecy* or** $\wp = s$] O *only observes the decision maker's decision and the true realization of the state,* (d,ω)*, and Bayes-updates expert* e*'s ability to* $\Pr(t_e \mid d,\omega)$*. The expected ability of* e *is*

$$E^{\wp=s}(t_e|d,\omega) = \Pr(t_e=\lambda \mid d,\omega)\lambda + \Pr(t_e=\xi \mid d,\omega)\xi. \tag{4}$$

This is also each expert's payoff.

D does not offer any explicit monetary rewards to the experts. The market pays the experts based on their expected absolute abilities. These expectations depend on what the market can observe, i.e., on the protocol of advice.

Let $\mu_i^\wp = \Pr(\omega = a|\tau_i)$ and $\mu_j^\wp = \Pr(\omega = a, \tau_i|\nu_i, \tau_j)$ denote expert i and j's beliefs about ω, and in the case of j also about i's type, conditional on the expert's private information. Let $\mu_D^\wp = \Pr(\omega = a|\nu_i, \nu_j)$ denote D's updated belief conditional on the recommendations. O's beliefs are denoted by $\mu_O^{\wp=t} = \Pr(t_e = \lambda|\nu_i, \nu_j, d, \omega)$ and $\mu_O^{\wp=s} = \Pr(t_e = \lambda|d, \omega)$ under transparency and secrecy, respectively.

This ends the description of the two games, conditional on the protocols – transparency and secrecy. We build our equilibrium solution starting from perfect Bayesian equilibrium.

Definition 1. *A* perfect Bayesian equilibrium (PBE) *of the game induced under the protocol* $\wp$ *is a profile of (pure) strategies and beliefs,*

$$\left(\nu_i^*(.),\ \nu_j^*(.),\ d^*(.,.);\ \mu_i^\wp,\ \mu_j^\wp,\ \mu_D^\wp, \mu_O^\wp\right),$$

for all histories such that the strategies are sequentially rational given beliefs, and the beliefs are derived applying Bayes' rule wherever possible.

Recall that $E_i^\wp$ and $E_j^\wp$ ($E_i^\wp = E_j^\wp$ in the case of secrecy) are the expectations over expert abilities t_i, t_j, respectively, that the outsider O estimates.[9] Note that $E_i^\wp$ and $E_j^\wp$ also define the experts' terminal payoffs in the game. These expectations are derived through $\mu_O^\wp$, which in turn is a part of equilibrium. Thus, our Bayesian game is quite different from standard

[9]Expected abilities of the first and second movers are the same under secrecy because the order of moves remains hidden.

games where players' payoffs in the terminal nodes are taken as given (or fixed), rather than endogenously determined in equilibrium.

Communication games always have equilibria, called babbling equilibria, under which recommendations are ignored. We will restrict our attention primarily to equilibria under which the expert moving first does not babble.

Definition 2. *An equilibrium of the sequential advice game, in short* DE*, induced under the protocol* $\wp$ *is a* PBE of $\wp$ *where the first expert does not babble.*

Under signal revealing strategies (i.e., an expert recommends A (B) when her signal is α (β)), we say that the expert reports truthfully. We now state a stronger version of DE.

Definition 3. *A strong sequential advice equilibrium, in short* SDE*, of the game induced under the protocol* $\wp$ *is a* DE of $\wp$ *where the first expert reports truthfully.*

D's problem is to choose an optimal protocol to maximize his ex-ante expected payoff from eventual decision making. Given the multiplicity of equilibria in communication games, we take a mechanism design approach.

Definition 4. D *chooses a protocol* $\wp$ *over* $\wp'$ *if for every profile of parameters* $(q, \theta, \xi, \lambda)$*, there exists some* PBE *in the game induced by* $\wp$ *under which the payoff of* D *is (weakly) greater than his payoff under all* PBE *in the game induced by* $\wp'$*, and strictly greater for some parameters.*

This ends the description of the decision maker's two-step decision problem.[10]

The following assumption will be maintained throughout the paper.[11]

Assumption 1. $1/2 < q < \xi < \lambda < 1$.

That is, even the low-ability expert's signal is more informative than the unrefined (prior) information. Assumption 1 along with the definition of k implies the following fact:

Fact 1. $\frac{\lambda}{1-\lambda} > \frac{k}{1-k} > \frac{\xi}{1-\xi} > \frac{q}{1-q} > 1$.

[10]The reader may be concerned that in the protocol chosen there may be an equilibrium which is worse than all equilibria in the other protocol and strictly worse under some parameter values. This won't happen in our environment because the worst equilibrium in both protocols is the babbling equilibrium.

[11]The assumption helps reducing the set of equilibria. Our qualitative results do not change if $q \geq \xi$.

3 Bias, beliefs, and partitioning k and q

Because the prior bias q and the lead opinion bias k will influence experts' strategies through various beliefs, this section develops some preliminaries to partition these variables.

We say that signal α (β) favors a corresponding state, if conditional on signal (and possibly other observables) the expert assigns to state a (b) a probability greater than $\frac{1}{2}$. In Appendix A we show that, both signals of the first expert, i.e. expert i, favor their respective states. This is due to Fact 1.

Now consider expert j who moves second. Under an *SDE,* she would deduce the first expert's signal, but not her ability, from the observed recommendation. If the second expert's signal *matches* that of the first, then the signal favors its corresponding state (see (A.2)).

When the first expert's signal is β and the second expert observes α, then the signal favors the corresponding state when the expert is of high ability λ. This "*non-herding*" belief occurs for two reasons: (i) prior favors a ($q > \frac{1}{2}$), and (ii) average precision k of the first expert is less than the high precision of the second expert (see (A.3)).

When the first expert's signal is deduced as α and the second expert receives β, the latter's signal *does not* favor the corresponding state when the expert is of low ability ξ (see (A.4)). This is because the prior bias favors a and the first expert's average precision exceeds that of the second expert. Thus, low-ability expert's signal not prevailing over the prior bias arises endogenously.

So far, the relative weights of q and k have not mattered. However, they *will* matter in two cases. In case (i), the first expert's signal (s_i) is deduced to be β, the second expert is of ability $t_j = \xi$ with signal $s_j = \alpha$. Then, conditional on (s_i, s_j, t_j),

$$\Pr(a \mid \beta, \alpha, \xi) = \frac{q\xi(1-k)}{q\xi(1-k)+(1-q)(1-\xi)k} \geq \frac{1}{2} \quad \text{if and only if} \quad \frac{\xi}{1-\xi} \geq \frac{k}{1-k}\frac{1-q}{q}. \tag{5}$$

In case (ii) we have $(s_i, s_j, t_j) = (\alpha, \beta, \lambda)$, and

$$\Pr(b \mid \alpha, \beta, \lambda) = \frac{(1-q)\lambda(1-k)}{q(1-\lambda)k+(1-q)\lambda(1-k)} \geq \frac{1}{2} \quad \text{if and only if} \quad \frac{\lambda}{1-\lambda} \geq \frac{k}{1-k}\frac{q}{1-q}. \tag{6}$$

To see when signals favor their respective states, or when the inequalities $\frac{\xi}{1-\xi} \geq \frac{k}{1-k}\frac{1-q}{q}$ and $\frac{\lambda}{1-\lambda} \geq \frac{k}{1-k}\frac{q}{1-q}$ are satisfied, let $k(\xi)$ and $k(\lambda)$ be the values of k when the first and second inequalities bind. Define:

$$k(\xi) = \frac{q\xi}{q\xi + (1-q)(1-\xi)}, \qquad k(\lambda) = \frac{(1-q)\lambda}{(1-q)\lambda + q(1-\lambda)}. \tag{7}$$

Given $q > \frac{1}{2}$ it is easy to see that $\xi < k(\xi)$ and $k(\lambda) < \lambda$, but $k(\lambda)$ may exceed or be less than $k(\xi)$, as listed below:[12]

$$k(\lambda) \leq \xi < \lambda \leq k(\xi), \tag{8}$$

$$\xi \leq k(\lambda) \leq k(\xi) < \lambda, \tag{9}$$

$$\xi < k(\xi) < k(\lambda) < \lambda. \tag{10}$$

These partitions will be important in deriving equilibria under secrecy and comparing across protocols.

Let

$$r \equiv (\frac{\lambda}{1-\lambda})/(\frac{\xi}{1-\xi}) > 1$$

be the "(relative) merit of the high ability expert". We now fix the following two definitions to be repeatedly referred in the rest of this paper:

Definition 5. *The* ***lead opinion bias****, k, is* <u>*small*</u> *or* <u>*large*</u> *if, respectively,*

$$k \leq k(\xi), \quad \textit{or} \quad k > k(\xi).$$

Definition 6. *The* ***prior bias****, q, will be called* <u>*small*</u>*,* <u>*medium*</u> *or* <u>*large*</u> *if, respectively,*

$$\frac{q}{1-q} < \sqrt{r}, \qquad \sqrt{r} \leq \frac{q}{1-q} < r, \qquad \textit{or} \qquad \frac{q}{1-q} \geq r.$$

These ranges are specified in terms of the prior odds ratio for ease of comparison with the ability index. An important point to note here is that as q increases, so does $k(\xi)$. This means, a strong prior bias is likely to weaken the importance of the lead opinion bias with k switching from being large to small.

The following lemma is easy to verify using Fig. 5 in the Appendix. Fig. 5 will be referred in the other proofs as well.

Lemma 1.

(i) *[**Panel 1**] Inequalities* (8) *will hold if and only if the prior bias is* <u>*large*</u>*;*

(ii) *[**Panel 2**] Inequalities* (9) *hold if and only if the prior bias is* <u>*medium*</u>*;*

(iii) *[**Panel 3**] Inequalities* (10) *hold if and only if the prior bias is* <u>*small*</u>*.*

[12]It is easy to rule out $k(\lambda) \leq \xi < k(\xi) < \lambda$. Panels 1-3 in Fig. 5 in the Appendix display these orderings.

4 Transparency: Revelation hurdles

D's objective is to maximize the probability that the decision chosen corresponds with the state. An expert's payoff derives from O's beliefs about the expert's ability. Under transparency, O gets to see not only the realized state but also who recommended what and when. Since O is only interested in the expert's ability, this information is sufficient for forming beliefs; D's decision becomes redundant.

We highlight the following two strategies from the strategy set $\mathbf{V}_e$ in (1). The complete list is included in the Appendix.

Truthful recommendation. An expert recommends according to her signal:

$$\mathbf{V}_e^s = \{(A, A, B, B)\}.$$

Babbling. An expert babbles if her recommendation is completely uninformative:

$$\mathbf{V}_e^b = \{(A, A, A, A), (B, B, B, B)\}.$$

With only two possible signals, the contrarian strategy (B, B, A, A) is equivalent to truthful recommendation.[13] Hence, without loss of generality, we drop this recommendation profile from the strategy sets. We can now start presenting the formal results.

Consider first the continuation equilibria involving the second expert. We first rule out certain strategies in equilibrium. The statements below in reference to the eliminated strategies may appear cryptic but the details are provided in the Appendix.

Lemma 2. *If $v_j(v_i)$ is an equilibrium strategy in the continuation game following a recommendation of v_i, then $v_j(v_i) \notin \mathbf{V}_j \backslash \{\mathbf{V}_j^s \cup \mathbf{V}_j^b\}$.*

Among the strategies eliminated in Lemma 2 are the following categories: one or both abilities get revealed, or experts of one ability recommend their signals while the other ability recommend contrary to signals. In the case of ability revelation, if disclosed ability is high then clearly the low-ability expert can mimic the high ability's strategy to make such disclosure impossible; if disclosed ability is low then such an expert would deviate to choosing a strategy so that such disclosure does not happen. Ruling out selectively contrarian recommendation cannot be done by an easy intuition, so we have to rely on the formal proof (see Case 3).

[13] Sometimes *contrarian* is used differently to mean the second expert recommending contrary to the first expert's recommendation. From the context the meaning should be clear.

Lemma 3. $\nu_j(\nu_i) \in \mathbf{V}_j^s = \{(A, A, B, B)\}$ *is a continuation equilibrium strategy following a recommendation of* ν_i *if and only if: the first expert babbles or*

(i) *the first expert observes* β *and* $\nu_i = B$*; and*

(ii) k *is small, i.e.,* $k \in [\xi, k(\xi)]$.

For a large prior bias, condition (ii) of Lemma 3 is always satisfied (see (i) of Lemma 1).

Equilibria, under transparency, are influenced by the lead opinion bias and solely driven by "learning". If the first expert babbles, the second expert learns nothing, so she recommends her signal. If the first expert were to reveal her signal and recommend A, the importance of the second expert's signal is washed out if she is of low ability with signal β. This in turn makes the low-ability second expert herd, making truthful recommendation impossible. Given Lemma 2, then the only equilibrium strategy for the second expert is to babble. Similarly, if the first expert were to reveal her signal and recommend B, then if the lead opinion bias, k, is large, the second expert of ability ξ with signal α would herd with the first expert. Given the impossibility of truthful recommendation, Lemma 2 then implies the second expert would babble. But if k is small, a signal of α would still be informative. Then, the second expert recommending truthfully would be an equilibrium strategy.

Lemma 4. $\nu_j(\nu_i) \in \mathbf{V}_j^b$ *is a continuation equilibrium strategy for all* $\nu_i \in \{A, B\}$.

That is, second expert would always babble if she is expected to behave that way.

Finally, consider the first expert's strategy.

Lemma 5.

(i) *There exists a* PBE *where the first expert truthfully recommends her signal.*

(ii) *There is no* PBE *with* $\nu_i \in \mathbf{V}_i \backslash \{\mathbf{V}_i^s \cup \mathbf{V}_i^b\}$.

(iii) *There is always a* PBE *in which expert* i*, who moves first, babbles.*

For part (ii), similar elimination arguments as in Lemma 2 are applicable. Thus there are *only* two kinds of equilibria involving the first expert, one where she truthfully recommends her signal and the other where she babbles. Under truthful recommendation, O's expectation about the expert's ability is greater if the recommendation were to match the state. Since signals are informative about the state for both precisions, the first expert reveals her signal.

Lemmas 2–5 characterize the set of equilibrium strategies for the experts. They imply that the following types of equilibria exist: (i) both experts babble; (ii) first expert babbles and second expert recommends truthfully; (iii) first expert recommends truthfully and second expert always babbles; and (iv) first expert recommends truthfully and second expert recommends truthfully only if the first recommendation is B. Equilibria (iii) and (iv) are SDE and will be of special interest to us. We state them explicitly in Proposition 1.

Combining Lemmas 2–5 we have:

Proposition 1 (Transparency: Equilibrium characterization)**.** *Under transparency an* SDE *equilibrium exists. The experts' strategies in* SDE *equilibria are as follows:*

1. *In all* SDE*:*
   - (i) *If* $\nu_i = \mathsf{A}$*, then* $\nu_j(\nu_i) \in \mathbf{V}_j^b$*.*
   - (ii) *If* $\nu_i = \mathsf{B}$ *and* k *is* <u>*small*</u> *(i.e.,* $k \leq k(\xi)$*), then* $\nu_j(\nu_i) \in \mathbf{V}_j^s \cup \mathbf{V}_j^b$*. If* $\nu_i = \mathsf{B}$ *and* k *is* <u>*large*</u> *(i.e.,* $k > k(\xi)$*), then* $\nu_j(\nu_i) \in \mathbf{V}_j^b$*.*
2. *In all* SDE*, following a recommendation of* A *by the first expert,* D *chooses* A*. Following a recommendation of* B*,* D*'s decision will depend on the continuation equilibrium. For the babbling equilibrium,* D *chooses* B*. For the truthful recommendation equilibrium,* $d(\mathsf{B},\mathsf{B}) = \mathsf{B}, \;\; d(\mathsf{B},\mathsf{A}) = \mathsf{A}$*.*

Under transparency there <u>does not exist</u> any equilibrium where both experts, regardless of their abilities, will recommend their signals. Also, the experts' abilities never get revealed. The only chance for the second expert to be pivotal is if the first expert recommends B. Proposition 1 thus indicates the likely limitations of transparency: recommendation of the prior-favorite A by the first expert or a high average ability of the first mover stifles any meaningful advice by the second mover. How it fares against secret advice will be the focus from here onwards.

5 Secrecy and revelation

In this section we present two classes of equilibria under secrecy.

Under transparency only two biases mattered: the prior bias q and the lead opinion bias k. Under secrecy potentially there is also a third bias, the conformity bias, that can arise endogenously due to O's beliefs. Now the experts cannot be judged directly by their recommendations but O will have to make inferences about their collective type based on decision accuracy. Specifically, O could believe that the low-ability type has a greater propensity to make a wrong prediction when the true state is a, the one favored by the prior, than when the true state is b, the less favored one. This biases the experts towards recommending A. Let x' (y'') be the experts' payoffs when D's decision matches (does not match) state a. Similarly, let x'' (y') be the experts' payoffs when D's decision matches (does not match) state b.[14] We say that there is ***conformity bias*** if

$$x' - y'' > x'' - y'.$$

In the first class of equilibria there will be no conformity bias (i.e., $x' - y'' = x'' - y'$), whereas the second class will exhibit this bias.

■ **Herding and partial type revelation.** We first present an equilibrium (see Fig. 1), where the second expert's advice is shaped *only* by the lead opinion bias and the prior bias. To help understand the figure, recall that an element of $\mathbf{V}_e$, $e \in \{i, j\}$ is a quadruple $(v_e(\alpha, \xi), v_e(\alpha, \lambda), v_e(\beta, \xi), v_e(\beta, \lambda))$. We call this a partial type revealing equilibrium because only the second expert is able to reveal her type and that too in certain cases. For example if the first expert were to recommend B, the second expert's *high* type gets revealed only if she were to recommend A (i.e. if she were to get signal α). A recommendation of B, neither reveals the signal nor ability.

Figure 1: Partial type revelation strategies

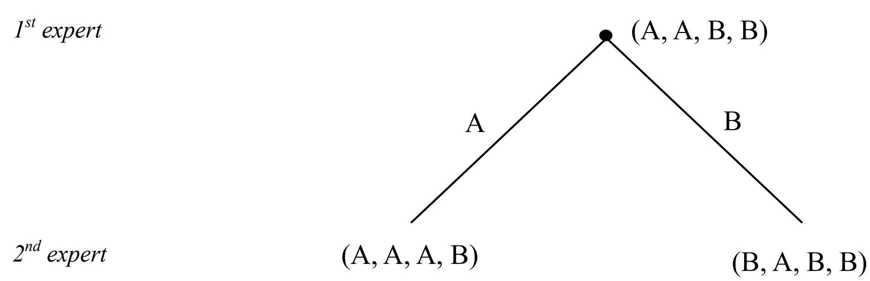

Proposition 2 (Partial type revelation).

(i) *Consider the case of small prior bias* q *as in Lemma 1 (Panel 3 in Fig. 5) and* $k \in [k(\xi), k(\lambda)]$ *(so* k *is large but not "too large"). Then the following recommendation*

[14] These payoffs, identical for both experts, will be derived in the proof of Proposition 3.

strategies can be supported as an SDE *under secrecy:* $\nu_j(A) = (A, A, A, B)$, $\nu_j(B) = (B, A, B, B)$.

(ii) *For the particular equilibrium in* (i), *conformity bias does not arise.*

In this equilibrium, the second expert's high ability, λ, is revealed when she recommends contrary to the first expert's recommendation. The beliefs, not specified here, will be derived in the proof.

On the left-hand branch of Fig. 1, the lead opinion bias induces the low-ability second expert who observes signal β to recommend non-truthfully the decision A. This is irrespective of the value of q, as even the worst prior ($q \approx 1/2$) agrees with the lead opinion (see (A.4)). Contrast this with a high-ability expert who gets signal β and sees a first recommendation of A. She knows that her ability is higher than the average ability of the first expert. This by itself is not sufficient for her to totally disregard the first recommendation. The weight she puts on the information content of the first expert depends on the prior bias q. When q is small (as in Panel 3), she puts relatively less weight on the information content and goes with her own signal β to recommend B; thus even a large lead opinion bias fails to prevail when the prior bias is weak. So for low enough q, ***failure*** *of full signal revelation* ***enables*** *partial type revelation.*

On the right-hand branch, given a small prior bias q, the low-ability second expert who observes signal α gets sufficiently influenced by a large lead opinion bias. For her, the expected type of the first expert is reasonably high (k above $k(\xi)$). Unable to ignore the first expert's recommendation of B, she then herds. For the high-ability expert who gets signal α, the expected precision of the first expert is not "too large" (k below $k(\lambda)$) relative to her own precision. Thus she recommends her signal. Therefore, again we see that a small prior bias q leads to partial type revelation.

Under transparency any separation between the high and low (second) expert types is impossible since otherwise the low type would deviate to mimic the high type's strategy.

Proposition 2 is quite significant in relation to the previous works on transparency of expert advice. The result is new and very different from what we know, say, from the one-shot voting model of Levy (2007a). In her framework, partial type revelation is *not* possible. Moreover, as can be seen from the proof, here partial type revelation happens with complete absence of the conformity bias, a bias known to play a prominent role in group decisions.

■ **Signal revelation.** Now consider the second class of equilibria where both experts recommend their signals. Then, for recommendation profiles (A, A), (A, B), (B, A) and

(B, B), D knows that the corresponding signals are (α, α), (α, β), (β, α) and (β, β). D's posteriors are then:

$$\begin{array}{rcll} \Pr(\omega = a \mid A, A) = & \frac{qk^2}{qk^2+(1-q)(1-k)^2} & > & \frac{1}{2}, \\ \Pr(\omega = a \mid A, B) = & q & > & \frac{1}{2}, \\ \Pr(\omega = a \mid B, A) = & q & > & \frac{1}{2}, \\ \Pr(\omega = a \mid B, B) = & \frac{q(1-k)^2}{q(1-k)^2+(1-q)k^2} & < & \frac{1}{2}. \end{array} \tag{11}$$

These probabilities are calculated using Tables A.2 and A.3 in Appendix A.[15]

Lemma 6 (D's decision under signal revelation)**.** *Let the recommendations reveal signals.* D *selects* B *if and only if both experts recommend* B*; otherwise* D *selects* A*.*

Recall that under transparency, the first expert is decisive when she recommends A. There, following a recommendation of A, the second expert babbled. Here, on the other hand, the first expert is decisive when she recommends A despite the second expert revealing her signal. This occurs because of two reasons: (i) the prior bias favors a; and (ii) the second expert does not reveal her ability, even partially.

We now proceed to determine the conditions under which the experts will recommend their signals. Since under secrecy experts are motivated by collective reputation (see, for example, McLennan (1998)), they would communicate their signals truthfully and thereby maximize the chance of a correct decision provided their signals meet a minimal *accuracy threshold*. Signal accuracy will be measured by the ratio of the likelihood of correct signal over that of incorrect signal, $t_i/(1-t_i)$ and $t_j/(1-t_j)$. The accuracy threshold will be shown to be linked to two bounds,

$$c(k) = \frac{1-k}{k}\frac{q}{1-q}\frac{1+k}{2-k}, \qquad c^{-1}(k) = \frac{k}{1-k}\frac{1-q}{q}\frac{2-k}{1+k}.$$

Here $\frac{k}{1-k}$ reflects lead opinion bias, $\frac{1+k}{2-k} > 1$ is to be shown as the scaling up factor for conformity bias, and $\frac{q}{1-q}$ is the prior bias. The higher the k, the higher would be both the lead opinion bias (stifling weak expert's revelation) and the conformity bias (stifling recommendation against favorite A). But these two biases appear differently in each bound,

[15] For example consider $\Pr(\omega = a \mid A, B)$. Since experts reveal their signals, $\Pr(\omega = a \mid A, B) = \Pr(\omega = a \mid \alpha, \beta)$. This probability is a ratio where the numerator is the sum of all column entries in row two of Table A.2. In the denominator we have this sum plus the sum of all column entries in row two of Table A.3. This ratio is q. The other probabilities can be derived similarly.

with one in the numerator and the other in the denominator. The overall magnitude of $c(k)$, decreasing in k, will place a lower bound on signal accuracy so that a β-signal expert with accuracy exceeding this bound will tell the truth. $c^{-1}(k)$, on the other hand, is increasing in k and places a lower bound on signal accuracy so that an α-signal expert with accuracy exceeding the bound will tell the truth. With

$$c(1) = 0, \ \lim_{k \to 0} c(k) = \infty, \text{ and } \ c^{-1}(0) = 0, \ \lim_{k \to 1} c^{-1}(k) = \infty,$$

we can thus define $\underline{k}(\xi)$ and $\bar{k}(\xi)$ such that

$$c(\underline{k}(\xi)) = \xi/(1-\xi) = c^{-1}(\bar{k}(\xi)),$$

where $\underline{k}(\xi) < \bar{k}(\xi)$, and

$$\lambda/(1-\lambda) > \xi/(1-\xi) \geq \max\{c(k), c^{-1}(k)\} \quad \text{for } \ \underline{k}(\xi) \leq k \leq \bar{k}(\xi).$$

Thus, for k in the range $[\underline{k}(\xi), \bar{k}(\xi)]$, the signal accuracy will always lie above the threshold $\max\{c(k), c^{-1}(k)\}$, as depicted in Fig. 2. In this range there will be truthful signal revelation by both experts. Full signal revelation will fail if k becomes too small or too large.

Figure 2: Critical k values

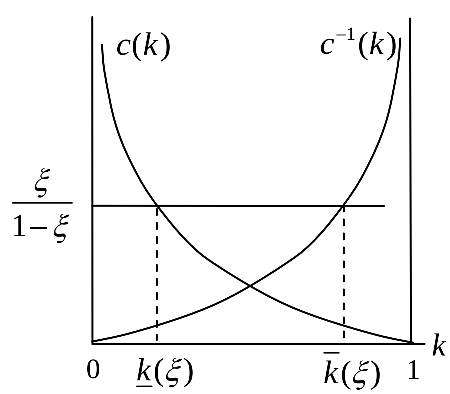

Let us now try to understand, as a first pass, what economics underlies k's influence on signal revelation (or non-revelation). Fixing $q > 1/2$, if k is too small (i.e., below $\underline{k}(\xi)$), reliance on the other expert's average quality signal to induce a correct decision cannot be very high. So for a low-ability (ξ) expert with a β-signal to reveal truthfully and go against prior bias and conformity bias, both favoring decision A, it has to be that the signal accuracy of the low ability, i.e. $\xi/(1-\xi)$, be sufficiently large. But this requirement is likely to fail because the signal accuracy threshold will become large due to $c(k)$ getting large (refer Fig. 2). On the other hand, if k becomes too large (i.e., above $\bar{k}(\xi)$), a low-ability expert with an α-signal becomes averse to communicating informatively as she would rather rely on the other expert's expected high quality signal to induce the right decision; note also, with large k conformity bias will become small and so the incentive to truthfully recommend A is not that high. Both these effects manifest in the relevant bound $c^{-1}(k)$ getting too large. Given that the two bounds $c(k)$ and $c^{-1}(k)$ move in opposite directions with respect to k, only an intermediate range of k will simultaneously satisfy the signal revelation incentives of the β- and α-signal experts. These are rough intuitions. Further explanations follow Proposition 3. The precise mechanism of how the three biases interact in signal revelation can be seen by studying the formal proof of Proposition 3. We need an intermediate lemma.

Lemma 7. $\underline{k}(\xi) < \xi$, *and* $\bar{k}(\xi) > k(\xi)$.

Since $k(\xi) > \xi$, we can now precisely break down k into two sets – revelation and non-revelation/babbling.

Proposition 3 (Signal revelation)**.** *Under secrecy, there exists an* SDE *where the second expert's strategies are as follows: For all* $v_i \in \{A, B\}$,

(i) *if* $\xi \leq k \leq \bar{k}(\xi)$, *then* $v_j(v_i) \in \mathbf{V}_j^s$; **(Revelation)**

(ii) *if* $k > \bar{k}(\xi)$, *then* $v_j(v_i) \in \mathbf{V}_j^b$. **(Babbling)**

The first expert, by definition of SDE, *will recommend truthfully according to her signal.*

Because $\bar{k}(\xi) > k(\xi)$ by Lemma 7, with the help of Fig. 5 in the Appendix the following result is immediate from Proposition 3:

Corollary 1. *If the bias* q *is large as in Lemma 1, i.e.* $\frac{q}{1-q} \geq r$, *then under secrecy there is always a full signal revelation equilibrium.*

Recall from Proposition 1 that under transparency, in all SDE, the second expert babbles following first expert's recommendation of A. Proposition 3 tells us that under secrecy, there exists an SDE for which this is not so. Under this SDE, when the prior bias is large, both types always recommend their signals. This is also true when the prior bias is medium or small, provided that $\xi \leq k \leq \bar{k}(\xi)$.

In the equilibrium stated in Proposition 3, conformity bias comes into play. Recall that x' (y'') are the experts' payoffs when D's decision matches (does not match) state a and x'' (y') be the experts' payoffs when D's decision matches (does not match) state b. An expert's expected payoff from recommending A will be in proportion to $(x' - y'')$, whereas her expected payoff from recommending B is in proportion to $(x'' - y')$, besides the parameters q, an expert's own type t_e and the other expert's average type k influencing the payoffs. Thus, the relative magnitudes $(x' - y'')$ vs. $(x'' - y')$ will be a key determinant of an expert's recommendation strategy. In the signal revelation equilibrium it follows from derivations in the Appendix (see the proof of Proposition 3) that $[x' - y''] = \frac{1+k}{2-k}[x'' - y']$ i.e. $(x' - y'') > (x'' - y')$, which is the conformity bias effect that makes experts predominantly biased towards recommending the prior-favored alternative A. Intuitively, the conformity bias arises because D's decision of A proving wrong leaves open the possibility in the minds of O that any of two experts privately could have made the right recommendation which is B, whereas a decision of B proving wrong allows no such positive accounting as both must have recommended B.

We can now understand the results which follow when the prior bias is not large. Consider the case where the first expert recommends B (so that the second expert is pivotal).[16] If the second expert who is of low ability observes α, her posterior on state a may go below $\frac{1}{2}$.[17] However, even though her posterior on state a is low, due to conformity bias her expected *payoff* from recommending A may be greater than the payoff from recommending B. In such situations, by recommending honestly and triggering A, she avails the insurance of conformity bias. On the other hand if she observes β, whether she is of low or high ability, she is emboldened by the first expert's opinion (lead opinion bias) and therefore recommends B despite the conformity bias. The safety net of the conformity bias when in possession of α signal and the support of the lead opinion bias in case of β signal are the two contributory factors for signal revelation under secrecy, that are not possible under the full glare of transparency.

What is rather surprising about Proposition 3, however, is that for signal revelation to happen k must not exceed $\bar{k}(\xi)$: having "too good" experts spoils it! Notice that

[16] After the first expert recommends A, the second expert's report is inconsequential for the experts' payoffs.

[17] If $\phi(k)\frac{1-q}{q}\frac{2-k}{1+k} \leq \frac{\xi}{1-\xi}$ so that $k \leq \bar{k}(\xi)$, it is possible that $\phi(k)\frac{1-q}{q} > \frac{\xi}{1-\xi}$; see (5) for the posterior. This may happen despite prior favoring a.

conformity bias again comes into play but this time it is washed out by the lead opinion bias. To see this, consider the case when k is large. As k increases, the conformity bias effect (as previously observed in the relation $[x' - y''] = \frac{1+k}{2-k}[x'' - y'])$ increases. But for k above $\bar{k}(\xi)$ the conformity bias is swamped by the lead opinion bias effect. So following a first expert recommendation of B, the second expert who observes signal α would recommend B too. The increase in k makes her updated beliefs in favor of state a very low. Conformity bias is no longer sufficient for her to recommend A. The following observation is already implicit in Proposition 3:

Corollary 2. *If the experts are "too good," the decision maker may as well do away with the second expert even if the expert is costless.*

The above corollary is meant only in reference to signal revelation.

When the prior bias q is small, by bringing together Propositions 2 and 3 one can see that there is an overlap of partial type revelation with full signal revelation in the region of k: $[k(\xi), k(\lambda)] \cap [\xi, \bar{k}(\xi)] \neq \emptyset$. Hence, under secrecy there are multiple SDE.[18] A natural question then is how do these equilibria rank? In partial type revelation equilibrium there is some amount of herding, whereas in full signal revelation D loses out on the important information that the second expert could be of high ability. An answer is provided in the next section.

6 Transparency or secrecy?

To evaluate the relative merits of the two protocols, it is sufficient to consider only one class of equilibria under secrecy, i.e., the one in Proposition 3. Consideration of type revelation equilibrium would only strengthen the main finding of this comparison.

Given that $\bar{k}(\xi) > k(\xi)$ we conclude, from Proposition 1 and Proposition 3, that signal revelation occurs under secrecy over a larger parameter space. The main difference between the two mechanisms comes in the form of revelation incentives of a low (ξ) ability expert who observes the signal α, when preceded by a B-recommendation by the first expert.[19] Under

[18]Multiplicity of equilibria is not surprising as there is always a babbling equilibrium.

[19]As we already know, under transparency if the first expert recommends A the only continuation equilibrium possible is that of babbling by the second expert. And under secrecy, the first expert recommending A makes the second expert's recommendation inconsequential for D's decision.

transparency, signal revelation requires:

$$\begin{aligned}\Pi_j^t(B,A,\alpha,\xi) = \Pr(a \mid \beta,\alpha,\xi)\gamma &+ \Pr(b \mid \beta,\alpha,\xi)\gamma' \\ &\geq \Pr(a \mid \beta,\alpha,\xi)\gamma' + \Pr(b \mid \beta,\alpha,\xi)\gamma = \Pi_j^t(B,B,\alpha,\xi) \\ \text{or,}\quad \Pr(a \mid \beta,\alpha,\xi)\big[\gamma-\gamma'\big] &\geq \Pr(b \mid \beta,\alpha,\xi)\big[\gamma-\gamma'\big],\end{aligned}$$

where $\gamma = E_j^t(B,A,a) = E_j^t(B,B,b)$, $\gamma' = E_j^t(B,A,b) = E_j^t(B,B,a)$ (see Lemma 3 proof). That is, the second expert will recommend truthfully if her chance of being right is higher than that of being wrong, i.e., $\Pr(a \mid \beta,\alpha,\xi) \geq 1/2$. The term $[\gamma-\gamma']$, which is the gain in one's perceived ability from being an accurate predictor over being inaccurate, is the <u>same</u> whether the revealed state is a or b and thus drops out. In our terminology, this means that there is no conformity bias. The experts are now evaluated by the accuracy of their individual recommendations, so there is no collective blame or benefit-of-doubt, i.e., no conformity bias.

In contrast, under secrecy a similar payoff comparison makes truthful recommendation optimal if (see Proposition 3 proof):

$$\Pr(a|s_i=\beta,s_j=\alpha,\xi)\big[x'-y''\big] \geq \Pr(b|s_i=\beta,s_j=\alpha,\xi)\big[x''-y'\big].$$

Here due to conformity bias (i.e., $\big[x'-y''\big] > \big[x''-y'\big]$) even if the state a is less likely, the low-ability expert j chooses to recommend decision A as she gets a higher payoff. This expands the truthful recommendation range of k beyond $k(\xi)$ to $\bar{k}(\xi)$.

Hence, restricting ourselves to only the class of secrecy equilibria in Proposition 3 which is about signal revelation, we have the following result.

Proposition 4 (Choice of protocol). *For* $k \in [\xi, k(\xi)]$ *and* $k \in (\bar{k}(\xi), \lambda]$*, transparency and secrecy provide the same payoff to* D*. For* $k \in (k(\xi), \bar{k}(\xi)]$*,* D *is strictly better off under secrecy. Thus by Definition 4,* D *prefers secrecy over transparency.*

■ **Summary of comparison**. From Definition 6 the prior bias q is *small*, *medium* or *large* if, respectively, $\frac{q}{1-q} < \sqrt{r}$, $\sqrt{r} \leq \frac{q}{1-q} < r$ and $\frac{q}{1-q} \geq r$. Below we provide a summary of our comparison of D's payoffs in terms of q and k. Table 1 collates this summary.

If $k \in [\xi, k(\xi))$, D's maximum equilibrium payoff under transparency and secrecy are the same (Proposition 4), so D is indifferent between the two protocols: there is only limited or full signal revelation. This is true whatever be the prior bias q.

When $k \in (k(\xi), \lambda]$, the second expert is redundant under transparency as she babbles (Proposition 1). Let us now look at a further breakdown of q.

Consider *medium* values of q. When $k \in (k(\xi), \bar{k}(\xi)]$, the second expert reveals her signal under secrecy (Proposition 3). This leads to D strictly preferring secrecy over transparency (Proposition 4). Further, for medium values of q, $k(\xi) \geq k(\lambda)$ (see (9) and Lemma 1), so under secrecy the partial type revelation equilibrium does not exist ((i) of Proposition 2).

Consider *small* values of q. For $k \in [k(\xi), k(\lambda)]$, the second expert partially reveals her type under secrecy (Proposition 2). As a result, D's maximum equilibrium payoff under secrecy is strictly larger than that under transparency: follows from Proposition 4 for signal revelation equilibrium; for partial type revelation equilibrium, the dominance can be easily verified (see Proposition 5 below). For $k \in (\xi, \bar{k}(\xi)]$, we also have the signal revelation equilibrium under secrecy.

Table 1: D's payoffs – secrecy vs. transparency

Prior\lead bias	$[\xi, k(\xi))$	$(k(\xi), \bar{k}(\xi)]$	$[k(\xi), k(\lambda)]$	$(\max\{\bar{k}(\xi), k(\lambda)\}, \lambda]$
small: $\frac{q}{1-q} < \sqrt{r}$	secrecy$\equiv$ transp; more signal reveln under secrecy, same decisions	secrecy$\succ$ transp; transp completely stifles 2nd expert above $k(\xi)$	secrecy$\succ$ transp; type reveln under secrecy	secrecy$\equiv$ transp;[a] 2nd expert always babbles under both protocols
medium: $\sqrt{r} \leq \frac{q}{1-q} < r$	same as above[b]	same as above	empty[c]	same as above
large: $\frac{q}{1-q} \geq r$	same as above	empty	empty	empty

[a] Whether some superior secrecy equilibrium exists we cannot ascertain.
[b] '$\equiv$' extends to $k(\xi)$.
[c] Cutoffs are endogenous and vary with q. As a result, some of the k-ranges are empty. See Fig. 5.

The above summary distills into another important observation on the value of secrecy:

Corollary 3. *Let* $z = \max\{k(\lambda), \bar{k}(\xi)\} < \lambda$. *Suppose the prior bias* q *is small, and the lead opinion bias (i.e., average quality of the first expert) is large but not too large:* $k(\xi) < k \leq z$. *Then the second expert can potentially have an impact on the decision only under the secrecy protocol.*

In the context of transparent (sequential) debates with heterogenous experts of *known* abilities, Ottaviani and Sorensen (2001) already pointed out why having too good a first expert might render the second expert's opinion meaningless. Our above observation goes well beyond Ottaviani-Sorensen, in answering the broader question of transparency vs. secrecy when the experts' abilities are private information. Also in contrast to Ottaviani-Sorensen's experts, our second expert is of the same expected quality as the first expert.

■ **Equilibria within secrecy**. We now turn to the question of ranking of possible equilibria within secrecy, posed at the end of Section 5.

Proposition 5 (Value of information: signal vs. type)**.** *Consider the secret advice protocol and let* q *be small, i.e.* $\frac{q}{1-q} < \sqrt{r}$.

(i) *For the parameter range* $\{k : k(\xi) \leq k \leq k(\lambda)\} \cap \{k : \xi < k \leq \bar{k}(\xi)\}$*,* D*'s payoff from the partial type revelation equilibrium strictly dominates the payoff from the full signal revelation equilibrium.*

(ii) *If* $\big(k(\xi), k(\lambda)\big) \cap \big(\bar{k}(\xi), \lambda\big] \neq \emptyset$*, then in this range the payoff under partial type revelation equilibrium strictly dominates the payoff when only the first expert reveals her signal.*

The reason for the first payoff dominance can be understood as follows. On the left-hand branch of Fig. 1 the decisions differ only for the recommendation sequence (A, B), with the final decision being A in the signal revealing equilibrium whereas the decision is B in the partial type revealing equilibrium. When the decisions differ, D not only learns the true signal of the second expert he also learns that it is coming from a high-ability expert; this lifts up partial type revelation for D. On the right-hand branch, decisions differ again in only one scenario: when low-ability second expert observes signal α she herds under partial type revelation, triggering decision B, whereas under full signal revelation she would have triggered decision A; given that in signal revealing equilibrium A is triggered by the low-ability second expert's α signal against the first expert's average quality (k) β signal, decisions will be poorer on average. Thus, overall, partial type revelation equilibrium yields higher expected payoff for D. Intuition for the second dominance is straightforward.

7 Deliberation and detailed advice

Under secrecy, we saw that partial type revelation was possible. But can *both* experts reveal their entire type (signal and precision level)? We show that this is possible if experts are allowed to deliberate through back-and-forth messages or simply engage in only a two-round dialogue (same as in the sequential advice model) but with more messages. We do not do any detailed analysis for transparency, given what we already know from Section 4 and the negative results (on information revelation) of Ottaviani and Sorensen (2001; 2006a,b,c), Meade and Stasavage (2008), and Fehrler and Hughes (2018).[20] We will informally argue, however, how making deliberation transparent can stymie the later stages of deliberation.

As a protocol, deliberation enriches our analysis of the secret sequential advice game. While some earlier works (cited in the Introduction) do model deliberations, to our knowledge, there is no general analysis of deliberation as an interactive, back-and-forth communication game (as opposed to simultaneous exchange of views/signals) readily applicable in decision making contexts. Aumann and Hart's (2003) long-drawn message game between two players (one informed and another uninformed), who have to take an action each in an uncertain bimatrix game at the end of the *talk* (or message) phase, is closest to modelling protracted cheap-talk communications. But both players in their model have intrinsic stakes in the *outcomes* of the game to be played. Directly related to organizational economics, there is a recognition that diverging opinions could in fact be beneficial for efficient decision making. See, for example, Landier, Sraer and Thesmar (2009) who consider how separation of an implementer and a proposer of projects, with the two parties' differing biases, can be a positive influence in decision making. In our secret sequential advice protocol, making an effective use of the second expert's differing views (about the suitable action) is in the same spirit of promoting diverse opinions for optimal decisions. But still the sequential advice mechanism fails to draw out any difference of opinions in a full-fledged manner. The deliberation game to be considered in this section examines the protocol's information aggregation potential via emerging conflicting/corroborating views, when the experts do not have any intrinsic interests or biases in decisions. Our main message dispels any negative view of deliberation necessarily compromising decisive actions by creating seeds of doubts, so long as deliberation is done secretly.

Formally, consider a communication format under secrecy where there are four stages. In each stage an expert recommends an action, i.e., an element from $\{\mathsf{A}, \mathsf{B}\}$. The experts

[20] Fehrler and Hughes study a committee of experts deliberating simultaneously and then voting simultaneously. They also establish that transparency hurts information aggregation.

and D observe all recommendations but O does not. A recommendation profile is a four-dimensional vector with coordinates belonging to $\{A, B\}$. Following the recommendation, D chooses $d \in \{A, B\}$. As before, O observes d and ω and forms expectations about the experts' types. These expectations are the payoffs of the experts. We shall be interested in the existence of an equilibrium under which *both* experts reveal their entire two dimensional type. We consider two communication formats: *deliberation* and *detailed recommendation*.[21]

Note that the experts get to send two messages, which we call recommendations. Since messages do not have any content in terms of meaning, we can let experts choose "reports" from $\{\alpha, \beta\} \times \{\xi, \lambda\}$ instead of "recommendations" from $\{A, B\} \times \{A, B\}$. Thus, unlike in previous sections, each expert now has access to messages with the richness to convey the entire content of information.

Under deliberation, in stages one and three, expert i reports and in stages two and four, expert j reports. The first report is about the signal and the second report is about the signal's precision. To literally accommodate the idea of back-and-forth deliberation, one can alternatively consider experts reinforcing their initial recommendation of an action by repeating later the same recommendation, in an attempt to signify high ability, or flipping the original recommendation to suggest that the initial recommendation stands but with a weaker force, which is effectively an admission of low ability. Thus, the recommendation at the second opportunity is a communication about the initial recommendation's precision.

Under detailed recommendation, in stage one expert i makes a recommendation together with an indication of its precision (i.e., first two stages rolled into one), followed by a similar communication in stage two (last two stages combined) by expert j.

The difference between the two protocols is that the first protocol allows the expert moving in stage 1 to choose her stage 3 report at the backdrop of the other expert's partial report in stage 2, whereas in the second protocol the first mover makes all her submission before hearing anything from the other expert.

We consider only truth-telling equilibria, so D can correctly deduce the signal and precision level of both the experts from the communication. Call such an equilibrium a *fully revealing equilibrium.*

[21] For the latter, the communication effectively reduces to only two stages.

Table 2: Deliberation/detailed recommendation

	Report Profile	Types Deduced by D	d	q large	q not large
1.	$(\alpha, \alpha, \lambda, \lambda)$	$\{(\alpha, \lambda), (\alpha, \lambda)\}$	A		
2.	$(\alpha, \alpha, \lambda, \xi)$	$\{(\alpha, \lambda), (\alpha, \xi)\}$	A		
3.	$(\alpha, \alpha, \xi, \lambda)$	$\{(\alpha, \xi), (\alpha, \lambda)\}$	A		
4.	$(\alpha, \alpha, \xi, \xi)$	$\{(\alpha, \xi), (\alpha, \xi)\}$	A		
5.	$(\alpha, \beta, \lambda, \xi)$	$\{(\alpha, \lambda), (\beta, \xi)\}$	A		
6.	$(\alpha, \beta, \lambda, \lambda)$	$\{(\alpha, \lambda), (\beta, \lambda)\}$	A		
7.	$(\alpha, \beta, \xi, \xi)$	$\{(\alpha, \xi), (\beta, \xi)\}$	A		
8.	$(\alpha, \beta, \xi, \lambda)$	$\{(\alpha, \xi), (\beta, \lambda)\}$		A	B
9.	$(\beta, \alpha, \xi, \lambda)$	$\{(\beta, \xi), (\alpha, \lambda)\}$	A		
10.	$(\beta, \alpha, \xi, \xi)$	$\{(\beta, \xi), (\alpha, \xi)\}$	A		
11.	$(\beta, \alpha, \lambda, \lambda)$	$\{(\beta, \lambda), (\alpha, \lambda)\}$	A		
12.	$(\beta, \alpha, \lambda, \xi)$	$\{(\beta, \lambda), (\alpha, \xi)\}$		A	B
13.	(β, β, ξ, ξ)	$\{(\beta, \xi), (\beta, \xi)\}$	B		
14.	$(\beta, \beta, \xi, \lambda)$	$\{(\beta, \xi), (\beta, \lambda)\}$	B		
15.	$(\beta, \beta, \lambda, \xi)$	$\{(\beta, \lambda), (\beta, \xi)\}$	B		
16.	$(\beta, \beta, \lambda, \lambda)$	$\{(\beta, \lambda), (\beta, \lambda)\}$	B		

Conditional on the posterior formed on the basis of a profile of reports, D makes the choice d. In the case of deliberation, Table 2 lists the report profiles under a fully revealing equilibrium (assuming that it exists). In the case of detailed recommendation, since strategies reveal types, one should ignore the "Report Profile" column and simply view the next column as the deduced profile of types. In both formats of communication, in the continuation game following recommendations (or reports), the decisions of D and beliefs of O would be the same in equilibrium. The optimal d is derived from "Types Deduced by D". Since this is a straightforward derivation, we simply state the optimal choice of d in the last two columns.

When q is large D's choice of B depends only on the first report of each expert, i.e., only on the experts' signals of β. Since this is similar to the two-stage case studied earlier, we focus on the case where q is not large. So let $\frac{q}{1-q} < r$.

We now come to O's beliefs. Consider the case where q is not large. Conditional on

(d, ω), O's beliefs are:

$$\begin{aligned}
\Pr(\lambda \mid A, a) &= \frac{\theta\lambda[1+\theta(1-\lambda)]}{\theta\lambda[1+\theta(1-\lambda)]+(1-\theta)[k+(1-\theta)\xi(1-\xi)]}, \\
\Pr(\lambda \mid A, b) &= \frac{\theta(1-\lambda)[1+\theta\lambda)]}{\theta(1-\lambda)[1+\theta\lambda)]+(1-\theta)[(1-k)+(1-\theta)\xi(1-\xi)]}, \\
\Pr(\lambda \mid B, b) &= \frac{\theta\lambda[(1-\theta)+\theta\lambda]}{\theta\lambda[(1-\theta)+\theta\lambda]+(1-\theta)[(1-\theta)\xi^2+\theta\lambda]}, \\
\Pr(\lambda \mid B, a) &= \frac{\theta(1-\lambda)[(1-\theta)+\theta(1-\lambda)]}{\theta(1-\lambda)[(1-\theta)+\theta(1-\lambda)]+(1-\theta)[(1-\theta)(1-\xi)^2+\theta(1-\lambda)]}.
\end{aligned} \tag{12}$$

Define

$$g' = \Pr(\lambda \mid A, a), \quad h'' = \Pr(\lambda \mid B, a), \quad g'' = \Pr(\lambda \mid B, b), \quad h' = \Pr(\lambda \mid A, b). \tag{13}$$

Let

$$Q = \max\left\{ \left(\frac{1-\xi}{\xi} \Big/ \frac{\xi}{1-\xi} \right), \left(\frac{\xi}{1-\xi} \Big/ \frac{\lambda}{1-\lambda} \right) \right\}.$$

The proof of the following Proposition, together with a numerical illustration of the sufficient condition, appears in a separate supplementary file.

Proposition 6 (Deliberation and detailed advice)**.** *Let the recommendation protocol be secrecy, and* $q/(1-q) < r$*. A fully revealing equilibrium exists under deliberation or detailed recommendation if and only if*

$$\frac{q}{1-q}(g' - h'') \geq (g'' - h') \geq \frac{q}{1-q} \cdot Q \cdot (g' - h''). \tag{14}$$

Proposition 6 informs us that for a fully revealing equilibrium to exist we need $g' - h'' \geq 0$ and $g'' - h' \geq 0$ (as $Q < 1$). However, conformity bias, i.e. $g' - h'' \geq g'' - h'$, is not necessary.

The result in Proposition 6, specifically the equivalence between the revelation sets under the two protocols, is quite surprising with no easy intuition. Aumann and Hart (2003) had observed in a two-player, bimatrix game with uncertainty about which of several bimatrix games is being played, allowing informed player to engage in a long cheap talk communication helps reveal his information, provided the revelation is done bit-by-bit over disconnected rounds with some agreements in between about how the eventual game will be played.

Making it "to reveal too much too soon, or to agree to too much too soon" destroys their revelation equilibrium.

In contrast to the above noted contrast, it might seem surprising that detailed recommendation does not improve upon deliberation as far as fully revealing equilibria are concerned. Casual intuition weighs in for deliberation to induce more herding. Under the veil of secrecy, this herding incentive can be negated. Our analysis, however, is only partial as we do not study other possible equilibria for these two communication protocols.

Before we conclude, let us review informally what might happen if deliberation were made transparent.[22] We will make the assumption that the outsider will hold *skeptical belief* about an expert if she is seen to flip her initial recommendation in a later round. Flipping is an indication that the expert must not be very confident of the quality of her signal, hence the expert must be of low ability. This skepticism will simply force expert 1 to repeat in stage 3 the advice she gave in stage 1, and expert 2 to repeat in stage 4 her advice made in stage 2 – a process we call *stymied deliberation*. Applying this logic it is easy to construct equilibria in the transparent deliberation game that are outcome equivalent to the equilibria in Proposition 1. Transparent deliberation will thus be *informationally inefficient*. Skeptical belief makes an expert's advice in subsequent moves completely subservient to her original recommendation, so inefficiency of Proposition 1 equilibrium is inescapable.

8 Simultaneous advice

Transparency or secrecy – the same question can be raised for simultaneous advice. It is reasonable to think that the herd behavior of the second mover causing inefficiency makes transparency worse. So replacing sequential advice with simultaneous advice might restore some parity. On this, below we provide an informal guidance.

Among two types of information, consider first how simultaneous advice impacts on signal revelation. Simple arguments show that transparency is superior to secrecy for $\xi > q > 1/2$, the environment considered in this paper. Under transparency even a low-ability expert would trust her contrarian signal β more to recommend B than recommend the prior-favored

[22] Making detailed advice transparent will make the second expert ineffectual similar to one in Proposition 1.

alternative of A; this would induce truth-telling by the (β, ξ) type, and thus full signal revelation,[23] because the experts are evaluated individually by the accuracy of their advice. Under secrecy, the signal revelation equilibrium of Proposition 3 can be easily replicated as the strategic considerations of the first mover in the sequential communication protocol is same as what an expert would face in the simultaneous advice game. Thus, under secrecy and simultaneous advice signal revelation happens only over $\xi \leq k \leq \bar{k}(\xi)$, whereas under transparency signal revelation obtains over the full range $\xi \leq k \leq \lambda$.

Let us consider next whether under secrecy and simultaneous advice experts' types can be communicated to the decision maker similar to Proposition 2. With the experts making simultaneous recommendations, the opportunity to contradict the other expert's recommendation and thereby conveying one's high-ability is no longer feasible. But still the experts may try to convey the high quality of their advice by selectively contradicting the prior. Whether such a strategy is immune to deviation and, if yes, likely to improve the quality of decision making is not a priori clear. With this mind, we consider the possibility of the following contrarian advice equilibrium:

$$\begin{aligned} & \nu_e(\alpha, \xi) = \nu_e(\alpha, \lambda) = \nu_e(\beta, \xi) = A, \quad \nu_e(\beta, \lambda) = B, \\ \text{and} \qquad & d(AA) = A, \quad d(AB) = d(BA) = d(BB) = B. \end{aligned}$$

The idea is that only high-ability experts with signal β will recommend B to contradict the prior, so seeing such recommendations D can be sure of their high quality while a recommendation of A could come from either high- or low-ability expert.[24] This would then prompt D to make decisions as specified above. In the Supplementary file we rule out such an equilibrium. There it is shown that an (α, ξ)-type would always do better to deviate to recommend B instead. This guarantees by triggering decision B that outsiders think that there is at least one high-ability expert, instead of the inference that both experts could be

[23] The types $(\beta, \lambda), (\alpha, \xi), (\alpha, \lambda)$ should also be truthful.

[24] Obviously, there cannot be an exact parallel with the sequential communication protocol because there is no second mover who can base her advice on first mover's advice. The posited strategy is a natural candidate for a partial type revelation equilibrium. It should be clear that full type revelation is never possible. Of the remaining possibilities, one could think of asymmetric strategies where only one expert adopts the selective contrarian strategy while the other expert always recommends her signal. We do not attempt an exhaustive analysis of all possible equilibria.

of low ability if A were induced.

Thus, for signal revelation simultaneous advice puts transparency ahead of secrecy and when it comes to (partial) type revelation simultaneity of advice may make secrecy less potent than under transparency. So simultaneous advice may change the ordering with transparency possibly dominating secrecy. We do not, however, want to make a definitive claim because a full analysis of secrecy under simultaneous advice requires a separate treatment.[25]

9 Some final remarks

We conclude by discussing some of the assumptions made, especially to emphasize the point that despite its simple looking structure extending the analysis in various directions would involve extraordinary complexities. Also, we point out the limitations and scope of some of the results.

First about the assumptions.

1. The number of experts is limited to two. While tractability is surely a consideration for this restriction, many decisions often involve a small number of experts. Much of the economic insight developed with two experts should remain valid for more than two experts. Under transparency herding is a pervasive force and that should be the case in our model with n experts, just like we have shown for two experts. Analysis of the secrecy game with more than two experts will encounter obvious technical hurdles. However, incentives for truthful recommendation (i.e., signal revelation) due to congruity of interests in building collective reputation remains in place.

2. The restriction to binary messages recommending in favor of one of two decisions while an expert's type includes ability related information is with a concern for modelling how experts might be expected to advise. As discussed in the Introduction, a career oriented expert may be reluctant to declare that the quality of her advice is not high. With expanded messages especially under secrecy, this concern does not have as much bite because experts are trying to impress upon outside evaluators collectively.

[25] Besides, our analysis of secrecy equilibria under sequential advice is not exhaustive.

3. The analysis has been restricted to only pure strategies. Consideration of mixed strategies will, once again, increase the complexity of analysis.

Coming to the results, the following points may be noted:

4. Possible equilibria under secrecy reported in Propositions 2 and 3 are not exhaustive. Despite this, using a plausible criterion of ranking between protocols (Definition 4), we are able to argue that secrecy should be preferred over transparency. We are also able to provide a strict ranking between partial type revelation equilibrium and full signal revelation equilibrium.

5. Section 7 is about how enriching the message space can give rise to full revelation of types and not just signals. The formal analysis becomes much more complex, with no general intuitions to guide us to attempt a characterization of other kinds of equilibria. We therefore restricted to only a possibility result.

6. While deliberation can give rise to full revelation (of types), it may also fail to transmit the desired information due to problems associated with the existence of a fully revealing equilibrium . Besides, time constraints on deliberations may render deliberations ineffective, in which case secret sequential advice, from constrained message sets, seems a natural procedure to follow.

7. We also indicate why the lessons from sequential communication cannot be taken to simultaneous advice, as the ranking between transparency and secrecy may well change.

This work fills a gap in the optimal design of transparency of decision making in organizations and political decisions by studying alternative disclosure protocols when career-concerned experts deliberate over issues.

A Appendix

■ **Section 3 materials.** The joint distribution of signal and state, given t_e, is as follows:

Table A.1: Joint distribution of signal and state

	a	b
α	qt_e	$(1-q)(1-t_e)$
β	$q(1-t_e)$	$(1-q)t_e$

An expert infers the state from her signal using Bayes' rule, e.g., $\Pr(\omega = a \mid s_e = \alpha, t_e) = \frac{qt_e}{qt_e+(1-q)(1-t_e)}$. We assume that the distribution of the experts' signals conditional on the state are independent. Below we report the joint distribution over state, signals and expert abilities. Given our assumption on independence (of both signals and abilities), we have (to facilitate reading we divide the distribution into two tables, and for only the following tables let $q' = (1-q)$):

Table A.2: Joint distribution of state, signals and abilities

	$t_i = \lambda, t_j = \lambda$	$t_i = \lambda, t_j = \xi$	$t_i = \xi, t_j = \lambda$	$t_i = \xi, t_j = \xi$
a, α, α	$q\theta^2\lambda^2$	$q\theta(1-\theta)\lambda\xi$	$q\theta(1-\theta)\lambda\xi$	$q(1-\theta)^2\xi^2$
a, α, β	$q\theta^2\lambda(1-\lambda)$	$q\theta(1-\theta)\lambda(1-\xi)$	$q\theta(1-\theta)\xi(1-\lambda)$	$q(1-\theta)^2\xi(1-\xi)$
a, β, α	$q\theta^2(1-\lambda)\lambda$	$q\theta(1-\theta)(1-\lambda)\xi$	$q\theta(1-\theta)(1-\xi)\lambda$	$q(1-\theta)^2(1-\xi)\xi$
a, β, β	$q\theta^2(1-\lambda)^2$	$q\theta(1-\theta)(1-\lambda)(1-\xi)$	$q\theta(1-\theta)(1-\xi)(1-\lambda)$	$q(1-\theta)^2(1-\xi)^2$

Table A.3: Joint distribution of state, signals and abilities

	$t_i = \lambda, t_j = \lambda$	$t_i = \lambda, t_j = \xi$	$t_i = \xi, t_j = \lambda$	$t_i = \xi, t_j = \xi$
b, α, α	$q'\theta^2(1-\lambda)^2$	$q'\theta(1-\theta)(1-\lambda)(1-\xi)$	$q'\theta(1-\theta)(1-\xi)(1-\lambda)$	$q'(1-\theta)^2(1-\xi)^2$
b, α, β	$q'\theta^2(1-\lambda)\lambda$	$q'\theta(1-\theta)(1-\lambda)\xi$	$q'\theta(1-\theta)(1-\xi)\lambda$	$q'(1-\theta)^2(1-\xi)\xi$
b, β, α	$q'\theta^2\lambda(1-\lambda)$	$q'\theta(1-\theta)\lambda(1-\xi)$	$q'\theta(1-\theta)\xi(1-\lambda)$	$q'(1-\theta)^2\xi(1-\xi)$
b, β, β	$q'\theta^2\lambda^2$	$q'\theta(1-\theta)\lambda\xi$	$q'\theta(1-\theta)\lambda\xi$	$q'(1-\theta)^2\xi^2$

The beliefs of expert i who moves first, conditional on her ability and signal, are given by $\Pr(\omega \mid s_i, t_i)$ as follows (refer Table A.1):

$$\begin{aligned}
\Pr(a \mid \alpha, t_i) &= \tfrac{qt_i}{qt_i+(1-q)(1-t_i)} &> \tfrac{1}{2},\\
\Pr(b \mid \alpha, t_i) &= \tfrac{(1-q)(1-t_i)}{qt_i+(1-q)(1-t_i)} &< \tfrac{1}{2},\\
\Pr(a \mid \beta, t_i) &= \tfrac{q(1-t_i)}{q(1-t_i)+(1-q)t_i} &< \tfrac{1}{2},\\
\Pr(b \mid \beta, t_i) &= \tfrac{(1-q)t_i}{q(1-t_i)+(1-q)t_i} &> \tfrac{1}{2}\,.
\end{aligned} \tag{A.1}$$

The inequalities follow from Assumption 1 and Fact 1. Note that despite the prior bias in favor of state a ($q > \frac{1}{2}$), signal β reverses this belief for either ability of expert.

Next consider expert j who moves second. Suppose she were to deduce the first expert's signal, but not her ability, from the observed recommendation. She would then update her beliefs conditional on (s_i, s_j, t_j), as follows (Assumption 1 and Fact 1 are used to establish the inequalities):

$$\begin{aligned}
\Pr(a \mid \alpha, \alpha, t_j) &= \tfrac{qt_jk}{qt_jk+(1-q)(1-t_j)(1-k)} &> \tfrac{1}{2},\\
\Pr(b \mid \alpha, \alpha, t_j) &= \tfrac{(1-q)(1-t_j)(1-k)}{qt_jk+(1-q)(1-t_j)(1-k)} &< \tfrac{1}{2},\\
\Pr(a \mid \beta, \beta, t_j) &= \tfrac{q(1-t_j)(1-k)}{(1-q)t_jk+q(1-t_j)(1-k)} &< \tfrac{1}{2},\\
\Pr(b \mid \beta, \beta, t_j) &= \tfrac{(1-q)t_jk}{(1-q)t_jk+q(1-t_j)(1-k)} &> \tfrac{1}{2}.
\end{aligned} \tag{A.2}$$

These four probabilities indicate that when the second expert's signal matches that of the first, then irrespective of ability, her posterior on state a (b) is higher than that on b (a) when she sees signal α (β).

When the first expert's signal is deduced as β and the second expert of ability λ receives signal α, we have:

$$\begin{aligned}
\Pr(a \mid \beta, \alpha, \lambda) &= \tfrac{q\lambda(1-k)}{q\lambda(1-k)+(1-q)(1-\lambda)k} &> \tfrac{1}{2},\\
\Pr(b \mid \beta, \alpha, \lambda) &= \tfrac{(1-q)(1-\lambda)k}{q\lambda(1-k)+(1-q)(1-\lambda)k} &< \tfrac{1}{2}.
\end{aligned} \tag{A.3}$$

When the first expert's signal is deduced as α and the second expert of ability ξ receives signal β, we have:

$$\begin{aligned}
\Pr(a \mid \alpha, \beta, \xi) &= \tfrac{q(1-\xi)k}{q(1-\xi)k+(1-q)\xi(1-k)} &> \tfrac{1}{2},\\
\Pr(b \mid \alpha, \beta, \xi) &= \tfrac{(1-q)\xi(1-k)}{q(1-\xi)k+(1-q)\xi(1-k)} &< \tfrac{1}{2}.
\end{aligned} \tag{A.4}$$

■ **List of strategies (relevant for Section 4 analysis).** $\mathbf{V}_e$ can be partitioned into the following class of strategies:

Partial type revelation. Strategies reveal the expert's ability only if she were to get signal s

but not s': $\mathbf{V}_e^{st} = \{v \mid v \in \mathbf{V}_i;\ \#A(v) = 1 \text{ or } \#B(v) = 1\}$.[26]

Type revealing. Strategies reveal the expert's ability: $\mathbf{V}_e^{t} = \{(A, B, A, B), (B, A, B, A)\}$.

Contrarian recommendation. Experts of both abilities recommend different from their signal: $\mathbf{V}_e^{c} = \{(B, B, A, A)\}$.

Revelation-Contrarian. Experts of only one ability recommend their signal while the other ability recommend contrarian: $\mathbf{V}_e^{sc} = \{(A, B, B, A), (B, A, A, B)\}$.

Truthful recommendation. $\mathbf{V}_e^{s} = \{(A, A, B, B)\}$.

Babbling. An expert's recommendation is completely uninformative:

$$\mathbf{V}_e^{b} = \{(A, A, A, A), (B, B, B, B)\}.$$

■ The proofs

Proof of Lemma 1. Let $\phi(k) = \frac{k}{1-k}$. Since $\frac{q}{1-q} > 1$, in Fig. 5 the graph of $\frac{q}{1-q}\phi(k)$ ($\frac{1-q}{q}\phi(k)$) lies above (below) $\phi(k)$.[27] Now the assertions follow from Fig. 5. ■

Figure 5: Partitioning k & q

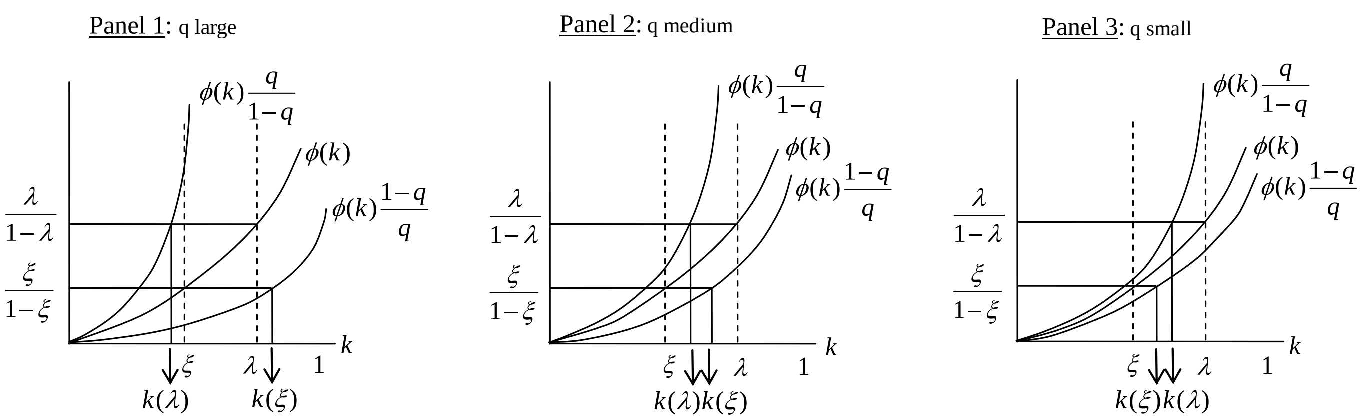

Proof of Lemma 2. **Case 1**: Let $v_j(v_i) \in \mathbf{V}_j^{st} = \{\mathbf{v} \mid \mathbf{v} \in \mathbf{V}_j;\ \#A(v) = 1 \text{ or } \#B(v) = 1\}$. Consider (i) $v_j(v_i) = (A, A, A, B)$, or (ii) $v_j(v_i) = (A, A, B, A)$. The other possibilities can be analyzed similarly.

[26] The cardinality of $\mathbf{V}_e^{st}$ is eight.

[27] Keeping ξ and λ constant, k increases from ξ to λ as θ varies from 0 to 1. The function $\phi(k)$ is increasing and convex with $\phi(0) = 0$ and $\lim_{k\to 1}\phi(k) = \infty$. Therefore, $k(\xi)$ and $k(\lambda)$ exist such that $\frac{\xi}{1-\xi} = \phi(k(\xi))\frac{1-q}{q}$ and $\frac{\lambda}{1-\lambda} = \phi(k(\lambda))\frac{q}{1-q}$.

(i) Suppose $\nu_j(\nu_i) = (A, A, A, B)$. Then, $E_j^t(\nu_i, B, a) = E_j^t(\nu_i, B, b) = \lambda$. So

$$\Pi_j^t(\nu_i, B, \alpha, t_j) = \lambda,$$

and

$$\begin{aligned} E_j^t(\nu_i, A, a) &= \left(\frac{\theta\lambda}{\theta\lambda + (1-\theta)\xi + (1-\theta)(1-\xi)}\right)\lambda + \left(\frac{(1-\theta)\xi + (1-\theta)(1-\xi)}{\theta\lambda + (1-\theta)\xi + (1-\theta)(1-\xi)}\right)\xi < \lambda, \\ E_j^t(\nu_i, A, b) &= \left(\frac{\theta(1-\lambda)}{(1-\theta\lambda-(1-\theta)\xi) + (1-\theta)\xi}\right)\lambda + \left(\frac{(1-\theta)\xi + (1-\theta)(1-\xi)}{(1-\theta\lambda-(1-\theta)\xi) + (1-\theta)\xi}\right)\xi < \lambda. \end{aligned}$$

So, for both abilities $t_j \in \{\xi, \lambda\}$ who observe α, recommending B is better than recommending A as:

$$\Pi_j^t(\nu_i, A, \alpha, t_j) = \Pr(a|\nu_i, \alpha, t_j) \cdot E_j^t(\nu_i, A, a) + \Pr(b|\nu_i, \alpha, t_j) \cdot E_j^t(\nu_i, A, b) < \lambda = \Pi_j^t(\nu_i, B, \alpha, t_j).$$

Thus, $\nu_j(\nu_i) = (A, A, A, B)$ cannot be an equilibrium strategy.

(ii) Now suppose $\nu_j(\nu_i) = (A, A, B, A)$. Then, $E_j^t(\nu_i, B, a) = E_j^t(\nu_i, B, b) = \xi$. Using arguments similar to above, one can see that

$$\Pi_j^t(\nu_i, B, \beta, \xi) = \xi < \Pi_j^t(\nu_i, A, \beta, \xi),$$

so $\nu_j(\nu_i) = (A, A, B, A)$ cannot be an equilibrium strategy.

Case 2: Consider $\nu_j(\nu_i) = (A, B, A, B)$. Then

$$\begin{aligned} &E_j^t(\nu_i, B, a) = E_j^t(\nu_i, B, b) = \lambda, \qquad E_j^t(\nu_i, A, a) = E_j^t(\nu_i, A, b) = \xi \\ \text{and} \qquad &\Pi_j^t(\nu_i, A, \alpha, \xi) = \xi < \lambda = \Pi_j^t(\nu_i, B, \alpha, \xi). \end{aligned}$$

So $\nu_j(\nu_i) = (A, B, A, B)$ cannot be an equilibrium strategy. The strategy $\nu_j(\nu_i) = (B, A, B, A)$ can be eliminated likewise.

Case 3: Let $\nu_j(\nu_i) \in \mathbf{V}_j^{sc} = \{(B, A, A, B), (A, B, B, A)\}$. Consider the case where $\nu_j(\nu_i) = (B, A, A, B)$. Then,

$$\begin{aligned} E_j^t(\nu_i, B, b) &= E_j^t(\nu_i, A, a) = \left(\frac{\theta\lambda}{\theta\lambda + (1-\theta)(1-\xi)}\right)\lambda + \left(\frac{(1-\theta)(1-\xi)}{\theta\lambda + (1-\theta)(1-\xi)}\right)\xi \equiv \eta, \\ E_j^t(\nu_i, B, a) &= E_j^t(\nu_i, A, b) = \left(\frac{\theta(1-\lambda)}{\theta(1-\lambda) + (1-\theta)\xi}\right)\lambda + \left(\frac{(1-\theta)\xi}{\theta(1-\lambda) + (1-\theta)\xi}\right)\xi \equiv \eta'. \end{aligned}$$

Note that

$$\eta > \eta',$$

$$\text{since} \qquad \frac{\theta\lambda}{\theta\lambda + (1-\theta)(1-\xi)} > \frac{\theta(1-\lambda)}{\theta(1-\lambda) + (1-\theta)\xi}, \quad \text{i.e.,} \quad \frac{\lambda}{1-\lambda} > \frac{1-\xi}{\xi},$$

given Assumption 1 and $\xi > q > 1/2$.

Also,

$$\begin{aligned}
\Pi_j^t(\nu_i, B, \alpha, \xi) &= \Pr(a \mid \nu_i, \alpha, \xi)\eta' + \Pr(b \mid \nu_i, \alpha, \xi)\eta, \\
\Pi_j^t(\nu_i, A, \alpha, \xi) &= \Pr(a \mid \nu_i, \alpha, \xi)\eta + \Pr(b \mid \nu_i, \alpha, \xi)\eta'.
\end{aligned}$$

In equilibrium we need $\Pi_j^t(\nu_i, B, \alpha, \xi) \geq \Pi_j^t(\nu_i, A, \alpha, \xi)$, which is true if and only if (as $\eta > \eta'$):

$$\Pr(b \mid \nu_i, \alpha, \xi) \geq \Pr(a \mid \nu_i, \alpha, \xi).$$

If the first expert babbled then $\Pr(\omega \mid \nu_i, \alpha, \xi) = \Pr(\omega \mid \alpha, \xi)$. So due to (A.1) the above cannot hold. Now suppose the first expert recommends truthfully. But then, due to (A.2) (see also (5)), it must be the case that the first expert observed β and recommended $\nu_i = B$ as she recommends her signal. (If the first expert observed α and recommended her signal, then that would have meant $\Pr(b \mid \alpha, \alpha, t_j) \geq \Pr(a \mid \alpha, \alpha, t_j)$, contradicting the first two inequalities in (A.2) for $t_j = \xi$.) Now,

$$\begin{aligned}
\Pi_j^t(\nu_i, A, \beta, \xi) &= \Pr(a \mid \beta, \beta, \xi)\eta + \Pr(b \mid \beta, \beta, \xi)\eta', \\
\Pi_j^t(\nu_i, B, \beta, \xi) &= \Pr(a \mid \beta, \beta, \xi)\eta' + \Pr(b \mid \beta, \beta, \xi)\eta.
\end{aligned}$$

In equilibrium we need $\Pi_j^t(\nu_i, A, \beta, \xi) \geq \Pi_j^t(\nu_i, B, \beta, \xi)$, which is true if and only if (as $\eta > \eta'$):

$$\Pr(a \mid \beta, \beta, \xi) \geq \Pr(b \mid \beta, \beta, \xi).$$

But then this contradicts the last two inequalities in (A.2) for $t_j = \xi$. So $\nu_j(\nu_i) = (B, A, A, B)$ cannot be an equilibrium strategy.

Next consider the case where $\nu_j(\nu_i) = (A, B, B, A)$. Then,

$$\begin{aligned}
E_j^t(\nu_i, B, a) &= E_j^t(\nu_i, A, b) = \left(\frac{\theta\lambda}{\theta\lambda + (1-\theta)(1-\xi)}\right)\lambda + \left(\frac{(1-\theta)(1-\xi)}{\theta\lambda + (1-\theta)\xi}\right)\xi \equiv \eta, \\
E_j^t(\nu_i, B, b) &= E_j^t(\nu_i, A, a) = \left(\frac{\theta(1-\lambda)}{\theta(1-\lambda) + (1-\theta)\xi}\right)\lambda + \left(\frac{(1-\theta)\xi}{\theta(1-\lambda) + (1-\theta)\xi}\right)\xi \equiv \eta'.
\end{aligned}$$

As shown previously,

$$\eta > \eta'.$$

Also,

$$\begin{aligned}\Pi_j^t(\nu_i, A, \alpha, \xi) &= \Pr(a \mid \nu_i, \alpha, \xi)\eta' + \Pr(b \mid \nu_i, \alpha, \xi)\eta, \\ \Pi_j^t(\nu_i, B, \alpha, \xi) &= \Pr(a \mid \nu_i, \alpha, \xi)\eta + \Pr(b \mid \nu_i, \alpha, \xi)\eta'.\end{aligned}$$

In equilibrium we need $\Pi_j^t(\nu_i, A, \alpha, \xi) \geq \Pi_j^t(\nu_i, B, \alpha, \xi)$, which is true if and only if (as $\eta > \eta'$):

$$\Pr(b \mid \nu_i, \alpha, \xi) \geq \Pr(a \mid \nu_i, \alpha, \xi).$$

Again, if the first expert babbled then $\Pr(\omega \mid \nu_i, \alpha, \xi) = \Pr(\omega \mid \alpha, \xi)$. So due to (A.1) the above cannot hold. Now suppose the first expert recommends truthfully. But then due to (A.2) it must be the case that the first expert observed β and recommended $\nu_i = B$. Now,

$$\begin{aligned}\Pi_j^t(\nu_i, B, \beta, \xi) &= \Pr(a \mid \beta, \beta, \xi)\eta + \Pr(b \mid \beta, \beta, \xi)\eta', \\ \Pi_j^t(\nu_i, A, \beta, \xi) &= \Pr(a \mid \beta, \beta, \xi)\eta' + \Pr(b \mid \beta, \beta, \xi)\eta.\end{aligned}$$

In equilibrium we need $\Pi_j^t(\nu_i, B, \beta, \xi) \geq \Pi_j^t(\nu_i, B, \beta, \xi)$, which is true if and only if (as $\eta > \eta'$):

$$\Pr(a \mid \beta, \beta, \xi) \geq \Pr(b \mid \beta, \beta, \xi).$$

But then again this contradicts the last two inequalities in (A.2) for $t_j = \xi$. So $\nu_j(\nu_i) = (A, B, B, A)$ cannot be an equilibrium strategy. ∎

Proof of Lemma 3. If the first expert babbled then $\Pr(\omega \mid \nu_i, s_j, t_j) = \Pr(\omega \mid s_j, t_j)$. Then, as in part (i) of Lemma 1, $\nu_j(\nu_i) = (A, A, B, B)$ is an equilibrium strategy.

Now let the first expert recommend truthfully and let $\nu_j(\nu_i) = (A, A, B, B)$ (we will have to show that this is an equilibrium). Then,

$$\begin{aligned}E_j^t(\nu_i, B, b) &= E_j^t(\nu_i, A, a) = \left(\frac{\theta\lambda}{\theta\lambda + (1-\theta)\xi}\right)\lambda + \left(\frac{(1-\theta)\xi}{\theta\lambda + (1-\theta)\xi}\right)\xi \equiv \gamma, \\ E_j^t(\nu_i, B, a) &= E_j^t(\nu_i, A, b) = \left(\frac{\theta(1-\lambda)}{\theta(1-\lambda) + (1-\theta)(1-\xi)}\right)\lambda + \left(\frac{(1-\theta)(1-\xi)}{\theta(1-\lambda) + (1-\theta)(1-\xi)}\right)\xi \equiv \gamma'.\end{aligned}$$

These outside evaluations of j's ability, under transparency, do not depend on the first expert's recommendation as per the projected strategy of j.

Note that,

$$\gamma > \gamma'.$$

Now first consider when j observes signal β and she is of ability ξ. Her expected payoffs are:

$$\begin{aligned}
\Pi_j^t(\nu_i, B, \beta, \xi) &= \Pr(a \mid \nu_i, \beta, \xi)\gamma' + \Pr(b \mid \nu_i, \beta, \xi)\gamma, \\
\Pi_j^t(\nu_i, A, \beta, \xi) &= \Pr(a \mid \nu_i, \beta, \xi)\gamma + \Pr(b \mid \nu_i, \beta, \xi)\gamma'.
\end{aligned}$$

In equilibrium we need $\Pi_j^t(\nu_i, B, \beta, \xi) \geq \Pi_j^t(\nu_i, A, \beta, \xi)$, which is true if and only if (as $\gamma > \gamma'$):

$$\Pr(b \mid \nu_i, \beta, \xi) \geq \Pr(a \mid \nu_i, \beta, \xi). \tag{A.5}$$

This can hold <u>only if</u> the first expert were to observe β and recommend B, otherwise (A.4) would be violated; this establishes the necessity of part (i) condition of the lemma. That (A.5) will be satisfied <u>if</u> $s_i = \beta$ and $\nu_i = B$ follows from (A.2).

Similarly, the condition for truthful recommendation by j for $t_j = \lambda$ and $s_j = \beta$ is given by

$$\Pr(b \mid \nu_i, \beta, \lambda) \geq \Pr(a \mid \nu_i, \beta, \lambda),$$

which, for $s_i = \beta$ and $\nu_i = B$, will be satisfied, given (A.2).

Next suppose j observes signal α. Then for $t_j = \xi$,

$$\begin{aligned}
\Pi_j^t(\nu_i, A, \alpha, \xi) &= \Pr(a \mid \beta, \alpha, \xi)\gamma + \Pr(b \mid \beta, \alpha, \xi)\gamma', \\
\Pi_j^t(\nu^1, B, \alpha, \xi) &= \Pr(a \mid \beta, \alpha, \xi)\gamma' + \Pr(b \mid \beta, \alpha, \xi)\gamma.
\end{aligned}$$

The equilibrium requirement $\Pi_j^t(\nu_i, A, \alpha, \xi) \geq \Pi_j^t(\nu_i, B, \alpha, \xi)$, given $\gamma > \gamma'$, is equivalent to:

$$\begin{aligned}
&\Pr(a \mid \beta, \alpha, \xi) \geq \Pr(b \mid \beta, \alpha, \xi) \quad \text{i.e.,} \quad \Pr(a \mid \beta, \alpha, \xi) \geq \frac{1}{2} \quad \text{i.e.,} \quad \frac{q\xi(1-k)}{q\xi(1-k) + (1-q)(1-\xi)k} \geq \frac{1}{2} \\
&\text{i.e.,} \qquad k \leq \frac{q\xi}{q\xi + (1-q)(1-\xi)} = k(\xi),
\end{aligned} \tag{A.6}$$

establishing the necessity requirement of part (ii) of the lemma. Sufficiency is completed with the next incentive compatibility verification.

Truthful recommendation by expert j for $t_j = \lambda$ and $s_j = \alpha$ requires

$$\Pr(a \mid \beta, \alpha, \lambda) \geq \Pr(b \mid \beta, \alpha, \lambda),$$

which is satisfied given (A.3). ■

Proof Lemma 5. (i) Let

$$\gamma_i = \frac{\theta\lambda}{k} \cdot \lambda + \frac{(1-\theta)\xi}{k} \cdot \xi, \qquad \gamma_i' = \frac{\theta(1-\lambda)}{1-k} \cdot \lambda + \frac{(1-\theta)(1-\xi)}{1-k} \cdot \xi.$$

By Assumption 1, we have:

$$\gamma_i > \gamma_i'.$$

Assume that expert 1 reveals her signal (we will have to show that this is indeed so in equilibrium). Since O's beliefs depend on i's recommendation, which he observes, the second expert's recommendation does not matter. Then,

$$E_i^t(A, \nu_j, a) = E_i^t(B, \nu_j, b) = E_i^t(A, a) = E_i^t(B, b) = \gamma_i,$$
$$E_i^t(A, \nu_j, b) = E_i^t(B, \nu_j, a) = E_i^t(A, b) = E_i^t(B, a) = \gamma_i'.$$

Hence, as the second expert's recommendation does not matter,

$$\begin{aligned}
\Pi_i^t(A, \nu_j(A), \alpha, t_i) &= \Pr(a|\alpha, t_i)\gamma_i + (1 - \Pr(a|\alpha, t_i))\gamma_i', \\
\Pi_i^t(B, \nu_j(B), \alpha, t_i) &= \Pr(a|\alpha, t_i)\gamma_i' + (1 - \Pr(a|\alpha, t_i))\gamma_i, \\
\Pi_i^t(A, \nu_j(A), \beta, t_i) &= \Pr(a|\beta, t_i)\gamma_i + (1 - \Pr(a|\beta, t_i))\gamma_i', \\
\Pi_i^t(B, \nu_j(B), \beta, t_i) &= \Pr(a|\beta, t_i)\gamma_i' + (1 - \Pr(a|\beta, t_i))\gamma_i.
\end{aligned}$$

For i to reveal her signal, we need to show that:

$$\Pi_i^t(A, \nu_j(A), \alpha, t_i) \geq \Pi_i^t(B, \nu_j(B), \alpha, t_i),$$
$$\Pi_i^t(B, \nu_j(B), \beta, t_i) \geq \Pi_i^t(A, \nu_j(A), \beta, t_i)$$

or,

$$\Pr(a|\alpha, t_i)\gamma_i + (1 - \Pr(a|\alpha, t_i))\gamma_i' \geq \Pr(a|\alpha, t_i)\gamma_i' + (1 - \Pr(a|\alpha, t_i))\gamma_i,$$
$$\Pr(a|\beta, t_i)\gamma_i' + (1 - \Pr(a|\beta, t_i))\gamma_i \geq \Pr(a|\beta, t_i)\gamma_i + (1 - \Pr(a|\beta, t_i))\gamma_i'$$

or,

$$\Pr(a|\alpha, t_i)(\gamma_i - \gamma_i') \geq \Pr(b|\alpha, t_i))(\gamma_i - \gamma_i'),$$
$$\Pr(b|\beta, t_i)(\gamma_i - \gamma_i') \geq \Pr(a|\beta, t_i)(\gamma_i - \gamma_i')$$

or, as $\gamma_i > \gamma'_i$,

$$\Pr(a|\alpha, t_i) \geq \Pr(b|\alpha, t_i),$$

$$\Pr(b|\beta, t_i) \geq \Pr(a|\beta, t_i)$$

which is true since $\Pr(a|\alpha, t_i) = \frac{t_i q}{t_i q + (1-t_i)(1-q)} > \frac{(1-t_i)(1-q)}{t_i q + (1-t_i)(1-q)} = \Pr(b|\alpha, t_i)$, or equivalently $t_i q > (1-t_i)(1-q)$, or equivalently $q + t_i > 1$ given that $q > 1/2$ and $t_i > q$ (Assumption 1), and likewise for the second inequality. So the first expert recommends her signal in equilibrium.

(ii) See the proof of Lemma 2 where strategies similar to $\nu_i \in \mathbf{V}_i \backslash \{\mathbf{V}_i^s \cup \mathbf{V}_i^b\}$ but by expert j (following any specific recommendation ν_i) were eliminated. The method of eliminations here will follow the same principles. The only difference is that for expert i there is no prior history except the null history, which simplifies the conditional probability calculations.

(iii) Expert i babbling can be supported in equilibrium by assuming that any deviation, say when i recommends B given that the posited equilibrium strategy is (A, A, A, A) (say), is punished by the outsider's belief that i must be of ability ξ. This is worse than the pooling ability k with which i will be perceived if she plays her equilibrium strategy. ■

Proof of Proposition 2. (i) The assumption $k(\xi) \leq k \leq k(\lambda)$ is equivalent to (see Panel 3 of Fig. 5):

$$\frac{\xi}{1-\xi}\frac{q}{1-q} \leq \frac{k}{1-k} \leq \frac{1-q}{q}\frac{\lambda}{1-\lambda}\,. \tag{A.7}$$

D's beliefs for the proposed strategies would be as follows, using Tables A.2 and A.3 (the tuple (AB) etc. are ordered pairs where the first coordinate denotes the first expert's advice):

$$\begin{aligned}
\Pr(a \mid AB) &= \frac{q\theta(1-\lambda)k}{q\theta(1-\lambda)k + (1-q)\theta\lambda(1-k)} \quad \leq \frac{1}{2}\,,\\
\Pr(a \mid AA) &= \frac{qk[k+(1-\theta)(1-\xi)]}{qk[k+(1-\theta)(1-\xi)] + (1-q)(1-k)[1-k+(1-\theta)\xi]} \quad > \frac{1}{2}\,,\\
\Pr(a \mid BA) &= \frac{q\theta\lambda(1-k)}{q\theta\lambda(1-k) + (1-q)\theta(1-\lambda)k} \quad > \frac{1}{2}\,,\\
\Pr(a \mid BB) &= \frac{q(1-k)[1-k+(1-\theta)\xi]}{q(1-k)[1-k+(1-\theta)\xi] + (1-q)k[k+(1-\theta)(1-\xi)]} \quad < \frac{1}{2}.
\end{aligned} \tag{A.8}$$

The first and third inequalities follow given the RHS inequality of (A.7), and the second and fourth inequalities follow given the LHS inequality of (A.7).

Let $d(\nu_i, \nu_j)$ be D's decision. Given D's beliefs (A.8), we have

$$d(AB) = B, \quad d(AA) = A, \quad d(BA) = A, \quad d(BB) = B.$$

Thus, if D's decision is A, then O knows that the recommendation profile is either (AA) or (BA). Conditional on observing D's decision d and state ω, O's beliefs about the experts' abilities are as follows:

$$\begin{aligned} \Pr(t=\lambda \mid A, a) &= \tfrac{\theta\lambda}{\theta\lambda+(1-\theta)k}, & \Pr(t=\lambda \mid A, b) &= \tfrac{\theta(1-\lambda)}{\theta(1-\lambda)+(1-\theta)(1-k)}, \\ \Pr(t=\lambda \mid B, b) &= \tfrac{\theta\lambda}{\theta\lambda+(1-\theta)k}, & \Pr(t=\lambda \mid B, a) &= \tfrac{\theta(1-\lambda)}{\theta(1-\lambda)+(1-\theta)(1-k)}, \end{aligned} \tag{A.9}$$

where $t=\lambda$ is the event of a randomly chosen expert being of ability λ.[28]

Let us consider the second expert's payoff:

$$\Pi_j^s(v_i, v_j, s_j, t_j) = \Pr(a|v_i, s_j, t_j)E_j^s(v_i, v_j, d(v_i, v_j), a) + \Pr(b|v_i, s_j, t_j)E_j^s(v_i, v_j, d(v_i, v_j), b),$$

where

$$\begin{aligned} E_j^s(A,A,A,a) &= \left[\tfrac{\theta\lambda}{\theta\lambda+(1-\theta)k}\right](\lambda-\xi)+\xi = E_j^s(B,A,A,a), \\ E_j^s(A,A,A,b) &= \left[\tfrac{\theta(1-\lambda)}{\theta(1-\lambda)+(1-\theta)(1-k)}\right](\lambda-\xi)+\xi = E_j^s(B,A,A,b), \\ E_j^s(B,B,B,b) &= \left[\tfrac{\theta\lambda}{\theta\lambda+(1-\theta)k}\right](\lambda-\xi)+\xi = E_j^s(A,B,B,b), \\ E_j^s(B,B,B,a) &= \left[\tfrac{\theta(1-\lambda)}{\theta(1-\lambda)+(1-\theta)(1-k)}\right](\lambda-\xi)+\xi = E_j^s(A,B,B,a). \end{aligned} \tag{A.10}$$

So, consider the second expert of ability t_j with signal s_j who sees a recommendation A. Then,

$$\begin{aligned} \Pi_j^s(A,A,s_j,t_j) &= \left[\Pr(a|\alpha,s_j,t_j)\left\{\frac{\theta\lambda}{\theta\lambda+(1-\theta)k}\right\} + \Pr(b|\alpha,s_j,t_j)\left\{\frac{\theta(1-\lambda)}{\theta(1-\lambda)+(1-\theta)(1-k)}\right\}\right](\lambda-\xi)+\xi, \\ \Pi_j^s(A,B,s_j,t_j) &= \left[\Pr(a|\alpha,s_j,t_j)\left\{\frac{\theta(1-\lambda)}{\theta(1-\lambda)+(1-\theta)(1-k)}\right\} + \Pr(b|\alpha,s_j,t_j)\left\{\frac{\theta\lambda}{\theta\lambda+(1-\theta)k}\right\}\right](\lambda-\xi)+\xi. \end{aligned}$$

Since $\frac{\theta\lambda}{\theta\lambda+(1-\theta)k} > \frac{\theta(1-\lambda)}{\theta(1-\lambda)+(1-\theta)(1-k)}$ and $\Pr(a|\alpha,\alpha,t_j) > 1/2$ (by (A.2)), following a recommendation of A it is optimal for both abilities of the second expert to recommend A, when she gets signal α. Due to (A.4), $\Pr(a|\alpha,\beta,\xi) > 1/2$. So, it is optimal for the second expert of ability ξ who gets signal β to herd and recommend A. Due to (6) and the condition that $\frac{\lambda}{1-\lambda} \geq \frac{q}{1-q}\frac{k}{1-k}$, we have $\Pr(b|\alpha,\beta,\lambda) > 1/2$. So, following an A-recommendation it is optimal

[28] To illustrate how the beliefs are derived, let us consider $\Pr\left(t=\lambda \mid A, a\right)$. Note that $d=A$ if and only if the recommendation profile is (AA) or (BA). Given $\omega=a$, the probability of the event (AA) is $k[k+(1-\theta)(1-\xi)]$ and that of (BA) is $(1-k)\theta\lambda$, so $\Pr\left(d=A \mid a\right) = \theta\lambda+(1-\theta)k$.

The probability of the event that the recommendation profile is (AA) or (BA) and a randomly selected expert is of ability λ, given $\omega=a$, is $\Pr\left((AA\cup BA)\cap\lambda\lambda \mid a\right)+\frac{1}{2}\left[\Pr\left((AA\cup BA)\cap\lambda\xi \mid a\right)+\Pr\left((AA\cup BA)\cap\xi\lambda \mid a\right)\right] = \left[\theta\lambda\,\theta\lambda+\theta(1-\lambda)\,\theta\lambda\right]+\frac{1}{2}\left[\left\{\theta\lambda\,(1-\theta)\xi+\theta\lambda\,(1-\theta)(1-\xi)\right\}+\left\{(1-\theta)\xi\,\theta\lambda+(1-\theta)(1-\xi)\,\theta\lambda\right\}\right] = \theta\lambda$. So, $\Pr\left(t=\lambda \mid A, a\right) = \frac{\theta\lambda}{\theta\lambda+(1-\theta)k}$. The other beliefs follow similarly.

for the second expert of ability λ, who gets signal β, to truthfully recommend B.

Now consider the second expert of ability t_j with signal s_j who sees a recommendation B. Then,

$$\Pi_j^s(B,B,s_j,t_j) = \left[\Pr(a|\beta,s_j,t_j)\left\{\frac{\theta(1-\lambda)}{\theta(1-\lambda)+(1-\theta)(1-k)}\right\} + \Pr(b|\beta,s_j,t_j)\left\{\frac{\theta\lambda}{\theta\lambda+(1-\theta)k}\right\}\right](\lambda-\xi)+\xi,$$
$$\Pi_j^s(B,A,s_j,t_j) = \left[\Pr(a|\beta,s_j,t_j)\left\{\frac{\theta\lambda}{\theta\lambda+(1-\theta)k}\right\} + \Pr(b|\beta,s_j,t_j)\left\{\frac{\theta(1-\lambda)}{\theta(1-\lambda)+(1-\theta)(1-k)}\right\}\right](\lambda-\xi)+\xi.$$

Again due to (A.2), $\Pr(b|\beta,\beta,t_j) > 1/2$. So, following a recommendation of B, it is optimal for the second expert to recommend B when she gets signal β. Due to (A.3), following a recommendation of B, it is optimal for the second expert to truthfully recommend A when she is of ability λ and gets signal α. Finally, consider the second expert of ability ξ who gets a signal α. From (5), she will herd and recommend B if and only if

$$\frac{(1-q)(1-\xi)k}{q\xi(1-k)+(1-q)(1-\xi)k} \geq \frac{1}{2} \quad \Leftrightarrow \quad \frac{k}{1-k} \geq \frac{q}{1-q}\frac{\xi}{1-\xi},$$

which is satisfied given (A.7).

Let us now consider the first expert and suppose she recommends A. Given the second expert's recommendation strategy, the first expert knows that if the second expert were to get signal α, then irrespective of her ability she would recommend A; if she were to get signal β and her ability were ξ she would recommend A, and if her ability were λ she would recommend B. Hence, if $\omega = a$, then the probability of the second expert recommending A would be $k+(1-\theta)(1-\xi)$, and of recommending B would be $\theta(1-\lambda)$. And if $\omega = b$, then the probability of the second expert recommending A would be $(1-k)+(1-\theta)\xi$ and of recommending B would be $\theta\lambda$. Hence, the first expert's expected payoff from recommending A is

$$\Pi_i^s(A,s_i,t_i) = \Pr(a|s_i,t_i)\left\{[k+(1-\theta)(1-\xi)]\rho+\theta(1-\lambda)\rho'\right\} + \Pr(b|s_i,t_i)\left\{[(1-k)+(1-\theta)\xi]\rho'+\theta\lambda\rho\right\},$$
where from (A.9), $\rho = \dfrac{\theta\lambda}{\theta\lambda+(1-\theta)k}$, $\quad \rho' = \dfrac{\theta(1-\lambda)}{\theta(1-\lambda)+(1-\theta)(1-k)}$.

Similarly, if the first expert were to recommend B, then conditional on $\omega = a$ the second expert would recommend A with probability $\theta\lambda$ and recommend B with probability $(1-k)+(1-\theta)\xi$; conditional on $\omega = b$ the second expert would recommend A with probability $\theta(1-\lambda)$ and B with probability $k+(1-\theta)(1-\xi)$. Thus,

$$\Pi_i^s(B,s_i,t_i) = \Pr(a|s_i,t_i)\left\{\theta\lambda\rho+[(1-k)+(1-\theta)\xi]\rho'\right\} + \Pr(b|s_i,t_i)\left\{\theta(1-\lambda)\rho'+[k+(1-\theta)(1-\xi)]\rho\right\}.$$

Hence,

$$\Pi_i^s(A, s_i, t_i) - \Pi_i^s(B, s_i, t_i) = \Pr(a|s_i, t_i)\{(1-\theta)(\rho-\rho')\} - \Pr(b|s_i, t_i)\{(1-\theta)(\rho-\rho')\}.$$

Therefore, as $\rho - \rho' > 0$ and $\theta < 1$, we have

$$\Pi_i^s(A, s_i, t_i) \geq \Pi_i^s(B, s_i, t_i) \quad \text{if and only if} \quad \Pr(a|s_i, t_i) \geq \Pr(b|s_i, t_i).$$

Since $\Pr(a|\alpha, t_i) > \Pr(b|\alpha, t_i)$ and $\Pr(b|\beta, t_i) > \Pr(a|\beta, t_i)$ for any $t_i \in \{\xi, \lambda\}$, the first expert recommends A when she gets signal α and recommends B when she gets signal β.

Thus, under the stated condition (A.7), the proposed equilibrium will exist. It follows from Lemma 1 that for *small* q, condition (A.7) is met.

(ii) It is easy to see from (A.10) that $E_j^s(.,.,A,a) - E_j^s(.,.,B,a) = E_j^s(.,.,B,b) - E_j^s(.,.,A,b)$, so there will be no conformity bias generated in the partial type revelation equilibrium. ■

Proof of Lemma 7. The proof will rely on Figs. 2 and 5.

$$\underline{k}(\xi) < \xi \underbrace{\Leftrightarrow}_{\text{Fig. 2}} c(\underline{k}(\xi)) > c(\xi) \Leftrightarrow \frac{\xi}{1-\xi} > \frac{1-\xi}{\xi}\frac{q}{1-q}\frac{1+\xi}{2-\xi}.$$

Start with the RHS of the above inequality to show that

$$\frac{1-\xi}{\xi}\frac{q}{1-q}\frac{1+\xi}{2-\xi} \underbrace{<}_{\text{since } \xi>1/2} \frac{1-\xi}{\xi}\frac{q}{1-q}\frac{\xi}{1-\xi} = \frac{q}{1-q} \underbrace{<}_{\text{since } \xi>q} \frac{\xi}{1-\xi},$$

establishing our first claim.

Next,

$$\begin{aligned}
&\bar{k}(\xi) > k(\xi) \underbrace{\Leftrightarrow}_{\text{Fig. 2}} c^{-1}(\bar{k}(\xi)) > c^{-1}(k(\xi)) \underbrace{\Leftrightarrow}_{\text{Fig. 2}} \frac{\xi}{1-\xi} > \frac{k(\xi)}{1-k(\xi)}\frac{1-q}{q}\frac{2-k(\xi)}{1+k(\xi)} \\
&\Leftrightarrow \frac{\xi}{1-\xi} > \phi(k(\xi))\frac{1-q}{q}\frac{2-k(\xi)}{1+k(\xi)} \underbrace{=}_{\text{Fig. 5}} \frac{\xi}{1-\xi}\frac{2-k(\xi)}{1+k(\xi)} \Leftrightarrow 1 > \frac{2-k(\xi)}{1+k(\xi)} \\
&\Leftrightarrow 1+k(\xi) > 2-k(\xi) \Leftrightarrow k(\xi) > \frac{1}{2},
\end{aligned}$$

which is true. This establishes the second claim. ■

Proof of Proposition 3. For the first part of the proof we shall assume that the experts recommend their signals in equilibrium and then show that such an equilibrium indeed exists under the stated parameter restrictions. The only information that O will have about the recommendations is through d. Recall, given that the experts recommend their signals, D selects B only if two recommendations are in favor of B; otherwise D selects A (Lemma 6). Therefore when $d = A$, O knows that one of three pairs of signals, (α, α), (α, β), (β, α), could have resulted. When $d = B$, O knows that (β, β) resulted. O's relevant posteriors are then calculated, using Tables A.2 and A.3, as follows:[29]

$$\begin{aligned}
\Pr(t = \lambda \mid A, a) &= \frac{\theta[\lambda + (1-\lambda)(\theta\lambda + (1-\theta)\xi)]}{\theta[\lambda + (1-\lambda)(\theta\lambda + (1-\theta)\xi)] + (1-\theta)[\xi + (1-\xi)(\theta\lambda + (1-\theta)\xi)]} = \frac{\theta[\lambda + (1-\lambda)k]}{k(2-k)},\\
\Pr(t = \lambda \mid A, b) &= \frac{\theta[(1-\lambda) + \lambda(1-\theta\lambda - (1-\theta)\xi)]}{\theta[(1-\lambda) + \lambda(1-\theta\lambda - (1-\theta)\xi)] + (1-\theta)[(1-\xi) + \xi(1-\theta\lambda - (1-\theta)\xi)]} = \frac{\theta(1-\lambda k)}{1-k^2},\\
\Pr(t = \lambda \mid B, b) &= \frac{\theta\lambda}{\theta\lambda + (1-\theta)\xi} = \frac{\theta\lambda}{k},\\
\Pr(t = \lambda \mid B, a) &= \frac{\theta(1-\lambda)}{\theta(1-\lambda) + (1-\theta)(1-\xi)} = \frac{\theta(1-\lambda)}{1-k}.
\end{aligned}$$

Also, $\Pr(t = \xi \mid A, a) = 1 - \Pr(t = \lambda \mid A, a)$, and likewise for the remaining posteriors.

Define

$$\begin{aligned}
x' &= \Pr(t = \lambda \mid A, a)\lambda + (1 - \Pr(t = \lambda \mid A, a))\xi,\\
y' &= \Pr(t = \lambda \mid A, b)\lambda + (1 - \Pr(t = \lambda \mid A, b))\xi,\\
x'' &= \Pr(t = \lambda \mid B, b)\lambda + (1 - \Pr(t = \lambda \mid B, b))\xi,\\
y'' &= \Pr(t = \lambda \mid B, a)\lambda + (1 - \Pr(t = \lambda \mid B, a))\xi.
\end{aligned}$$

Now, these are O's expectations of expert abilities given any state and D's decision.

Let the second expert j recommend her signal. Consider the first expert i's strategy who has observed signal α. If she recommends A then irrespective of what the second expert recommends, $d = A$. So expert i receives a payoff:[30]

$$\Pi_i^s(A, \alpha, t_i) = \Pr(a|\alpha, t_i)x' + \Pr(b|\alpha, t_i)y' = \frac{qt_i}{H(t_i)}x' + \frac{(1-q)(1-t_i)}{H(t_i)}y',$$

where $H(t_i) \equiv qt_i + (1-q)(1-t_i)$.

If she recommends B then d depends on whether j observes α or β. This payoff can be

[29]To alert the reader, here the first conditioning variable is the decision d.

[30]$\Pr(a|\alpha, t_i) = \frac{\Pr(\alpha|a,t_i)\cdot\Pr(a,t_i)}{\Pr(\alpha|a,t_i)\cdot\Pr(a,t_i)+\Pr(\alpha|b,t_i)\cdot\Pr(b,t_i)} = \frac{qt_i\cdot\Pr(t_i)}{qt_i\cdot\Pr(t_i)+(1-q)(1-t_i)\cdot\Pr(t_i)} = \frac{qt_i}{qt_i+(1-q)(1-t_i)}$.

written as:

$$\begin{aligned}
\Pi_i^s(B,\alpha,t_i) &= \Pr(a|s_i=\alpha,t_i)\Big[\sum_{t_j=\lambda,\xi}\Pr(s_j=\alpha|a,t_j)\cdot\Pr(t_j)x' + \sum_{t_j=\lambda,\xi}\Pr(s_j=\beta|a,t_j)\cdot\Pr(t_j)y''\Big] \\
&\quad + \Pr(b|s_i=\alpha,t_i)\Big[\sum_{t_j=\lambda,\xi}\Pr(s_j=\alpha|b,t_j)\cdot\Pr(t_j)y' + \sum_{t_j=\lambda,\xi}\Pr(s_j=\beta|b,t_j)\cdot\Pr(t_j)x''\Big] \\
&= \Pr(a|\alpha,t_i)\Big[\Big\{\Pr(\alpha|a,\lambda)\theta+\Pr(\alpha|a,\xi)(1-\theta)\Big\}x' + \Big\{\Pr(\beta|a,\lambda)\theta+\Pr(\beta|a,\xi)(1-\theta)\Big\}y''\Big] \\
&\quad + \Pr(b|\alpha,t_i)\Big[\Big\{\Pr(\alpha|b,\lambda)\theta+\Pr(\alpha|b,\xi)(1-\theta)\Big\}y' + \Big\{\Pr(\beta|b,\lambda)\theta+\Pr(\beta|b,\xi)(1-\theta)\Big\}x''\Big] \\
&= \Pr(a|\alpha,t_i)\Big[\big\{\lambda\theta+\xi(1-\theta)\big\}x' + \big\{(1-\lambda)\theta+(1-\xi)(1-\theta)\big\}y''\Big] \\
&\quad + \Pr(b|\alpha,t_i)\Big[\big\{(1-\lambda)\theta+(1-\xi)(1-\theta)\big\}y' + \big\{\lambda\theta+\xi(1-\theta)\big\}x''\Big] \\
&= \frac{qt_i}{H(t_i)}\big[kx'+(1-k)y''\big] + \frac{(1-q)(1-t_i)}{H(t_i)}\big[kx''+(1-k)y'\big].
\end{aligned}$$

Substituting terms and with some algebra we obtain:

$$\begin{aligned}
&\Pi_i^s(A,\alpha,t_i) \geq \Pi_i^s(B,\alpha,t_i) \Leftrightarrow qt_i\big[(1-k)(x'-y'')\big] \geq (1-q)(1-t_i)\big[k(x''-y')\big] \\
&\Leftrightarrow \frac{t_i}{1-t_i} \geq \frac{k}{1-k}\frac{1-q}{q}\frac{x''-y'}{x'-y''} \Leftrightarrow \frac{t_i}{1-t_i} \geq c^{-1}(k),
\end{aligned}$$

since

$$\frac{x''-y'}{x'-y''} = \frac{2-k}{1+k} > 0. \tag{A.11}$$

Let i observe signal β. Then,

$$\begin{aligned}
\Pi_i^s(B,\beta,t_i) &= \Pr(a|s_i=\beta,t_i)\Big[\sum_{t_j=\lambda,\xi}\Pr(s_j=\alpha|a,t_j)\cdot\Pr(t_j)x' + \sum_{t_j=\lambda,\xi}\Pr(s_j=\beta|a,t_j)\cdot\Pr(t_j)y''\Big] \\
&\quad + \Pr(b|s_i=\beta,t_i)\Big[\sum_{t_j=\lambda,\xi}\Pr(s_j=\alpha|b,t_j)\cdot\Pr(t_j)y' + \sum_{t_j=\lambda,\xi}\Pr(s_j=\beta|b,t_j)\cdot\Pr(t_j)x''\Big] \\
&= \Pr(a|\beta,t_i)\Big[\Big\{\Pr(\alpha|a,\lambda)\theta+\Pr(\alpha|a,\xi)(1-\theta)\Big\}x' + \Big\{\Pr(\beta|a,\lambda)\theta+\Pr(\beta|a,\xi)(1-\theta)\Big\}y''\Big] \\
&\quad + \Pr(b|\beta,t_i)\Big[\Big\{\Pr(\alpha|b,\lambda)\theta+\Pr(\alpha|b,\xi)(1-\theta)\Big\}y' + \Big\{\Pr(\beta|b,\lambda)\theta+\Pr(\beta|b,\xi)(1-\theta)\Big\}x''\Big] \\
&= \Pr(a|\beta,t_i)\Big[\big\{\lambda\theta+\xi(1-\theta)\big\}x' + \big\{(1-\lambda)\theta+(1-\xi)(1-\theta)\big\}y''\Big] \\
&\quad + \Pr(b|\beta,t_i)\Big[\big\{(1-\lambda)\theta+(1-\xi)(1-\theta)\big\}y' + \big\{\lambda\theta+\xi(1-\theta)\big\}x''\Big] \\
&= \frac{q(1-t_i)}{J(t_i)}[kx'+(1-k)y''] + \frac{(1-q)t_i}{J(t_i)}[(1-k)y'+kx''],
\end{aligned}$$

where $J(t_i) \equiv (1-q)t_i + q(1-t_i)$. On the other hand,

$$\Pi_i^s(A,\beta,t_i) = \Pr(a|\beta,t_i)x' + \Pr(b|\beta,t_i)y' = \frac{q(1-t_i)}{J(t_i)}x' + \frac{(1-q)t_i}{J(t_i)}y'.$$

It is easy to check that

$$\begin{aligned}\Pi_i^s(B,\beta,t_i) \geq \Pi_i^s(A,\beta,t_i) \;&\Leftrightarrow\; (1-q)t_i\big[k(x''-y')\big] \geq q(1-t_i)\big[(1-k)(x'-y'')\big] \\ &\Leftrightarrow\; \frac{t_i}{1-t_i} \geq c(k).\end{aligned}$$

But $\frac{t_i}{1-t_i} \geq c^{-1}(k)$ and $\frac{t_i}{1-t_i} \geq c(k)$ if and only if $\underline{k}(\xi) \leq k \leq \bar{k}(\xi)$.

Now consider the second expert j. If she sees a first period recommendation of A then she knows that $d = A$, so she is indifferent between recommending A and recommending B. Thus, recommending her signal is a best response for j.

Suppose j sees recommendation B. She knows that i recommended her signal, so $s_i = \beta$. Let j observe signal α. Then,

$$\begin{aligned}\Pi_j^s(A, s_i=\beta, s_j=\alpha, t_j) &= \Pr(a|s_i=\beta, s_j=\alpha, t_j)x' + \Pr(b|s_i=\beta, s_j=\alpha, t_j)y', \\ \Pi_j^s(B, s_i=\beta, s_j=\alpha, t_j) &= \Pr(a|s_i=\beta, s_j=\alpha, t_j)y'' + \Pr(b|s_i=\beta, s_j=\alpha, t_j)x'',\end{aligned}$$

where $\Pr(a|s_i = \beta, s_j = \alpha, t_j) = \frac{qt_j(1-k)}{qt_j(1-k)+(1-q)(1-t_j)k}$ and $\Pr(b|s_i = \beta, s_j = \alpha, t_j) = 1 - \Pr(a|s_i = \beta, s_j = \alpha, t_j)$ (see (A.3) and (5)). It is then easy to verify that

$$\begin{aligned}&\Pi_j^s(A, s_i=\beta, s_j=\alpha, t_j) \geq \Pi_j^s(B, s_i=\beta, s_j=\alpha, t_j) \\ \Leftrightarrow\; &\Pr(a|s_i=\beta, s_j=\alpha, t_j)\big[x'-y''\big] \geq \Pr(b|s_i=\beta, s_j=\alpha, t_j)\big[x''-y'\big] \\ \Leftrightarrow\; &qt_j(1-k)\big[x'-y''\big] \geq (1-q)(1-t_j)k\big[x''-y'\big] \\ \Leftrightarrow\; &\frac{t_j}{1-t_j} \geq c^{-1}(k). \end{aligned} \tag{A.12}$$

Let j observe signal β. Then,

$$\begin{aligned}\Pi_j^s(B, s_i=\beta, s_j=\beta, t_j) &= \Pr(a|s_i=\beta, s_j=\beta, t_j)y'' + \Pr(b|s_i=\beta, s_j=\beta, t_j)x'', \\ \Pi_j^s(A, s_i=\beta, s_j=\beta, t_j) &= \Pr(a|s_i=\beta, s_j=\beta, t_j)x' + \Pr(b|s_i=\beta, s_j=\beta, t_j)y',\end{aligned}$$

where $\Pr(a|s_i = \beta, s_j = \beta, t_j) = \frac{q(1-t_j)(1-k)}{(1-q)t_jk+q(1-t_j)(1-k)}$ and $\Pr(b|s_i = \beta, s_j = \beta, t_j) = 1 -$

$\Pr(a|s_i = \beta, s_j = \beta, t_j)$ (see (A.2)). Now it can be verified that

$$\begin{aligned}
& \Pi_j^s(B, s_i = \beta, s_j = \beta, t_j) \geq \Pi_j^s(A, s_i = \beta, s_j = \beta, t_j) \\
\Leftrightarrow \quad & \Pr(a|s_i = \beta, s_j = \beta, t_j)[y'' - x'] \geq \Pr(b|s_i = \beta, s_j = \beta, t_j)[y' - x''] \\
\Leftrightarrow \quad & q(1-t_j)(1-k)[y'' - x'] \geq (1-q)t_j k[y' - x''] \\
\Leftrightarrow \quad & \frac{t_j}{1-t_j} \geq \frac{1-k}{k}\frac{q}{1-q}\frac{x' - y''}{x'' - y'} \Leftrightarrow \frac{t_j}{1-t_j} \geq c(k).
\end{aligned}$$

But then again, $\frac{t_j}{1-t_j} \geq c^{-1}(k)$ and $\frac{t_j}{1-t_j} \geq c(k)$ if and only if $\underline{k}(\xi) \leq k \leq \bar{k}(\xi)$.

When $k > \bar{k}(\xi)$ we will construct an SDE under which the second expert babbles. In fact, such an SDE exists for all parameters k.

Let

$$\begin{aligned}
x &\equiv \frac{\theta\lambda}{\theta\lambda + (1-\theta)\xi}\lambda + \frac{(1-\theta)\xi}{\theta\lambda + (1-\theta)\xi}\xi, \\
y &\equiv \frac{\theta(1-\lambda)}{\theta(1-\lambda) + (1-\theta)(1-\xi)}\lambda + \frac{(1-\theta)(1-\xi)}{\theta(1-\lambda) + (1-\theta)(1-\xi)}\xi.
\end{aligned}$$

Here x is the expected accuracy of an expert's signal who has made a correct recommendation based on her signal alone, and y is similarly defined for an expert who has made an inaccurate recommendation. Thus, x and y are the market's imputed value of an expert's skill based on the accuracy of her recommendation, when the recommendation is based only on her own signal.

Due to Assumption 1,

$$x > k > y.$$

To calculate equilibrium beliefs about the experts' abilities, suppose the experts follow their respective strategies as specified above (the optimality of strategies to be verified later). Then, for recommendation pairs (A, A) and (A, B), the decision maker will choose $d = A$ (Lemma 6). If the first expert recommends B (i.e. for recommendation pairs (B, A) and (B, B)), the decision maker will select $d = B$ based only on the first expert's recommendation;

the second recommendation is uninformative. Hence for O the beliefs are as follows:

$$\begin{aligned}
\Pr(t=\lambda \mid A,a) = \quad \Pr(t=\lambda \mid B,b) &= \tfrac{\theta\lambda}{\theta\lambda+(1-\theta)\xi},\\
\Pr(t=\xi \mid A,a) = \quad \Pr(t=\xi \mid B,b) &= \tfrac{(1-\theta)\xi}{\theta\lambda+(1-\theta)\xi},\\
\Pr(t=\lambda \mid A,b) = \quad \Pr(t=\lambda \mid B,a) &= \tfrac{\theta(1-\lambda)}{\theta(1-\lambda)+(1-\theta)(1-\xi)},\\
\Pr(t=\xi \mid A,b) = \quad \Pr(t=\xi \mid B,a) &= \tfrac{(1-\theta)(1-\xi)}{\theta(1-\lambda)+(1-\theta)(1-\xi)}.
\end{aligned}$$

Again, the beliefs are applicable to both experts, given that neither the experts' identities nor the timing of moves are revealed.

Now consider expert i who moves first. Let her observe α. If she recommends A, she receives

$$\Pi_i^s(A,\alpha,t_i) = \frac{qt_i}{qt_i+(1-q)(1-t_i)}x + \frac{(1-q)(1-t_i)}{qt_i+(1-q)(1-t_i)}y.$$

If she recommends B, her payoff is

$$\Pi_i^s(B,\alpha,t_i) = \frac{(1-q)(1-t_i)}{qt_i+(1-q)(1-t_i)}x + \frac{qt_i}{qt_i+(1-q)(1-t_i)}y.$$

Since $qt_i > (1-q)(1-t_i)$ and $x > y$, we have $\Pi_i^s(A,\alpha,t_i) > \Pi_i^s(B,\alpha,t_i)$. Now suppose she observes β. If she recommends B, her payoff is

$$\Pi_i^s(B,\beta,t_i) = \frac{(1-q)t_i}{(1-q)t_i+q(1-t_i)}x + \frac{q(1-t_i)}{(1-q)t_i+q(1-t_i)}y.$$

If she recommends A, she receives

$$\Pi_i^s(A,\beta,t_i) = \frac{q(1-t_i)}{(1-q)t_i+q(1-t_i)}x + \frac{(1-q)t_i}{(1-q)t_i+q(1-t_i)}y.$$

As $\frac{t_i}{1-t_i} > \frac{q}{1-q}$, and $x > y$, we have $\Pi_i^s(B,\beta,t_i) > \Pi_i^s(A,\beta,t_i)$. Hence, it is strictly optimal for i to recommend her signal, irrespective of her ability. Now, given a first period recommendation of A, the second expert j knows that $d = A$ (Lemma 6). Her payoff remains unchanged whether she recommends A or B. Similarly, if the first recommendation is B, j's recommendation is immaterial and $d = B$. Hence again, j's payoff remains unchanged whether she recommends A or B. So it is optimal for j to babble. ■

Proof of Proposition 4. Under the transparency protocol, Lemmas 2 through 5 state that the following PBE exist (i.e., a case of multiple equilibria): (i) both experts babble; (ii) one expert babbles and one recommends truthfully; (iii) the first expert recommends truthfully and the second expert babbles if the first recommendation is A and recommends truthfully if the first recommendation is B, provided that either the bias q is large, or the bias q is medium/small and $k \in [\xi, k(\xi)]$. The payoff to D under (i) is q, under (ii) is k, and under (iii) is $q[k + (1-k)k] + (1-q)k^2$ (this one follows applying the equilibrium strategies in Proposition 1). By Assumption 1, $k > q$. Since $2q > 1$ and $k < 1$, we have $q[k+(1-k)k] + (1-q)k^2 > k$.

Thus for large q, or medium/small q and $k \in [\xi, k(\xi)]$, the maximum equilibrium payoff of D is $q[k + (1-k)k] + (1-q)k^2$. When q is medium/small and $k > k(\xi)$, D's maximum equilibrium payoff is k.

Under the secrecy SDE in Proposition 3 the payoff of D is $q[k^2+2k(1-k)]+(1-q)k^2$, same as $q[k + (1-k)k] + (1-q)k^2$, when q is either large, or medium/small and $k \in [\xi, \bar{k}(\xi)]$. When q is medium/small and $k > \bar{k}(\xi)$, the first expert recommends her signal and the second expert babbles. The payoff of D then is k.

Since by Lemma 7 we have $\bar{k}(\xi) > k(\xi)$, on overall comparison D should choose secrecy over transparency: secrecy and transparency yield identical payoffs except for q medium/small and $k \in (k(\xi), \bar{k}(\xi)]$, when secrecy yields strictly higher payoff. ■

Proof of Proposition 5. (i) Recall that

$$k(\xi) = \frac{q\xi}{q\xi + (1-q)(1-\xi)}, \quad k(\lambda) = \frac{(1-q)\lambda}{(1-q)\lambda + q(1-\lambda)}$$

and k is a function of θ. We change variables and write k as a function of x,

$$k(x) = \frac{(x + 2q - 1)\xi\lambda}{(2q-1)\lambda + x\xi} \quad \text{for} \quad x \in [\underline{x}, \bar{x}],$$

where

$$\underline{x} \equiv \frac{(2q-1)^2\lambda(1-\xi)}{(1-q)\lambda(1-\xi) - q\xi(1-\lambda)}, \quad \bar{x} \equiv \frac{(1-q)\lambda(1-\xi) - q\xi(1-\lambda)}{\xi(1-\lambda)}.$$

Note that $k(\underline{x}) = k(\xi)$ and $k(\bar{x}) = k(\lambda)$. The denominator of $\underline{x}$ and the numerator of $\bar{x}$ are the same and the value is positive because $\frac{1-q}{q}\frac{\lambda}{1-\lambda} > \frac{q}{1-q}\frac{\xi}{1-\xi}$ (follows from $k(\xi) < k(\lambda)$) and $\frac{q}{1-q}\frac{\xi}{1-\xi} > \frac{\xi}{1-\xi}$ (as $q > \frac{1}{2}$). Furthermore

$$k'(x) = \frac{(2q-1)\xi\lambda(\lambda - \xi)}{[(2q-1)\lambda + x\xi]^2} > 0.$$

Hence $\{k : k(\xi) \le k \le k(\lambda)\} = \{k : k(\underline{x}) \le k \le k(\bar{x})\}$ and $x > 0$ for all $x \in [\underline{x}, \bar{x}]$.

We now start our proof by noting that D's payoff from the partial type revealing equilibrium,

$$k^2 + k(1-\theta)(1-\xi) + (1-k)\theta\lambda,$$

dominates the payoff from the signal revealing equilibrium,

$$q[k^2 + 2k(1-k)] + (1-q)k^2,$$

if and only if

$$(1-\theta)(1-\xi)k + \theta\lambda(1-k) > 2qk(1-k), \tag{A.13}$$

or equivalently,

$$(2q-1)\frac{\theta\lambda}{(1-\theta)(1-\xi)} + 2q\frac{\xi}{1-\xi} < \frac{k}{1-k}. \tag{A.14}$$

The following claim (which relates θ to x) will be used in the proof.

Claim 1: $(2q-1)\frac{\theta\lambda}{(1-\theta)(1-\xi)} = x\frac{\xi}{1-\xi}$ if and only if $k = \frac{(x+2q-1)\xi\lambda}{(2q-1)\lambda + x\xi}$.

This is because

$$\begin{aligned}
&(2q-1)\frac{\theta\lambda}{(1-\theta)(1-\xi)} = x\frac{\xi}{1-\xi} \iff (2q-1)\theta\lambda = x\xi(1-\theta)\\
\iff\ & \theta = \frac{x\xi}{(2q-1)\lambda + x\xi} \iff \theta(\lambda-\xi) = \frac{x\xi(\lambda-\xi)}{(2q-1)\lambda + x\xi}\\
\iff\ & \theta(\lambda-\xi) + \xi = \frac{x\xi(\lambda-\xi)}{(2q-1)\lambda + x\xi} + \xi \iff \theta\lambda + (1-\theta)\xi = \frac{(x+2q-1)\xi\lambda}{(2q-1)\lambda + x\xi}\\
\iff\ & k = \frac{(x+2q-1)\xi\lambda}{(2q-1)\lambda + x\xi}.
\end{aligned}$$

Hence, it follows from Claim 1 and (A.14) that D's payoff from the partial type revealing equilibrium dominates that from the signal revealing equilibrium if and only if

$$(x+2q)\frac{\xi}{1-\xi} < \frac{k}{1-k} \quad \left(\equiv \frac{(x+2q-1)\xi\lambda}{(2q-1)\lambda(1-\xi) + x\xi(1-\lambda)}\right) \tag{A.15}$$

or, as $x > 0$,

$$x^2\xi(1-\lambda) - 2x[(1-q)\lambda(1-\xi) - q\xi(1-\lambda)] + (2q-1)^2\lambda(1-\xi) < 0. \tag{A.16}$$

Call the LHS of (A.16) $f(x)$. So,

$$f(x) = x^2\xi(1-\lambda) - 2x[(1-q)\lambda(1-\xi) - q\xi(1-\lambda)] + (2q-1)^2\lambda(1-\xi).$$

Claim 2: $f'(x) < 0$ for all $x < (1-q)\frac{\lambda}{1-\lambda}\frac{(1-\xi)}{\xi} + q$.

Because

$$f'(x) = 2x\xi(1-\lambda) - 2[(1-q)\lambda(1-\xi) - q\xi(1-\lambda)],$$

$$\begin{aligned} f'(x) < 0 \quad &\Leftrightarrow \quad 2x\xi(1-\lambda) - 2[(1-q)\lambda(1-\xi) - q\xi(1-\lambda)] < 0 \\ &\Leftrightarrow \quad x < (1-q)\frac{\lambda}{1-\lambda}\frac{(1-\xi)}{\xi} + q. \end{aligned}$$

We now check whether (A.15), and therefore (A.16), holds at $x = \underline{x}$.

Claim 3:

$$(\underline{x} + 2q)\frac{\xi}{1-\xi} < \frac{k(\underline{x})}{1-k(\underline{x})}. \tag{A.17}$$

We know that the RHS of (A.17) is equal to $\frac{k(\xi)}{1-k(\xi)} = \frac{q}{1-q}\frac{\xi}{1-\xi}$. So we need to show that $(\underline{x} + 2q)\frac{\xi}{1-\xi} < \frac{q}{1-q}\frac{\xi}{1-\xi}$, or $(\underline{x} + 2q) < \frac{q}{1-q}$. Now,

$$\begin{aligned} &(\underline{x} + 2q) < \frac{q}{1-q} \quad \Leftrightarrow \quad \underline{x} < \frac{q(2q-1)}{1-q} \quad \Leftrightarrow \quad \frac{(2q-1)^2\lambda(1-\xi)}{(1-q)\lambda(1-\xi) - q\xi(1-\lambda)} < \frac{q(2q-1)}{1-q} \\ &\Leftrightarrow \quad (2q-1)(1-q)\lambda(1-\xi) < q(1-q)\lambda(1-\xi) - q^2\xi(1-\lambda) \\ &\Leftrightarrow \quad q^2\xi(1-\lambda) < (q-2q+1)(1-q)\lambda(1-\xi) \quad \Leftrightarrow \quad \frac{q^2}{(1-q)^2} < \frac{\lambda}{1-\lambda}\frac{1-\xi}{\xi} \\ &\Leftrightarrow \quad \frac{q}{1-q} < \sqrt{r}\ \text{, which is true.} \end{aligned}$$

So (A.15), and therefore (A.16), is satisfied at $x = \underline{x}$. Now by Claim 2, $f'(x)$ is negative for all x such that $\underline{x} < x < (1-q)\frac{\lambda}{1-\lambda}\frac{1-\xi}{\xi} + q$. So (A.15), and (A.16), will be satisfied for all x such that $\underline{x} < x < (1-q)\frac{\lambda}{1-\lambda}\frac{1-\xi}{\xi} + q$. Therefore, if $\bar{x} < (1-q)\frac{\lambda}{1-\lambda}\frac{1-\xi}{\xi} + q$, then (A.15) and (A.16) would hold for all $x \in [\underline{x}, \bar{x}]$. The condition is verified as follows:

$$\bar{x} = \frac{(1-q)\lambda(1-\xi) - q\xi(1-\lambda)}{\xi(1-\lambda)} = (1-q)\frac{\lambda}{1-\lambda}\frac{1-\xi}{\xi} - q < (1-q)\frac{\lambda}{1-\lambda}\frac{(1-\xi)}{\xi} + q.$$

This completes the proof of part (i).

(ii) When $\big(k(\xi), k(\lambda)\big) \cap \big(\bar{k}(\xi), \lambda\big] \neq \emptyset$, by subtracting the payoff of k when only the first expert reveals her signal from the partial type revelation equilibrium payoff we obtain:

$$\begin{aligned} & k^2 + k(1-\theta)(1-\xi) + (1-k)\theta\lambda - k \\ = \; & k^2 + k(1-\theta)(1-\xi) + (1-k)\theta\lambda - k(1+1-k) \\ = \; & k(1-\theta)(1-\xi) + (1-k)[\theta\lambda - \theta\lambda - (1-\theta)\xi] \\ = \; & (1-\theta)[k(1-\xi) - \xi(1-k)] > 0, \end{aligned}$$

establishing our claim. ■

References

Aumann, R.J. and Hart, S., "Long Cheap Talk," *Econometrica*, 71 (2003), 1619-1660.

Austen-Smith, D., "Interested Experts and Policy Advice: Multiple Referrals under Open Rule," *Games and Economic Behavior*, 5 (1993), 3-43.

Diermeier, D., Keane, M. and Merlo, A., "A Political Economy Model of Congressional Careers," *American Economic Review*, 95 (2005), 347-373.

Elster, J. (edited), *Secrecy and Publicity in Votes and Debates,* Cambridge University Press, (2015), ISBN 978-1-107-08336-3.

Fehrler, S. and Hughes, N., "How Transparency Kills Information Aggregation: Theory and Experiment," *American Economic Journal: Microeconomics*, 10 (2018), 181-209.

Forges, F. and Koessler, F., "Long Persuasion Games," *Journal of Economic Theory*, 143 (2008), 1-35.

Gersbach, H. and Hahn, V., "Should Individual Voting Records of Central Bankers be Published?" *Social Choice Welfare*, 30 (2008), 655-683.

Holmstrom, B., "Managerial Incentive Problems: A Dynamic Perspective," *Review of Economic Studies*, 66 (1999), 169-182.

Iaryczower, M., Shi, X. and Shum, M., "Can Words Get in the Way? The Effect of Deliberation in Collective Decision-Making," *Journal of Political Economy*, 2016, forthcoming.

Jackson, M.O. and Tan, X., "Deliberation, Disclosure of Information, and Voting," *Journal of Economic Theory*, 148 (2013), 2-30.

Krishna, V. and Morgan, J., "The Art of Conversation: Eliciting Information from Experts through Multistage Communication," *Journal of Economic Theory*, 117 (2004), 147-179.

Landier, A., Sraer, D. and Thesmar, D., "Optimal Dissent in Organizations," *Review of Economic Studies*, 76 (2009), 761-794.

Levy, G., "Decision Making in Committees: Transparency, Reputation and Voting Rules," *American Economic Review*, 97 (2007a), 150-168.

Levy, G., "Decision Making Procedures for Careerist Experts," *American Economic Review Papers and Proceedings*, 97 (2007b), 306-310.

McLennan, A., "Consequences of the Condorcet Jury Theorem for Beneficial Information Aggregation by Rational Agents," *American Political Science Review* 92 (1998), 413-418.

Meade, E. and Stasavage, D., "Publicity of Debate and the Incentive to Dissent: Evidence from the US Federal Reserve," *Economic Journal*, 118 (2008), 695-717.

Ottaviani, M. and Sorensen, P., "Information Aggregation in Debate: Who Should Speak First?" *Journal of Public Economics*, 81 (2001), 393-421.

Ottaviani, M. and Sorensen, P., "Professional Advice," *Journal of Economic Theory*, 126 (2006a), 120-142.

Ottaviani, M. and Sorensen, P., "Reputational Cheap Talk," *Rand Journal of Economics*, 37 (2006b), 155-175.

Ottaviani, M. and Sorensen, P., "The Strategy of Professional Forecasting," *Journal of Financial Economics*, 81 (2006c), 441-466.

Scharfstein, D. and Stein, J., "Herd Behavior and Investment," *American Economic Review*, 90 (2000), 465-479.

Seidmann, D., "A Theory of Voting Patterns and Performance in Private and Public Committees," *Social Choice Welfare*, 36 (2011), 49-74.

Sibert, A., "Monetary Policy Committees: Individual and Collective Reputations," *Review of Economic Studies*, 70 (2003), 649-665.

Visser, B. and Swank, O., "On Committees of Experts," *Quarterly Journal of Economics*, 122 (2007), 337-372.

Wolinsky, A., "Eliciting Information from Multiple Experts," *Games and Economic Behavior*, 41 (2002), 141-160.