A delegation-based theory of expertise*

Attila Ambrus[†], Volodymyr Baranovskyi[‡] and Aaron Kolb[§]

November 24, 2020

Abstract

We investigate information aggregation and competition in a delegation framework. An uninformed principal is unable to perform a task herself and must choose between one of two biased and imperfectly informed experts. In the focal equilibrium, experts exaggerate their biases, anticipating an ideological winner's curse. We show that having a second expert can benefit the principal, even when equally or more biased than the first expert. The principal can benefit from commitment to an "element of surprise," and prefers experts with equal rather than opposite biases.

1 Introduction

There are many situations in which a principal lacks the knowledge and expertise to perform a certain task, and therefore has to delegate the job to a qualified expert. Examples include a candidate running for office who has to hire an expert to work out her economic agenda, or the CEO of a pharmaceutical company who must delegate building a research and development division to a scientist. Further complicating the principal's situation is that experts tend to have systemic biases, preferring suboptimal actions from the principal's perspective.

In this paper we investigate a model in which a principal has to delegate a task to one of two experts. The need to delegate differentiates our model from models of expertise in which experts send cheap talk recommendations to the principal, such as Krishna and Morgan (2001b). In particular, we consider the following game. First, experts receive noisy and conditionally independent signals of a single dimensional state variable. The principal's ideal action is equal to the state, but each expert has a constant bias (either positive or negative) and a resulting ideal point different from the sender's. Next, the experts simultaneously

*We thank Ricardo Alonso, Alan Jaske, Navin Kartik, Scott Kominers, Peter Landry, David McAdams, Wolfgang Pesendorfer, Curtis Taylor, Xiao Yu Wang and participants of the 2016 Wallis Conference on Political Economy at Rochester for useful feedback. All errors are our own.

[†]Duke University, *attila.ambrus@duke.edu*

[‡]University of Rochester, *vbaranovskyi@gmail.com*

[§]Indiana University Kelley School of Business, *kolba@indiana.edu*

propose actions. A proposal is assumed to bind the expert to perform the given action whenever the principal delegates the task to him.[1] The principal then chooses one of the two offers, and the corresponding action is taken by the given expert. We are motivated by situations in which the principal originally has much less knowledge about the state than the experts, and correspondingly we assume that the principal's prior is improper uniform (diffuse) over a state space represented by the real line.

The particular game form we investigate is motivated by various applications. In general, our model best applies to situations in which the principal lacks the knowledge to implement or initiate changes in the proposed actions, so all she can do is solicit different proposals and choose one of them. The assumption that proposals commit the experts corresponds to common law, according to which an offer is a statement of terms on which the offeror is willing to be bound, and it shall become binding as soon as it is accepted by the person to whom it is addressed.[2] A different application for our model is political competition: starting from Downs (1957), most papers on political competition in a Hotelling (1929) framework assume that candidates are committed to the policies they announce in the campaign, and the electorate can only choose between the policies announced by the candidates.[3] In this context, the principal corresponds to the median voter, and the bonus corresponds to the rents from being in office.[4] Yet another application for our model is a setting where a legislative body (floor) seeks legislative proposals for the same bill from multiple committees, using a modified rule (see Gilligan and Krehbiel (1989), Krishna and Morgan (2001a)), meaning that the floor cannot amend the proposals and can only accept one of the proposed bills without modification, conforming to the basic assumptions of our model.[5]

Experts have preferences over both the policy outcome and whether they are selected. In particular, to model the latter, we allow for a bonus to the chosen expert, either as a monetary payment or as a non-monetary benefit, such as increased prestige in his profession. We investigate two cases, with the bonus amount given exogenously in one case and optimally chosen by the principal in the other.

The above game is very complex in general, due to the size of the strategy space. In

[1] Even if the principal might not have the knowledge to verify whether the expert indeed chose the action that he proposed, outside experts might be able to verify if that was the case and hence penalties can be imposed on experts deviating from their proposals.

[2] See Treitel (1999), p8.

[3] For theoretical motivations for this assumption, and empirical relevance in the political competition context, see Pétry and Collette (2009), Kartik et al. (2015), and papers cited therein.

[4] We make this connection explicit in Remark 1.

[5] Gilligan and Krehbiel (1989) analyze this situation with an additional option to the floor, in the form of not accepting either of the proposals and opting for a status quo outcome. As opposed to our model, Gilligan and Krehbiel (1989) assume perfectly informed experts (committees), which fundamentally changes the strategic interaction.

this paper we restrict attention to equilibria in which the experts' strategies are stationary with respect to signals, meaning that each expert's proposal is equal to his signal plus a constant. We consider a focus on such strategies, which treat states symmetrically, to be natural in a game with diffuse prior and preferences that are relatively stationary in the state, since in such an environment all states are perfectly symmetric (nothing distinguishes them besides the labeling, which can be arbitrarily rescaled). A further motivation comes from a companion paper (Ambrus and Kolb, 2020), in which we examine the possibility of extending the concept of ex ante expected payoffs to a larger class of games with diffuse prior (and hence bringing them into the realm of traditional game theory, in which payoffs have to be well-defined for any strategy profile). The companion paper shows that in our game, expert strategies need to be restricted to be stationary in order for well-defined ex ante expected payoffs to exist. In particular, given some weak conditions on the principal's set of strategies, essentially stationary strategies are the only strategies for which well-defined limit expected payoffs exist for any strategy profile when taking a sequence of proper priors diffusing (converging in a formal sense to the diffuse prior), with the limit not depending on the particular choice of sequence. This result shows that in order to obtain well-defined ex ante expected payoffs corresponding to all strategy profiles, one would need to restrict experts' strategies to stationary ones. In the current paper, stationary strategies are either characterized by a constant *markup* — the expert adds a fixed markup to whatever his signal is — or mixtures thereof. Here, we do not restrict experts' strategies, but simply focus on equilibria in which they play stationary strategies, and similarly to existing game theoretical models of improper prior (Friedman (1991), Klemperer (1999), Morris and Shin (2002, 2003), Myatt and Wallace (2014)), we only evaluate payoffs in the interim stage (after signal realizations).

Our main results show that having a second expert as a possible choice can benefit the principal, even when he is equally or more biased than the first expert. There are two effects driving this result. The first is informational: if both experts get selected with positive probability then both experts' signals are utilized, reducing the variance of the bias of the implemented action. The second one is a competition effect: if the bonus for being selected is strictly positive, experts competing to be selected reduce the markups on their offers, which for a range of bonus levels reduces the bias of the implemented action from the principal's perspective.

The equilibrium that can outperform simple delegation of the task to the less biased expert has a simple structure. Assuming without loss of generality that the sum of the experts' biases is nonnegative, it involves the principal selecting the minimum of the two proposals, for which reason we refer to it as the min equilibrium.[6] When the bonus is not too high relative to the noise in the experts' signals, such an equilibrium exists and it involves experts applying markups above their biases; if both biases are positive, this means that both experts exaggerate their biases (propose actions strictly about their ideal actions based purely on their private signals). This is because, similarly to the winner's curse phenomenon in common value auctions, being selected by the principal contains information on the other expert's signal (namely that his signal is higher), changing the optimal action of the expert. In equilibrium, proposals have to be optimal conditional on the event that the other expert's action proposal is higher. The feature of the equilibrium that similarly biased experts exaggerate in a particular direction (and hence the principal should choose the proposal least in the direction of the exaggeration) is in line with empirical evidence. For example, Zitzewitz (2001), Bernhardt et al. (2006) and Chen and Jiang (2006) find that financial analysts systematically exaggerate their forecasts relative to unbiased forecasts based on the analysts' information sets, while Iezzoni et al. (2012) report that 55% of doctors in a survey said that in the previous year they had been more positive about patients' prognoses than their medical histories warranted.

Whether the min equilibrium or simple delegation yields a higher payoff to the principal depends on the parameters of the model, but there are several salient cases in which the principal prefers the min equilibrium, including (i) when the experts have the same biases (as in settings where all available experts have similar agendas), and (ii) when the experts have biases of opposite sign and equal magnitude (as in settings where experts are chosen from two opposing camps).[7] These comparisons hold even when the bonus is zero, in which case the benefit to the principal is driven solely by the informational effect alluded to earlier rather than the competition effect. In case (i), the reduction in variance from choosing the lowest offer is maximized, and moreover, in this case, the bias in simple delegation is relatively large. In case (ii), the min equilibrium can feature both reduced variance and lower bias of the chosen action, since the winner's curse causes experts (and in particular, the expert with the negative bias) to choose markups above their biases. The principal also prefers min equilibrium to simple delegation when one expert's bias is positive and the other's is zero, despite the fact that under simple delegation the unbiased expert's incentives are perfectly aligned with the principal's. Here again, the min equilibrium reduces variance and has very

[6] This strategy is also feasible for a principal who can only process information in a coarse way, being only able to make binary comparisons between two offers and lacking the ability to measure the difference between them, as consumers in Kamenica (2008). Therefore such information processing constraints would not hurt the principal in the equilibria we investigate.

[7] The principal's payoffs are continuous in the parameters of the model for a particular type of equilibrium, hence the above comparisons are the same when the absolute values of the biases are close to each other but not exactly equal.

little negative effect on the expected bias of the chosen offer, since that bias is already very small. Applied to a political setting, the latter comparison between min equilibrium and simple delegation helps explain what voters may otherwise perceive as corruption — a politician may want to seek advice from a biased expert, even if it is common knowledge that she already has access to an unbiased expert.

We also compare the principal's payoffs when experts have equal versus opposite biases, and our results here are in contrast with some of the existing literature. In our model, assuming the min equilibrium is played, having two experts with identical biases yields a higher payoff than having two antagonist experts with opposite biases. In general, the expected bias of the implemented action is smaller with antagonist experts than with experts having the same bias, but this benefit is outweighed by a higher variance of the implemented action that arises because the expert with the lower bias is selected most of the time, and so the information from the other expert's signal is only utilized to a limited extent. This result contrasts with models of competition in persuasion (Milgrom and Roberts (1986), Gentzkow and Kamenica (2017)), in which antagonist experts benefit the principal by pressing each other to reveal more information,[8] and with the multi-sender cheap talk model of Krishna and Morgan (2001b), in which having a second sender with the same bias does not benefit the receiver.[9] See also Shin (1998) and Dewatripont and Tirole (1999) for different types of models making the case for adversarial procedures.

Our final main result shows that the principal's welfare is nonmonotonic in the magnitude of the bonus payment.[10] Increasing the bonus reduces the absolute values of the experts' markups, hence bringing their proposals closer to truthful reporting. Intuitively, a higher bonus increases competition among experts, leading them to decrease their proposals in the min equilibrium. This initially improves the principal's expected payoff by decreasing the expected bias of the implemented action. Eventually, however, the expected bias becomes negative, and further increases in the bonus then *increase* the magnitude of this negative expected bias, reducing the principal's payoff.[11] When the bonus comes from exogenous sources, the optimal bonus from the principal's perspective is always strictly positive, and

[8] Experts with identical agendas can be better for the principal than experts with opposing agendas in the persuasion model of Bhattacharya and Mukherjee (2013). The mechanism is rather different than in our paper, though: with similar experts an undesirable default action can provide strong incentives for both experts to reveal information.

[9] As opposed to cheap talk models with multiple senders and one receiver, where it tends to be better for information revelation if the senders are oppositely biased from the point of view of the receiver, in committee settings, where committee members can reveal information to each other, it helps information revelation if members have more similar preferences — see for example Li and Suen (2009).

[10] We also show nonmonotonicity in the precision of the experts' signals.

[11] For bonus values in this region, the principal still prefers to choose the lower offer because choosing the higher offer would result in a positive expected bias of larger magnitude.

is in the interior of the interval of bonuses for which the min equilibrium exists. When the bonus is paid by the principal, the optimal bonus amount is always strictly smaller than in the previous case, and depending on the parameters it can be either strictly positive or zero. In the political competition application of the model, the nonmonotonicity result implies that a small amount of office-seeking motivation can be beneficial for voters, but at higher levels a further increase in office-seek motivation can adversely affect voters' welfare.

We consider several extensions of our model, for equally biased experts. In the first one we allow the principal to commit ex ante to any mixture over a set of simple strategies, and show that when there is no bonus, such commitment leads to the same outcome as in the min equilibrium of the original game, hence the ability to commit does not improve the principal's welfare. On the other hand, we show how committing to choosing an inferior (in expectation) offer with some small probability — introducing an element of surprise — can improve the principal's welfare in the case of opposite biases of large magnitude by inducing the expert with positive bias to drastically reduce his large positive markup. The finding that the principal can benefit from committing to a mixed strategy in certain situations is consistent with an observed pattern of regulatory uncertainty. Ederer et al. (2014) show similarly that commitment to an opaque reward scheme reduces temptation to game the system in a principal-agent environment.

In the second extension we drop the dependence of the unselected expert's payoff on the implemented action, and instead assume that the expert gets a fixed outside option payoff. This variant of the model is more realistic in market transaction situations, such as when experts are car mechanics or doctors. The analysis of this version of the model is more involved, but we show that under some parameter restrictions similar min and max equilibria exist as in the baseline model. The fact that the unselected expert gets a fixed outside payment increases the experts' proposals in the min equilibrium.

We also discuss how our results extend to correlated signals, more than two experts, and dispensing with the functional form assumptions that give us analytical tractability in the main model.

2 Related Literature

The literature on delegation so far mainly focused on either the question of delegating the action choice versus retaining the right to take the action (Dessein, 2002; Li and Suen, 2004) or on optimally constraining the action choices of a particular expert (Holmström, 1977; Melumad and Shibano, 1991; Alonso and Matouschek, 2008). An exception is a literature on final offer arbitration (Gibbons, 1988; Mylovanov and Zapechelnyuk, 2013); however, in those papers, agents have misaligned preferences, while our model allows for a range of preferences including full alignment between experts. Krishna and Morgan (2008) investigate how monetary incentives can be used optimally in delegation to a single agent. More related to our investigation are papers introducing policy-relevant private information on the part of candidates into the context of the classic Downs (1957) model of political competition: Heidhues and Lagerlöf (2003), Laslier and Van der Straeten (2004), Loertscher et al. (2012), Gratton (2014), and closest to our model Kartik et al. (2015), as the latter focuses on cases when voters are relatively uninformed. Similarly to our setting, in the above papers politicians receive independent private signals about the state of the world and hence the optimal policy from the electorate's point of view. The main difference relative to our model is that the politicians in the above models do not have policy preferences, and they are purely office-motivated. For this reason neither their own private information nor the rival's private information directly affects their expected payoffs, and the candidates play a zero-sum game. In contrast, in our model the experts' signals are directly payoff-relevant for them, and their interests are partially aligned, as in higher states they would both like to induce higher actions. This leads to different equilibrium dynamics than in Kartik et al. (2015), and to distinct conclusions: in particular, in their model the electorate can never strictly benefit from the presence of a second candidate (relative to just a single one).[12]

There is also a line of literature extending the Downs (1957) model framework to politicians having mixed motivation (having both policy preferences and wanting to win), as in our model, starting with Wittman (1983) and Calvert (1985). Schultz (1996) and Martinelli and Matsui (2002) introduce asymmetric information in this context, but as opposed to our paper, with perfectly informed politicians. This leads to different conclusions, including that full revelation of information is possible in equilibrium when policy preferences are not too extreme. Martinelli (2001) allows for imperfectly informed politicians, but he considers the limiting case in which politicians' preferences become state-independent; in contrast, the experts' preferences in our model are state-dependent, which introduces and allows us to study the ideological winner's curse. Callander (2011) considers a model of sequential elections with imperfectly informed politicians having mixed motivations, but the issues investigated

[12]Our model approximates the model in Kartik et al. (2015) when the bonus payment is very large and so the agents mainly care about being selected. We find that for very large bonuses the only equilibria in our model involve delegating the action choice to a single agent, which is in line with the result on maximum informativeness of political competition in Kartik et al. (2015). Correspondingly, Kartik et al. (2015) discuss an extension of their model in which they show that allowing a small amount of ideological motivation for the candidates, and assuming that they are close to unbiased from the electorate's point of view implies that in equilibrium one candidate must be winning ex ante with probability close to 1. These results suggest that there is no discontinuity between no policy preference versus a small amount of policy preference for the agents. Our paper mainly focuses on cases in which agents' policy preferences are relatively important.

are different from ours and inherently dynamic: searching for a good policy in a complex environment, by trial and error.

Outside the delegation literature, Prendergast (1993) considers a context in which both a worker and a manager observes the state with noise, and the worker also receives a noisy signal about the manager's observation. In this model the worker has an incentive to cater to the manager and bias her report toward what she thinks the manager's observation is. Gerardi et al. (2009) investigates aggregation of expert opinions through a particular mechanism that approximates the first best outcome if signals are very accurate. Pesendorfer and Wolinsky (2003) investigate the effects of being able to solicit a second opinion from a different expert, in a dynamic model in which experts are not biased but it is costly for them to gather information.

Another line of literature investigates multi-sender extensions of the cheap talk model of Crawford and Sobel (1982), and finds that under certain conditions there can be equilibria in which the receiver can extract full or almost full information from the senders (Gilligan and Krehbiel (1989), Austen-Smith (1993), Wolinsky (2002), Battaglini (2002, 2004), Ambrus and Takahashi (2008), Ambrus and Lu (2014)). As opposed to the above papers, we investigate settings in which the principal cannot solicit information from experts and then take the action choice herself. Ottaviani and Sørensen (2006) consider a model with multiple experts with reputational concerns reporting sequentially on privately observed signals. The issues they focus on (potential herding behavior of experts) are very different than in the current paper.

Finally, while private information is exogenous in our model, recent work has considered strategic information acquisition in models of communication or delegation. These papers include Di Pei (2015), Argenziano et al. (2016), Migrow and Squintani (2018) and Deimen and Szalay (2018). Endogenous information is made tractable in these models by other model features such as cheap talk communication or the presence of a single expert.

3 Base Model

We consider the following multi-stage game with incomplete information. There are three players: a principal and two experts. The set of states of the world is $\mathbb{R}$, and we assume that the common prior distribution of states is diffuse (improper uniform).

In stage 0 state $\theta \in \mathbb{R}$ realizes. In stage 1 each expert $i = 1, 2$ receives a noisy private signal about the state of the world $s_i = \theta + \epsilon_i$, where $\epsilon_i \sim N(0, \sigma^2)$, and ϵ_1 and ϵ_2 are independent. In stage 2 each expert i proposes an action $a_i \in \mathbb{R}$ to the principal. In stage 3 the principal chooses one of the two experts, who then implements the action he proposed in stage 2.

Let real-valued functions $a_1(s_1)$ and $a_2(s_2)$ denote the strategies of expert 1 and expert 2 respectively, while $C(a_1, a_2) \in \{1, 2\}$ is the principal's choice strategy. If action $a = a_{C(a_1,a_2)}$ is taken then the principal's payoff is $V(a, \theta) = -(a - \theta)^2$, and the payoff of expert $i = 1, 2$ is $U_i(a, \theta, C(a_1, a_2)) = \mathbb{1}\{i = C(a_1, a_2)\}B - (a - \theta - b_i)^2$, where the indicator function $\mathbb{1}\{i = C(a_1, a_2)\}$ equals 1 if $i = C(a_1, a_2)$ and 0 otherwise. We call b_i the *bias* of expert i. Note that each expert experiences a quadratic loss which depends on the action chosen but not the identity of the chosen expert; in Section 6, we analyze the case where the quadratic loss applies only to the chosen expert. Without loss of generality, we assume that $b_1 + b_2 \geq 0$ and $b_1 \geq b_2$. We further assume that all parameters of the game are common knowledge.

In the analysis below, we focus on perfect Bayesian equilibria in which experts' strategies are stationary in the following sense: $a_i(s_i) = s_i + k_i$, where $k_1, k_2 \in \mathbb{R}$. In words, each expert applies a constant *markup* to his signal when forming a proposed action. With slight abuse of notation, we use simply $(k_1, k_2, C(a_1, a_2))$ to denote a strategy profile with constant markup strategies.

3.1 Best Responses to Stationary Expert Strategies

Here we analyze the principal's best response to constant markup strategies, where each expert's offer is his signal translated by a constant, and show that the best response only depends on the sum of the markups. As the principal has a quadratic loss function, her expected payoff can be decomposed into losses from the uncertainty about the true state (which is independent of her action) and the losses from the expected difference between the chosen action and the true state. Therefore the principal prefers the offer which is closer to her posterior expectation of the true state. After observing the offers, the principal's expectation about the true state is lower (higher) than the average of the experts' offers if and only if the sum of the markups is positive (negative). Figure 1 illustrates a case where the markups have positive sum.

Let $\arg\min\{a_1, a_2\}$ be defined as $\{1\}$ if $a_1 < a_2$, $\{2\}$ if $a_1 > a_2$, and $\{1, 2\}$ if $a_1 = a_2$. Similarly, let $\arg\max\{a_1, a_2\}$ be defined as $\{1\}$ if $a_1 > a_2$, $\{2\}$ if $a_1 < a_2$, and $\{1, 2\}$ if $a_1 = a_2$.

Theorem 1. *If experts follow constant markup strategies $a_i(s_i) = s_i + k_i$, then*

- *if $k_1+k_2 > 0$, the principal strictly prefers the lower offer, and $C(a_1, a_2) \in \arg\min\{a_1, a_2\}$;*
- *if $k_1+k_2 < 0$, the principal strictly prefers the higher offer, and $C(a_1, a_2) \in \arg\max\{a_1, a_2\}$;*
- *if $k_1 + k_2 = 0$, the principal is indifferent between the offers.*

Proof. After observing both offers, the principal updates her belief: $\theta|a_1, a_2 \sim N(\frac{1}{2}(a_1 + a_2) - \frac{1}{2}(k_1 + k_2), \frac{\sigma^2}{2})$. Therefore the principal's expected utility from choosing offer a_1 or a_2 is:

$$V(a_1) = \mathbb{E}[-(\theta - a_1)^2] = -\text{Var}(\theta) - (\mathbb{E}[\theta] - a_1)^2 = -\frac{\sigma^2}{2} - \left[\frac{1}{2}(a_2 - a_1) - \frac{1}{2}(k_1 + k_2)\right]^2$$
$$V(a_2) = \mathbb{E}[-(\theta - a_2)^2] = -\text{Var}(\theta) - (\mathbb{E}[\theta] - a_2)^2 = -\frac{\sigma^2}{2} - \left[\frac{1}{2}(a_1 - a_2) - \frac{1}{2}(k_1 + k_2)\right]^2.$$

Hence, $V(a_1) - V(a_2) = (a_2 - a_1)(k_1 + k_2)$, which immediately implies the statements in the theorem. □

Theorem 1 illustrates a contrast between our paper and the standard communication literature. In the latter, where proposals are mere communication and the principal has the ability to choose any action, the principal would back out the signals $s_i = a_i - k_i$ and would optimally take action $\frac{s_1+s_2}{2}$. Here, the principal knows that this would be the best action but is unable to take an action other than the two offers provided.

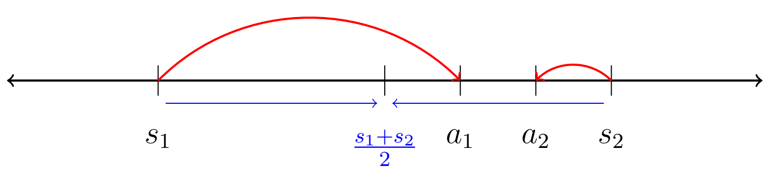

Figure 1: $k_1 > 0, k_2 < 0, k_1 + k_2 > 0$. The principal infers the signals s_i from the offers a_i given the markups and chooses the lower offer a_1, which lies closer to her expectation $\frac{s_1+s_2}{2}$.

An equilibrium $(k_1, k_2, C(a_1, a_2))$ is said to be a *min equilibrium*, if the minimum proposal is accepted and on average experts adjust their signals upwards: $k_1 + k_2 \geq 0$ and $C(a_1, a_2) \in \arg\min(a_1, a_2)$. In the min equilibrium, the principal's updated expectation of the state of the world is lower than the average of the two offers. Her best response is to choose the lower offer, which is closer to her expectation, as demonstrated in Figure 1. Likewise, an equilibrium $(k_1, k_2, C(a_1, a_2))$ is said to be a *max equilibrium* if $k_1 + k_2 \leq 0$ and $C(a_1, a_2) \in \arg\max(a_1, a_2)$.

When $k_1 + k_2 = 0$ then the principal's posterior expectation of θ is exactly the average of the two offers, and she is indifferent between the two.[13]

Remark 1 (Median Voter Interpretation)**.** *Although we refer to a single principal, the model readily captures voting applications, in which case the principal is the median voter choosing*

[13]This raises the possibility of equilibria in which the experts play constant markup strategies and the principal mixes between the lower and the higher offer with some fixed probability. Details on these equilibria are available from the authors upon request.

between two candidates for office with positions a_1 and a_2 on an ideological spectrum. To make this interpretation precise, consider a continuum of voters indexed by a preference parameter $\chi \in (-\infty, +\infty)$ with massless cumulative distribution function G and normalized so that $G(0) = \frac{1}{2}$. Voter χ obtains a payoff $-(a - \theta - \chi)^2$ when action a is taken, and by generalizing the argument above, voter χ prefers the candidate with the lower position if and only if $\chi \leq \frac{k_1+k_2}{2}$. Assuming voters vote for their most preferred candidate, and that the candidate with the most votes wins, the election is decided by the median voter $\chi = 0$, whose vote is the same as the best response of the principal described previously.

3.2 Simple delegation

In our game there always exist simple pure strategy equilibria in which the principal always chooses the same expert, independently of two offers, in effect delegating the decision to him. In particular, Theorem 1 implies that if expert i chooses constant markup $k_i = b_i$ and the other expert j chooses constant markup $k_j = -b_i$ then the principal is always indifferent between the two offers, and she might as well always choose expert i. Given this strategy of the principal, expert i's best response is choosing exactly markup b_i, which in expectation implements his ideal action. The other expert has no profitable deviation since his proposal is never accepted. While such an equilibrium exists for each of the two experts, it is more natural to consider the one in which the principal always chooses the expert with the smaller absolute bias, who is expert 2 by convention. These observations are summarized in the next proposition.

Proposition 1. *For $i \in \{1, 2\}$, an equilibrium exists in which the principal always chooses expert i and markups are $k_i = b_i$ and $k_j = -b_i$ for $j \neq i$. The principal's expected payoff in this equilibrium is $\mathbb{E}[-(s + b_i)^2] = -\sigma^2 - b_i^2$, where $s \sim N(0, \sigma^2)$.*

4 Symmetric Biases

In this section we characterize stationary equilibria in pure strategies of the delegation game for the case of symmetrically biased experts: $b_1 = b_2 = b > 0$. Although we analyze general biases in Section 5, we start with the symmetric case for two reasons. First, in many contexts, experts' biases are fairly similar, reflecting similar background and financial interests. Second, we can derive closed form expressions for equilibrium strategies and comparative statics in this case, which helps to provide intuition for the main qualitative results.

4.1 Min Equilibrium

We begin by investigating equilibria, in which the principal always chooses the minimum of the two offers: $\{(k_1, k_2, a \in \arg\min\{a_1, a_2\}) : k_1 + k_2 \geq 0\}$. We call these *min equilibria*. Here and throughout the rest of the paper, let f and F denote the PDF and the CDF of the distribution $N(0, 2\sigma^2)$. We use the shorthand L (or H) to denote the principal's strategy of choosing the lower (or higher) of the two offers.

Using the fact that the expected state given signal realizations (s_1, s_2) is $\frac{s_1+s_2}{2}$, expert 1's expected payoff is

$$U_1(k_1, k_2, L) = \Pr(s_1 + k_1 \leq s_2 + k_2)\mathbb{E}\left[B - \left(s_1 + k_1 - \frac{s_1 + s_2}{2} - b\right)^2 | s_1 + k_1 \leq s_2 + k_2\right] + \Pr(s_1 + k_1 > s_2 + k_2)\mathbb{E}\left[-\left(s_2 + k_2 - \frac{s_1 + s_2}{2} - b\right)^2 | s_1 + k_1 > s_2 + k_2\right] - \sigma^2/2,$$

which he must maximize over choices of k_1. Now consider the effects of a marginal increase in k_1. For cases where expert 1 has a strictly higher offer, there is no change in expert 1's utility since the implemented action of expert 2 remains unchanged. For cases where expert 1 has a strictly lower offer, expert 1 is still selected, but the change in his action influences his quadratic loss. For the marginal case where $s_1 + k_1 \approx s_2 + k_2$, expert 1 goes from being selected to unselected, and thus loses the associated bonus B. However, the quadratic loss term is the same regardless of which expert is chosen, and so causing expert 2 to be chosen at the margin does not influence expert 1's loss term. In equilibrium, expert 1 is best responding, and thus the amount of foregone bonus must be offset by a helpful reduction in quadratic loss. Specifically, expert 1's first order condition is

$$k_1 = b + \left(\sigma^2 - \frac{B}{2}\right) \frac{f(k_1 - k_2)}{1 - F(k_1 - k_2)}. \tag{1}$$

Combining this with the first order condition for expert 2, we obtain a unique equilibrium candidate in which both experts choose a markup $k_m = b + \frac{\sigma^2 - \frac{B}{2}}{\sigma\sqrt{\pi}}$.[14] This expression combines two effects: a *winner's curse* effect and a *bonus-stealing* effect. To isolate the former, suppose for a moment that $B = 0$. Clearly $k_m > b$, that is, each expert's offer is above his optimal offer were he guaranteed to be chosen, as in the case of simple delegation. The reason is that, conditional on being selected, an expert learns that his rival had a sufficiently high signal and the expert would prefer to revise his offer upward. The expert's equilibrium markup must

[14]Here and throughout we use subscript or superscript m to denote min equilibrium; we will use M for max equilibrium, introduced in the next section.

account for this. The bonus-stealing effect acts in the opposite direction and is proportional to the size of the bonus.

As the bonus increases, experts compete more aggressively by reducing their markups. For these markups to remain part of an equilibrium, the principal must be best-responding by choosing the lower offer, which requires that each markup is positive. Equivalently, the bonus must not be too large: $B \leq B_m := 2\sigma^2 + 2\sqrt{\pi}\sigma b$. In what follows, it is useful to define $\rho := \sigma^2 - \frac{B}{2}$.

Proposition 2. *Suppose $b_1 = b_2 = b > 0$. Then a min equilibrium exists if and only if $B \leq B_m$. When it exists, it is unique, and the markups are $k_1^m = k_2^m = k_m = b + \frac{\rho}{\sigma\sqrt{\pi}}$.*

It is worth noting that if $B < 2\sigma^2$, the winner's curse effect outweighs the bonus-stealing effect and $k_m > b$, or in other words, experts make offers that exceed the sum of observed signal and bias (which would be the optimal proposal solely based on own information). In the opposite case the bonus-stealing effect dominates, with $k_m \leq b$.

4.2 Max Equilibrium

If the sum of markups is negative, then the principal's best response is to choose the maximum of the two offers. We now consider *max equilibria*.[15] The forces of the previous section are reversed. Here, an expert's winner's curse is that he would have preferred a lower offer, since being selected (with the higher offer) indicates that the rival expert had a sufficiently low signal. On the other hand, experts compete for the bonus by raising their offers. In a max equilibrium, the experts' offers are $k_M = b - \frac{\sigma^2 - \frac{B}{2}}{\sigma\sqrt{\pi}}$. For the principal to best-respond by choosing the maximum offer, it must be that each markup is nonpositive, which again places an upper bound on the bonus: $B \leq B_M := 2\sigma^2 - 2\sqrt{\pi}\sigma b$.

Proposition 3. *Suppose $b_1 = b_2 = b > 0$. Then a max equilibrium exists if and only if $B \leq B_M$. When it exists, it is unique, and the markups are $k_1^M = k_2^M = k_M = b - \frac{\rho}{\sigma\sqrt{\pi}}$.*

By inspection, the upper bound on the bonus for a max equilibrium is strictly lower than that for a min equilibrium, and in fact, it can be negative. In this case, no max equilibrium exists. The reason is that the assumption $b_1 = b_2 = b > 0$ anchors markups at a positive value, and this works in the same direction as bonus-stealing in the case of a max equilibrium. Since markups must be negative, the sum of these forces must be outweighed by the winner's curse. If b is sufficiently large, then markups would be positive even with zero bonus, which is inconsistent with max equilibrium.

[15] A detailed analysis of max equilibrium with general biases is contained in the supplementary appendix.

4.3 Principal-Optimal Equilibrium

For the case of symmetric biases $b_1 = b_2 > 0$, the min equilibrium is optimal among the pure-strategy, stationary equilibria discussed so far; if one flips the sign of both biases, the ranking between min and max equilibria reverses. To demonstrate this, we make use of the quadratic loss utility form, noting that the principal's expected utility $\mathbb{E}(a-\theta)^2$ (net of the bonus paid, and conditional on θ) is equal to $-\bar{b}^2 - Var$ where $\bar{b} := \mathbb{E}(a-\theta)$ is the expected bias of the chosen offer and $Var := \text{Var}(a-\theta)$ is the variance of the bias. For the equilibria discussed, we have

- simple delegation: $\bar{b} = b,\ Var = \sigma^2$
- min equilibrium: $\bar{b} = b - \frac{B}{2\sqrt{\pi}\sigma},\ Var = (1-\frac{1}{\pi})\sigma^2$
- max equilibrium: $\bar{b} = b + \frac{B}{2\sqrt{\pi}\sigma},\ Var = (1-\frac{1}{\pi})\sigma^2$.

Comparing min and max equilibria, we see that the variance of the bias is the same due to symmetry, while $\bar{b}$ has smaller magnitude for the min equilibrium because $b > 0$. Intuitively, when $b > 0$, the principal's ability to select the lower of the two offers offsets the inherent bias of the agents, whereas selecting the higher offer does not. Comparing simple delegation with the min equilibrium, we see that Var is lower for the latter. The reason is that the inclusion of the second expert primarily is binding when he has the lower offer, which is more likely when the first expert's signal is high. This truncates the principal's loss in such situations and reduces variance. On the other hand, expected bias can have larger magnitude in the min equilibrium than under simple delegation, as can be seen easily in the case of $b = 0$, but the condition $B \leq B_m$ for the existence of a min equilibrium ensures any such disadvantage is outweighed by the variance reduction. These comparisons show that the min equilibrium is principal-optimal among pure-strategy, stationary equilibria.

Proposition 4. *Suppose $b_1 = b_2 = b > 0$. Then whenever the min equilibrium exists, the principal prefers it to simple delegation. If, in addition, the max equilibrium also exists, the principal prefers the min equilibrium to the max equilibrium.*

In Section 6.1, we show further that in the case of equal biases and zero bonus, the min equilibrium is optimal if the principal is given commitment power and has the ability to choose a mixed strategy. In particular, commitment by the principal results in the same outcome as in the min equilibrium of the game without commitment.

4.4 Comparative Statics

Given that for symmetric biases the min equilibrium is optimal among pure strategy stationary equilibria, and optimal under commitment for a large range of bonus values, from here on we investigate comparative statics of the principal's expected payoff in the min equilibrium. We highlight three nonmonotonicities with respect to input parameters: B, b, and σ.

First, the principal's payoff is nonmonotonic in the bonus B, even when this bonus is provided externally (not from the principal's pocket). Formally, suppose that the game is preceded by a stage in which the principal publicly commits to a bonus level B, and that (for now) the principal's payoff is still $-(a-\theta)^2$. While the variance of the action $Var(k_m, k_m, L) = \sigma^2 - \frac{\sigma^2}{\pi}$ does not depend on B, the expected bias $b_m = b - \frac{B}{2\sqrt{\pi}\sigma}$ is decreasing in B. The principal prefers the expected bias to be as close to 0 as possible, so her expected payoff in the min equilibrium is maximized at $B = 2\sqrt{\pi}\sigma b$, where it is equal to $-\sigma^2[1-\frac{1}{\pi}]$. At $B = 0$, $\bar{b} = b_m := b - \frac{B}{2\sqrt{\pi}\sigma} = b \geq 0$ and a small increase in bonus decreases the experts' markup and benefits the principal. However, at $B = B_m$, $b_m = b - \frac{B}{2\sqrt{\pi}\sigma} = -\frac{\sigma}{\sqrt{\pi}} < 0$ and principal prefers to increase markups and correspondingly lower the bonus. As a consequence, an intermediate point $B = 2\sqrt{\pi}\sigma b$ is optimal. Figure 2 below illustrates the relationship between the expected bias and the principal's payoff.

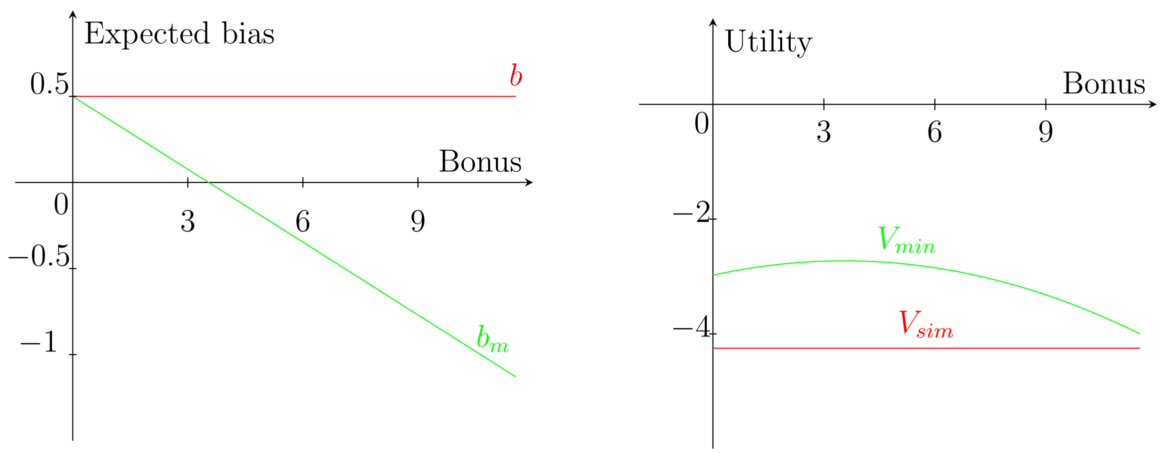

Figure 2: The expected bias and the principal's payoff in the min equilibrium as functions of B (exogenous sources case) for $b = 0.5$ and $\sigma = 2$, compared to simple delegation.

The principal's payoff is also non-monotonic in the bonus B when it is paid by the principal. While increasing the bonus has a direct cost to the principal, it increases the bonus-stealing incentive which reduces markups. This reduces the magnitude of the expected bias when the expected bias is positive. A simple condition determines when a positive B is optimal. Recall that the principal's payoff is $-B - b_m^2 - Var$. Starting from $B = 0$, an increase in B reduces b_m^2 at a rate $\frac{b}{\sqrt{\pi}\sigma}$, which is increasing in b due to quadratic loss and which is decreasing in σ since the bonus-stealing incentive is weaker for larger σ. The net change in the principal's payoff is positive if and only if $\frac{b}{\sqrt{\pi}\sigma} > 1$.

The relationship between the principal's payoff and the bonus is illustrated in Figure 3. The principal chooses a strictly positive bonus if the decrease in expected bias exceeds the marginal disutility from increasing bonus: $\frac{b}{\sqrt{\pi}\sigma} > 1$. This holds for $b = 2$, $\sigma = 0.5$, but not for $b = 1$, $\sigma = 1$.

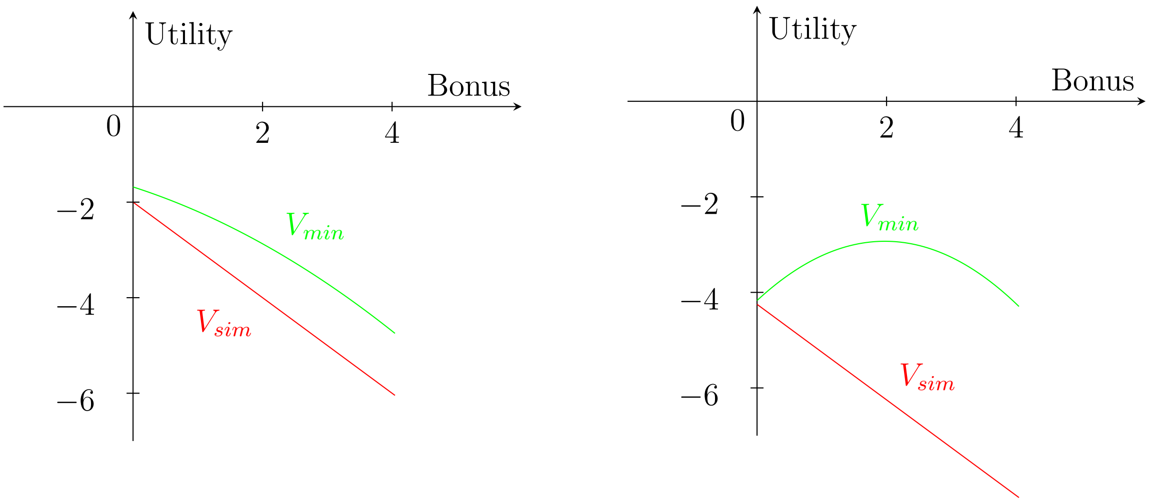

Figure 3: The principal's payoff in the min equilibrium and simple delegation as functions of the bonus (paid by principal) for $b = 1$, $\sigma = 1$ on the left and for $b = 2$, $\sigma = 0.5$ on the right.

The principal's payoff is also nonmonotonic in the common bias b of the experts and in the noise level σ in the experts' signals. If the bonus is positive and b is small, then the expected bias of the chosen offer is negative, and increasing b in this case brings it closer to zero; but once the expected bias is zero, further increases in b hurt the principal. The nonmonotonicity with respect to σ stems from two effects: increasing σ unambiguously increases the variance of the chosen offer but it can help the principal by reducing the magnitude of the expected bias when the expected bias is negative.

Proposition 5 summarizes the various nonmonotonicities discussed above.

Proposition 5. *Suppose $b_1 = b_2 = b > 0$. The principal's expected payoff in the min equilibrium is nonmonotonic in B, both when it is paid from exogenous sources and when it is paid by the principal. It is also nonmonotonic in b and σ.*

5 General Biases

Many of the economic insights obtained in the case of symmetric biases extend to general biases. The characterization of min equilibrium is contained in Section A.1, and that of max equilibrium is contained in the supplementary appendix. In this section, we compare the principal's payoffs across equilibria and discuss the underlying forces, and then we compare the principal's payoffs for equal versus opposite biases of the experts.

5.1 Principal-Optimal Equilibrium

In this section we compare the principal's expected utility in the min and max equilibria, and in the case of simple delegation. We also investigate how the principal's expected payoff in equilibrium depends on the biases of the experts. The main finding is that the min equilibrium remains always preferred by the principal to the max equilibrium, and for a large range of parameter values (including the cases of equal biases, opposite biases, or one bias equal to zero) the min equilibrium outperforms simple delegation as well.

First we establish that whenever both min and max equilibria exist, the principal always prefers the former.

Proposition 6. *For any fixed $B \geq 0$, the principal prefers the min equilibrium to the max equilibrium whenever both exist.*

The intuition for the above result can be summarized as follows. Let $z^* \geq 0$ denote the difference between the experts' markups. Now z^* is the same in both the min and max equilibria, so in turn the variance of the chosen offer, $\text{Var}(a - \theta)$, is the same. However, in the max equilibrium, expert 1's offer is chosen with probability $F(z^*)$, which is weakly higher than $1 - F(z^*)$ in the min equilibrium; this, along with the desire for this shifts the expected bias higher. Furthermore, in the max equilibrium the bonus motivates experts to increase their markups, as opposed to min equilibrium in which the bonus motivates experts to decrease their markups: the expected biases b_M and b_m in max and min equilibria, respectively, satisfy $b_M = b_1 F(z^*) + b_2(1 - F(z^*)) + Bf(z^*) \geq b_1(1 - F(z^*)) + b_2 F(z^*) - Bf(z^*) = b_m$. Hence, to conclude that the principal is better off in the min equilibrium it is enough to show that $|b_M| \geq |b_m|$ or, taking into account the above, $b_m + b_M \geq 0$. By inspection, $b_m + b_M = b_1 + b_2 \geq 0$, implying that the principal is better off in the min equilibrium.

We can further argue that in case the two experts are not equally biased, the expert with the lower bias also prefers min equilibrium to max equilibrium, while the expert with the higher bias has the opposite preferences. This is both because the expected action is closer to expert 2's ideal point in the min equilibrium, and closer to expert 1's ideal point in max equilibrium, and because expert 2 is chosen (and hence receives the bonus) with higher probability in the min equilibrium, while expert 1 is chosen with higher probability in the max equilibrium.

Given Proposition 6, we now compare the principal's utility in the min equilibrium and simple delegation to the expert with the smaller absolute bias. In general this comparison is complicated, but in the next proposition we show that if one fixes the ratio of biases to be positive and less than unity, then for sufficiently large scale of bias, the principal prefers simple delegation. On the other hand, when the experts are equally biased, or one of the experts is unbiased, or biases have the opposite sign, the principal prefers the min equilibrium to simple delegation. The principal also prefers the min equilibrium when the noise in the experts' signals becomes large. Figure 4 illustrates the comparison between simple delegation and min equilibrium.

Proposition 7. *Assume $B = 0$, and parameterize biases as $b_1 = b$ and $b_2 = rb$ for some $b \in [0, \infty)$, $r \in [-1, 1]$. For each $r \in (0, 1)$, there exists a threshold $b^* > 0$ such that the principal prefers simple delegation to the min equilibrium if $b > b^*$. For each $r \in [-1, 0] \cup \{1\}$, the principal prefers min equilibrium for all $b \in [0, \infty)$. For any fixed biases, if σ is sufficiently large, the principal prefers the min equilibrium.*

In particular, Proposition 7 says that when biases are equal or opposites or when one of the biases is zero, the principal prefers the min equilibrium. Note that the equal biases case was treated in Proposition 4, where we showed that for any common bias b, and any bonus B such that the min equilibrium exists, the principal prefers the min equilibrium to simple delegation.

For intuition behind the result above, it helps to decompose the effects of min equilibrium (relative to simple delegation) into bias-shifting and variance-reduction effects and examine how these effects depend on parameter values. First consider equal biases ($r = 1$). In this case, the expected bias in either the min equilibrium or under simple delegation is exactly b, so the bias-shifting effect is zero, while the (conditional) variance of the chosen offer is much lower in the min equilibrium; hence the principal unambiguously prefers the min equilibrium. Next consider unequal biases $b_1 > b_2$. Due to the winner's curse in the min equilibrium and the fact that the expert's markups are above their biases, the expected bias is strictly above b_2, while the expected bias under simple delegation is exactly b_2. Hence, if both biases are positive and their difference is large but b_2 is not too close to zero, min equilibrium features a larger expected bias than simple delegation; however, since expert 2 is still chosen most of the time, the variance of the chosen offer is not much smaller in min equilibrium than under simple delegation; the bias-shifting effect dominates and simple delegation is preferred. But if b_2 is sufficiently close to zero, the increase in expected bias in min equilibrium has very little effect on the principal's payoff due to the concavity of the loss function — the variance-reduction benefit of min equilibrium dominates. Finally if $b_2 < 0$ and b_2 is not too close to 0, then the bias-shifting effect helps the principal (in addition to the variance-reduction effect): higher equilibrium markups in min equilibrium mean that the expected bias is closer to 0 than under simple delegation.

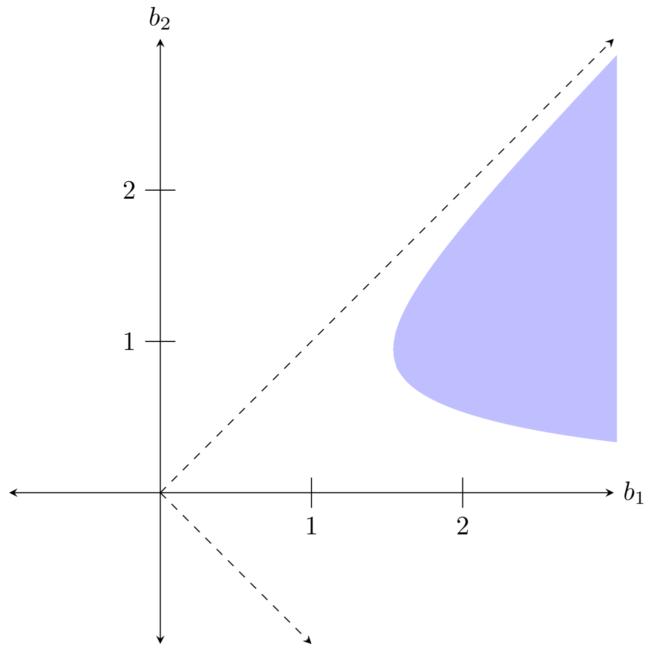

Figure 4: Depending on the biases (b_1, b_2), the principal-optimal equilibrium is either simple delegation to expert 2 (shaded region) or choosing the lower offer (unshaded). Parameter settings are $B = 0$ and $\sigma = 1$.

Since the principal is able to (partially) utilize two signals when she has access to two experts, one might wonder whether the advantage of having a second expert is simply that there is more information available. To disentangle the role of competition from the role of better information when the principal has access to two experts, we now compare the principal's payoff in the min equilibrium to that of simple delegation to a single expert ("super expert") who receives *two* signals. For simplicity, assume both experts have the same bias $b > 0$. We show that it is still possible for the min equilibrium to be preferred, provided the bias is sufficiently large (relative to the noise in experts' signals) and the bonus level is intermediate. Intuitively, while the min equilibrium utilizes information less effectively than simple delegation, as measured by the variance terms $\left(1 - \frac{1}{\pi}\right)\sigma^2 > \frac{\sigma^2}{2}$ in the principal's payoff, the combination of a positive bonus and competition between experts reduces the expected bias in the min equilibrium, and the maximum benefit of this reduction is precisely b^2. The bonus, on the other hand, cannot be too large, otherwise competition is so strong that the expected bias in the min equilibrium is negative with a magnitude that is increasing in the bonus.

Proposition 8. *For each $b > \sigma\sqrt{\frac{1}{2} - \frac{1}{\pi}}$, there exists a nonempty interval of bonus levels $(\underline{B}, \overline{B}) \subset (0, +\infty)$ such that the principal strictly prefers competition between two experts relative to simple delegation to a super expert if and only if $B \in (\underline{B}, \overline{B})$. If $b < \sigma\sqrt{\frac{1}{2} - \frac{1}{\pi}}$, simple delegation to a super expert is strictly preferred.*

5.2 Equal vs. Opposite Biases

Here we compare the expected payoff the principal can achieve with two equally biased experts, $b_1 = b_2 = b$, to the expected payoff she can achieve with two oppositely biased experts, $b_1 = -b_2 = b$. Theorem A.1 implies that the expected bias of the action in the case of equally biased experts is equal to $b - Bf(0)$. In the case of oppositely biased experts the expected bias of the action is $b(1 - 2F(z^*)) - Bf(z^*)$, simplifying to $b(1 - 2F(z^*))$ when $B = 0$. Hence, with $B = 0$ the absolute value of the expected bias in the case of oppositely biased experts is lower than in the case of equally biased experts (where it is b). However, the next proposition shows that the variance of the action is lower in the symmetric case, and in fact this effect dominates, resulting in the principal preferring to have two equally biased experts. Let $V_{symm}(b)$ be the principal's expected payoff in the min equilibrium when $b_1 = b_2 = b$, let $V_{opp}(b)$ be the principal's expected payoff in the min equilibrium when $b_1 = -b_2 = b$, and let $V_{sim}(b) = -\sigma^2 - b^2$ be the principal's expected payoff in the case of simple delegation to an expert with absolute bias b.

Proposition 9. *For any $b > 0$, $V_{symm}(b) \geq V_{opp}(b) \geq V_{sim}(b)$ whenever the min equilibrium exists for both symmetric and opposite biases.*

6 Extensions

In this section, we consider several extensions of the model.

6.1 Strengthening or Relaxing Commitment

Motivated by applications, a key ingredient of our model is that the principal can only select one of the two actions proposed by the experts, but she cannot ex ante commit to which one. A naturally arising question is how the conclusions would change in two models in which we change this assumption. The first model keeps the assumption that the principal can only choose between two alternatives, but allows for ex ante commitment on her part. The second one departs from the principal having to choose between the two proposals and allows her to choose any action after observing the proposals, meaning that the proposals become cheap talk messages.

We first extend the model to allow the principal to publicly commit to a mixed strategy. In particular, suppose that the principal can credibly commit to randomizing with probabilities (p_m, p_M, p_1, p_2) over the following four pure strategies, respectively: always choose the lowest offer, always choose the highest offer, always choose expert 1, always choose expert 2. For simplicity, we consider only the case $B = 0$.

We show that when experts have equal biases $b > 0$, commitment power does not help the principal. First, in any equilibrium, the experts' markups are such that the bias of their action conditional on being chosen is exactly b, and hence from the principal's ex ante perspective, the bias of the chosen offer is b. Second, we show that for any constant markup strategies of the experts, the variance of the chosen offer is at least $\sigma^2(1-1/\pi)$, the variance of the chosen offer in the min equilibrium. Thus, the principal can do no better than to commit to choosing the lowest offer. But commitment power is not necessary, since the min equilibrium always exists for $B = 0$ and $b > 0$.

On the other hand, in some instances, the principal benefits from introducing an "element of surprise" and optimally commits to a mixed strategy. To illustrate this point, it suffices to consider oppositely biased experts and randomization by choosing the lowest offer with probability p and the highest offer with probability $1 - p$. As the common magnitude b of the biases increases, expert 1's winner's curse in the min equilibrium becomes more severe, and his markup very large. By introducing a small probability of choosing the higher offer, expert 1 is incentivized to reduce his markup, benefiting the principal. The cost to the principal of doing so is in choosing the wrong offer. When b is large, the markup-reduction benefit is large and outweighs the cost, and thus the principal can profitably deviate from $p = 1$ to an interior p. As stated in Proposition 10, this result generalizes beyond biases which are exact opposites: the principal benefits from commitment to mixing if the biases have opposite sign and they are scaled by a sufficiently large common factor.

Proposition 10. *Assume $B = 0$. If $b_1 = b_2 = b > 0$, then the principal weakly prefers the min equilibrium to commitment to any vector of probabilities (p_m, p_M, p_1, p_2). If $(b_1, b_2) = (b, rb)$ for some $r \in [-1, 0)$, then there exists a threshold $b^* > 0$ depending on r such that if $b > b^*$, the principal strictly benefits from committing to a mixture between choosing the lowest offer and choosing the highest offer, relative to commitment to a pure strategy.*

Next, consider a model in which the proposals are only cheap talk messages and the principal is allowed to choose any action after observing a pair of proposals. In this model it is difficult to make comparisons between having access to one versus two experts, for multiple reasons. One is that not much is known about equilibria in a one-sender, one-receiver cheap talk game with the state space being the whole real line (an assumption of our main model that we make for tractability). Some of the results from Filipovich (2008), investigating a cheap talk model with the state space being a circle, suggest that there might not be any informative equilibria when the state space is the real line.[16] There are more results in the literature about two-sender cheap talk games with noisy state observations when the state space is the whole real line, but mainly regarding the limit as the noise in state observations vanishes. In particular Ambrus and Lu (2014) show that as the noise vanishes, there is a sequence of equilibria converging to full revelation of information. To summarize, the existing literature in cheap talk suggests that with imperfect state observations, the presence of a second expert should benefit the principal, just like in our context, although we are not aware of any paper that makes this point explicitly.

6.2 Unselected Expert Indifferent over Actions

In the baseline model we assumed that an expert whose offer is not selected is still affected by principal's action. While this is a reasonable assumption in some contexts, in other situations it is more realistic to assume that the expert not selected by the principal receives an outside payoff that is independent of the state and the implemented action. For instance, a car mechanic is unlikely to care what kind of maintenance is done if he is not the one selected for the job. In this extension we assume that if expert i is chosen then expert j's realized payoff is normalized to be 0. We restrict attention the case of equally biased experts: $b_1 = b_2 = b > 0$.

In this version of the model we assume that B is large enough ($B \geq \sigma^2$) that an expert's expected payoff under simple delegation to that expert is nonnegative; that is, he prefers simple delegation to not being selected at all. Under this condition, the same simple delegation equilibria exist in this version of the model as in the baseline model. Below we show that under some conditions there also exist symmetric pure strategy equilibria that are similar to the ones characterized in the baseline model. For this to be the case, the bonus payment must be neither too small nor too large.

Using the same notation as before, we investigate strategy profiles $\{(k_1, k_2, C(a_1, a_2) \in \arg\min\{a_1, a_2\}) : k_1 + k_2 \geq 0\}$. While for any such profile the principal's payoff does not change, the experts' expected payoffs must be be recalculated.

[16] In Filipovich (2008) the sender observes the state without noise, but with only one sender the presence of noise in the sender's observation does not change the substance of the analysis.

Notice that for a fixed constant markup strategy of the other expert, an expert can choose arbitrarily high constant markup and guarantee an expected payoff arbitrarily close to 0.[17] Hence, 0 is a lower bound for experts' equilibrium payoffs.

In what follows, define $\beta := \sqrt{\pi + \frac{B}{\sigma^2} - \frac{5}{2}}$, $B_1 := \left(\frac{5}{2} - \frac{3\pi(8\pi-11)}{16(\pi-1)^2}\right)\sigma^2$ and $B_2 := \frac{5}{2}\sigma^2 + 2\sqrt{\pi}b\sigma + b^2$.

Proposition 11. *A symmetric min equilibrium $k_1^m = k_2^m = k_m$ exists if and only if $B \in [B_1, B_2]$. When it exists, it is unique and characterized by $k_1^m = k_2^m = k_m = b + (\sqrt{\pi} - \beta)\,\sigma$.*

In the Appendix we show that $k_m = b + (\sqrt{\pi} - \beta)\sigma > k_m^{bas.} = b + (1 - \frac{B}{2\sigma^2})\frac{\sigma}{\sqrt{\pi}}$, hence in this version of the model experts select higher markups in the min equilibrium than in the baseline model (for parameter values for which min equilibrium exists in both model versions). The intuition behind this result is that in this alternative version of the model, the relative gain from being selected is reduced by the policy loss (that is not imposed on the expert if not selected). Since we consider B sufficiently large that expected payoffs are nonnegative, the resulting "net bonus" is still nonnegative; being selected is still preferable, conditional on having made the lower offer. It follows that an expert's equilibrium offer in either version is lower than what is ex-post optimal for that expert — that is, optimal after conditioning on both the expert's signal and having the lower offer. The smaller net bonus in the alternative version reduces the expert's incentive to marginally lower his offer in order to more frequently earn the net bonus. This reduction must be met by an offsetting reduction in his incentive to raise his offer, which is enforced through his bidding higher and thus closer to his ex-post optimum; due to quadratic losses, marginal movements toward the ex-post optimum have decreasing marginal benefits.

We note that the qualitative comparison between this extension and the baseline model is dependent upon the modeling of preferences over policy outcomes through losses. Such a model is appropriate for applications where an expert would prefer not to be associated with the project if his action would be sufficiently far from the true state; for example, this would be the case if the expert has a reputation at stake. Alternatively, one could model preferences through gains, using some single-peaked, nonnegative utility function of the distance between the action and the true state. In that model, the comparison above would be reversed, as being selected enhances the bonus and thus experts compete more aggressively by lowering their offers.

[17] For this reason, we do not introduce an explicit participation constraint in this version of the model, even though such a constraint would be natural in many applications.

6.3 Correlated signals

We have assumed that the expert's signals are conditionally i.i.d. given the true state θ — the signal errors are independent. However, the existence and uniqueness of min equilibrium is robust to correlation — positive or negative — in these errors.

Suppose the signal pair (s_1, s_2) is multivariate normal with mean (θ, θ) and covariance matrix $\Sigma = \begin{pmatrix} \sigma^2 & p\sigma^2 \\ p\sigma^2 & \sigma^2 \end{pmatrix}$, where $p \in [-1, 1]$ is the correlation in the error terms $\epsilon_i = s_i - \theta$. Setting $p = 0$ nests the baseline model.

For simplicity, suppose $B = 0$ and the experts are symmetric. Proposition 12 shows that there is a unique min equilibrium, and the common markup can be characterized in closed form. By inspection, the markup is decreasing in p. Intuitively, as p increases, the winner's curse in the min equilibrium becomes less severe: since the other expert's signal carries less additional information beyond an expert's own signal, there is less upward belief revision upon being selected (and thereby learning that the other expert's action was higher). In the limit as $p \to 1$, the winner's curse vanishes, and the markup converges to the bias. The variance of the chosen offer remains lower than under simple delegation, while the expected bias of the chosen offer is b in either case; hence, the principal continues to prefer the min equilibrium to simple delegation.

Proposition 12. *Suppose $B = 0$ and experts have a common bias $b \geq 0$. For each $p \in [-1, 1]$, there exists a unique min equilibrium, and it is characterized by a common markup $k^* = b + \sigma\sqrt{(1-p)/\pi}$. The principal prefers this min equilibrium to simple delegation.*

6.4 N Experts

Our model gets analytically difficult when there are more than two experts. But we can show that if the bonus is zero and experts are symmetric, then for any number N of experts, a min equilibrium exists, as long as the principal can only observe ordinal information, that is the ordering of the N proposals (as in Kamenica (2008)).[18] In fact, in this case it is possible to analytically characterize this equilibrium. Intuitively, the fact that a particular expert is chosen implies that *all* other experts' signals were higher, and this implies that the winner's curse is more severe when there are more experts. Thus, in the symmetric min equilibrium, the common markup is increasing in the number of experts.

Proposition 13. *Suppose there are $N \geq 2$ experts with equal biases $b > 0$, and set $B = 0$. There exists a unique symmetric min equilibrium, and it is characterized by a common markup $k^* = b + \frac{N-1}{N}\frac{2\sigma}{\sqrt{\pi}}$. The principal prefers this min equilibrium to simple delegation.*

[18] By continuity it follows that a min equilibrium also exists when the bonus is positive but small enough.

The assumption that the principal can only observe the ordinal ranking of the proposals is used in that we do not allow the principal to condition her decision, of whether to choose the lowest proposal or one of the other proposals, on the exact vector of proposals. That is, we only require the principal to prefer the lowest proposal to any other proposal "in expectation." Allowing the principal to understand cardinal information would require considering more complicated equilibria with more than two experts. For example if the lowest proposal were a clear outlier (much lower than all other proposals), it could suggest to the principal that the expert making this proposal observed the state with a negative noise term of large magnitude, making her prefer the second lowest proposal. Characterization of equilibria in such settings in which the principal pays attention to all experts becomes very challenging.

6.5 Relaxing the Functional Form Assumptions

We have exploited the tractability of quadratic loss functions with Gaussian signals; in this section, we show that the existence of min equilibrium can be extended beyond these functional forms. For simplicity, we assume $B = 0$.

Let $\ell : \mathbb{R} \to \mathbb{R}$ denote the loss function for each player, taking as its argument $a - \theta$ for the principal and $a - \theta - b_i$ for expert i; the baseline model corresponds to the special case $\ell(x) = -x^2$. Assume that ℓ is symmetric about 0, strictly concave and continuously differentiable. Let the experts observe private signals $s_i = \theta + \epsilon_i$, where the ϵ_i are i.i.d. with density g. Assume that g is symmetric about 0 and has full support on $\mathbb{R}$. Finally, we require that experts' feasible markups are uniformly bounded: there is some $M \in (0, \infty)$ such that $b_1, b_2 \in (-M, M)$, and for each $i \in \{1, 2\}$ and all s_i, each expert's action must lie in $[s_i - M, s_i + M]$.

As the following proposition shows, the existence of min equilibrium extends from the baseline model, and markups remain higher than biases due to the winner's curse.

Proposition 14. *There exists a min equilibrium, and it satisfies $k_i > b_i$ for $i = 1, 2$.*

7 Conclusion

We have proposed a model in which a principal can choose between two imperfectly informed experts, introducing the possibility of competition in a delegation framework. We have showed that a principal with limited knowledge of the decision environment can benefit from the presence of two experts, relative to a simple unconstrained delegation to one of them, even if the experts have exactly the same bias. The main reason is that in equilibria in which the selection of the expert depends nontrivially on the experts' proposals, information is utilized from both experts' private signals. The option of offering a bonus payment to the selected expert can improve the principal's payoff, by inducing the experts to report more truthfully, but only to a certain point. Lastly, committing with a small probability to choose the (in expectation) inferior proposal can benefit the principal.

As this is the first step in investigating the benefits of multiple choices of experts in a delegation problem, there are many avenues of future research. One is examining multi-dimensional environments, in which different experts differ in their dimensions of specialization. Another direction would be investigating the problem of choosing an expert to whom to delegate a task using a more general mechanism design approach.

A Appendix

This appendix is arranged as follows. In Section A.1, we characterize min equilibria for general biases; analogous results for max equilibria are presented and proved in the supplementary appendix. In Section A.2, we prove the results of Section A.1 and Section 5. In Section A.3, we prove the results of Section 4 as special cases of these, and in Section A.4 we prove the results of Section 6.

A.1 Min Equilibria for General Biases

Recall the definition $\rho := \sigma^2 - \frac{B}{2}$. In min equilibrium, the first order conditions for the players are

$$k_1 = b_1 + \rho \frac{f(k_1 - k_2)}{1 - F(k_1 - k_2)} \tag{2}$$

$$k_2 = b_2 + \rho \frac{f(k_1 - k_2)}{F(k_1 - k_2)}. \tag{3}$$

After defining $z := k_1 - k_2$ and functions $v(z) := \frac{f(z)}{1-F(z)}$ and $w(z) := \frac{f(z)}{F(z)}$, subtracting these equations yields

$$z - \rho\,[v(z) - w(z)] = b_1 - b_2. \tag{4}$$

In the proof of Theorem A.1, there is a unique solution to (4) which we denote z^*.

Theorem A.1. *There exists a threshold $B_m > 0$ such that a min equilibrium exists if and only if $B \leq B_m$. When it exists, it is unique and characterized by markups $k_1^m = b_1 + \rho v(z^*)$ and $k_2^m = b_2 + \rho w(z^*)$ with $k_1^m - k_2^m = z^* \geq 0$. Moreover, B_m lies in the interval $[2\sigma^2 + 2\sqrt{\pi}\sigma \max(0, b_2), 2\sigma^2 + \sqrt{\pi}\sigma(b_1 + b_2)]$. For $B \leq B_m$,*

- $z^* \geq b_1 - b_2 \iff B \leq 2\sigma^2 \iff \rho \geq 0$;
- $\bar{b}(k_1^m, k_2^m, L) = b_m := b_1(1 - F(z^*)) + b_2 F(z^*) - Bf(z^*)$;
- $Var(k_1^m, k_2^m, L) = \sigma^2 - 4\sigma^4 f^2(z^*) - 2\sigma^2 z^* f(z^*)(2F(z^*) - 1) + (z^*)^2 F(z^*)(1 - F(z^*))$.

Lemma A.1 provides key quantities given arbitrary (constant markup) expert strategies when the principal chooses the min of the two offers.

Lemma A.1. *If both experts follow constant markup strategies $a_j(s_j) = s_j + k_j$ and the principal always chooses the lower offer, then*

$$
\begin{aligned}
\bar{b}(k_1, k_2, L) &= -2\sigma^2 f(z) + k_1(1 - F(z)) + k_2 F(z); \\
Var(k_1, k_2, L) &= \sigma^2 - 4\sigma^4 f^2(z) - 2\sigma^2 z f(z)(2F(z) - 1) + z^2 F(z)(1 - F(z)); \\
V(k_1, k_2, L) &= -\sigma^2 + 2(k_1 + k_2)\sigma^2 f(z) - k_1^2(1 - F(z)) - k_2^2 F(z); \\
U_i(k_1, k_2, L) &= -\sigma^2 + 2\sigma^2(k_i + k_j - 2b_i)f(z) - (k_j - b_i)^2 F(k_i - k_j) - (k_i - b_i)^2 F(k_j - k_i) \\
&\quad + B(1 - F(k_i - k_j)).
\end{aligned}
$$

Notice that the equilibrium markup difference z^* as well as $\mathrm{Var}(a - \theta)$ depend on the biases of the experts only through $b_1 - b_2$. In Corollary A.1 of Section A.2, we give the full expansion of the players' utilities in the min equilibrium.

A.2 Proofs for Section A.1 and Section 5

We prove the general results of Section A.1 and Section 5, after which many of the results of Section 4 can be handled readily as special cases.

We first provide an auxiliary lemma.

Lemma A.2. *Let $\epsilon_1, \epsilon_2 \sim N(0, \sigma^2)$, where ϵ_1 and ϵ_2 are independent, and define* $\xi(k_1, k_2) := \min(\epsilon_1 + k_1, \epsilon_2 + k_2)$, $\eta(k_1, k_2) := \max(\epsilon_1 + k_1, \epsilon_2 + k_2)$. *Then*

$$
\begin{aligned}
\mathbb{E}\xi(k_1, k_2) &= -2\sigma^2 f(k_1 - k_2) + k_1(1 - F(k_1 - k_2)) + k_2 F(k_1 - k_2); \\
\mathbb{E}\eta(k_1, k_2) &= 2\sigma^2 f(k_1 - k_2) + k_1 F(k_1 - k_2) + k_2(1 - F(k_1 - k_2)); \\
\mathbb{E}\xi^2(k_1, k_2) &= \sigma^2 - 2(k_1 + k_2)\sigma^2 f(k_1 - k_2) + k_1^2(1 - F(k_1 - k_2)) + k_2^2 F(k_1 - k_2); \\
\mathbb{E}\eta^2(k_1, k_2) &= \sigma^2 + 2(k_1 + k_2)\sigma^2 f(k_1 - k_2) + k_1^2 F(k_1 - k_2) + k_2^2(1 - F(k_1 - k_2)).
\end{aligned}
$$

Proof. These expressions follow from Nadarajah and Kotz (2008), setting the correlation coefficient to 0. □

Proof of Lemma A.1. After observing the signal s_i, expert i does a Bayesian update of his beliefs: $\theta|s_i \sim N(s_i, \sigma^2)$ and $s_j|s_i \sim N(s_i, 2\sigma^2)$. Since the principal chooses the lower offer, she accepts a_i iff $s_j > s_i + k_i - k_j$. Denote by g the PDF of $N(s_i, 2\sigma^2)$. Hence, the expected utility of expert i is

$$\begin{aligned}U_i(k_1, k_2, L) = &\int_{s_i+k_i-k_j}^{\infty} \mathbb{E}\left[B - (a_i - \theta - b_i)^2 | s_i, s_j\right] g(s_j)\, ds_j \\ &+ \int_{-\infty}^{s_i+k_i-k_j} \mathbb{E}\left[-(a_j - \theta - b_i)^2 | s_i, s_j\right] g(s_j)\, ds_j.\end{aligned}$$

As $(\theta|s_i, s_j) \sim N(\frac{s_j+s_j}{2}, \frac{\sigma^2}{2})$, $a_i = s_i + k_i$, $a_j = s_j + k_j$, we obtain

$$\begin{aligned}U_i(k_1, k_2, L) = &\int_{s_i+k_i-k_j}^{\infty} \left[B - \left(k_i - b_i - \frac{s_j - s_i}{2}\right)^2 - \frac{\sigma^2}{2}\right] g(s_j)\, ds_j \\ &+ \int_{-\infty}^{s_i+k_i-k_j} \left[-\left(k_j - b_i + \frac{s_j - s_i}{2}\right)^2 - \frac{\sigma^2}{2}\right] g(s_j)\, ds_j.\end{aligned}$$

Now make a substitution $t = s_j - s_i$ and denote by f and F the PDF and CDF of $N(0, 2\sigma^2)$.

$$U_i(k_1, k_2, L) = \int_{k_i-k_j}^{\infty} \left[B - \left(k_i - b_i - \frac{t}{2}\right)^2 - \frac{\sigma^2}{2}\right] f(t)\, dt + \int_{-\infty}^{k_i-k_j} \left[-\left(k_j - b_i + \frac{t}{2}\right)^2 - \frac{\sigma^2}{2}\right] f(t)\, dt.$$

Note that $U_i(k_1, k_2, L)$ does not depend on signal s_i, which is intuitive for the improper prior. As $\int_a^{\infty} t f(t)\, dt = 2\sigma^2 f(a)$ and $\int_{-\infty}^{\infty} t^2 f(t)\, dt = 2\sigma^2$, we get the expression for $U_i(k_i, k_j, L)$.

Now in state θ, the principal's action a is distributed as $\theta + \xi$, where $\xi = \min(\epsilon_1 + k_1, \epsilon_2 + k_2)$; $\epsilon_1, \epsilon_2 \sim N(0, \sigma^2)$, ϵ_1 and ϵ_2 are independent.

Therefore, from Lemma A.2 the expected bias of the accepted offer is

$$\bar{b}(k_1, k_2, L) = \mathbb{E}\xi(k_1, k_2) = -2\sigma^2 f(k_1 - k_2) + k_2 F(k_1 - k_2) + k_1(1 - F(k_1 - k_2))$$

and the expected utility of the principal is

$$\begin{aligned}V(k_1, k_2, L) &= -\mathbb{E}(a - \theta)^2 = -\mathbb{E}(\theta + \xi - \theta)^2 = -\mathbb{E}\xi^2(k_1, k_2) \\ &= -\sigma^2 + 2(k_1 + k_2)\sigma^2 f(k_1 - k_2) - k_2^2 - (k_1^2 - k_2^2)(1 - F(k_1 - k_2)).\end{aligned}$$

Finally, the variance of the chosen offer is

$$\begin{aligned}Var(k_1, k_2, L) =& - V(k_1, k_2, L) - \bar{b}^2(k_1, k_2, L)\\ =&\sigma^2 - 4\sigma^4 f^2(z) - 2\sigma^2 z f(z)(2F(z) - 1) + z^2 F(z)(1 - F(z)).\end{aligned}$$

□

The following lemma provides several useful bounds. We only prove the final statements, as the rest are immediate corollaries of Sampford (1953) after adjusting for the variance $2\sigma^2$.

Lemma A.3. *The following inequalities hold for all $x \in \mathbb{R}$:*

- $0 < v'(x) < \frac{1}{2\sigma^2}$*;*
- $0 > w'(x) > -\frac{1}{2\sigma^2}$*;*
- $v''(x) > 0$.

Moreover, $v(x)w(x)$ is decreasing in x^2, $\lim_{x\to\pm\infty} v(x)w(x) = 0$, and $\lim_{x\to\infty} x/v(x) = 2\sigma^2$.

Proof. Only the limit claims require proof. For the first, by symmetry, consider $x \to +\infty$. We have $v(x)w(x) = \frac{f(x)^2}{F(x)(1-F(x))}$ which has limiting form $0/0$, and by L'Hopital's rule, the limit is $\lim_{x\to\infty} \frac{2f(x)f'(x)}{f(x)[1-2F(x)]} = 0$. For the second, write $x/v(x) = x(1 - F(x))/f(x)$ which has the limiting form $0/0$, and by L'Hopital's rule, that limit equals $\lim_{x\to\infty} \frac{1-F(x)-xf(x)}{-xf(x)/(2\sigma^2)} = \lim_{x\to\infty} \left[-\frac{1}{xv(x)/(2\sigma^2)} + 2\sigma^2\right] = 2\sigma^2$. □

Proof of Theorem A.1. We start by showing that $U_i(k_1, k_2, L)$ is a single-peaked function of k_i. Taking a derivative w.r.t. k_i yields

$$\begin{aligned}U_i'(k_i) &= -2\left[(k_i - b_i)(1 - F(k_i - k_j)) - \rho f(k_i - k_j)\right]\\ &= -2(1 - F(k_i - k_j))\left[k_i - b_i - \rho v(k_i - k_j)\right].\end{aligned}$$

Let $g(k_i)$ denote the term in square brackets above; $U_i'(k_i) = 0$ if and only if $g(k_i) = 0$. But Lemma A.3 implies that $g'(k_i) = 1 - \rho v'(k_i - k_j) \geq 1 - \sigma^2 v'(k_i - k_j) > 1/2$. It follows that g has at most one root. Additionally, $\lim_{k_i\to-\infty} v(k_i - k_j) = 0$, so $\lim_{k_i\to-\infty} g(k_i) = -\infty$. Hence, a root exists, and g is negative (positive) to the left (right) of this root. Hence, there is a unique solution to $U_i'(k_i) = 0$ and U_i is single-peaked. Combining these facts, U_i has a unique critical point, which is a global maximum.

We now look for min equilibria. The FOCs for the experts are equivalent to:

$$k_1 - b_1 - \rho \frac{f(k_1 - k_2)}{1 - F(k_1 - k_2)} = 0 \tag{5}$$

$$k_2 - b_2 - \rho \frac{f(k_1 - k_2)}{F(k_1 - k_2)} = 0. \tag{6}$$

Subtracting (6) from (5) substituting $z = k_1 - k_2$, we get

$$z - \rho \left[\frac{f(z)}{1 - F(z)} - \frac{f(z)}{F(z)} \right] = b_1 - b_2. \tag{7}$$

Denote $l(z) = z - \rho \left[\frac{f(z)}{1-F(z)} - \frac{f(z)}{F(z)} \right]$. Using inequalities from Lemma A.3, we obtain

$$l'(z) = 1 - \rho \left[v'(z) - w'(z) \right] \geq 1 - \sigma^2 [v'(z) - w'(z)] > 0.$$

Now $l(z)$ is continuous, strictly increasing on $\mathbb{R}$, and ranges from $-\infty$ to $+\infty$. Therefore (4) has a unique solution, z^*; we use $z(B)$ to denote explicitly the dependence on B.
Using this solution, we get (k_1^m, k_2^m) as the only critical point and check that this point satisfies both initial FOCs. As $U_i(k_i, k_j, L)$ is a single-peaked function of k_i, (k_1^m, k_2^m) is a pair of best responses.

As it was shown in Theorem 1, choosing the lower offer is the BR strategy for the principal iff $k_1 + k_2 \geq 0$, or equivalently

$$b_1 + b_2 + \rho \left[\frac{f(z^*)}{1 - F(z^*)} + \frac{f(z^*)}{F(z^*)} \right] \geq 0.$$

Also the LHS of (4) is equal to 0 at $z = 0$, and therefore $z^* \geq 0$ and $k_1^m - k_2^m \geq 0$.

Define a function $m(B) = b_1 + b_2 + \rho[v(z(B)) + w(z(B))]$; the min equilibrium exists if and only if $m(B) \geq 0$.
1) For $B \leq 2\sigma^2$: $m(B) \geq 0$, therefore the min equilibrium exists.
2) Next, we show that $m(B)$ is decreasing in B in the region $B \geq 2\sigma^2$.

$$m'(B) = -\frac{1}{2}[v(z(B)) + w(z(B))] + \rho[v'(z(B)) - v'(-z(B))]z'(B). \tag{8}$$

Differentiating equation (4) at point B, we get:

$$z'(B) - \rho[v'(z(B)) + v'(-z(B))]z'(B) + \frac{1}{2}[v(z(B)) - w(z(B))] = 0. \tag{9}$$

By substituting (9), the second term of (8) becomes

$$
\begin{aligned}
&+\rho[v'(z(B)) - v'(-z(B))]\frac{-\frac{1}{2}[v(z(B)) - w(z(B))]}{1-\rho\,[v'(z(B)) + v'(-z(B))]} \\
&\leq -\rho[v'(z(B)) - v'(-z(B))]\frac{\frac{1}{2}[v(z(B)) - w(z(B))]}{-\rho\,[v'(z(B)) + v'(-z(B))]} \\
&= \frac{1}{2}\frac{v'(z(B)) - v'(-z(B))}{v'(z(B)) + v'(-z(B))}[v(z(B)) - w(z(B))] \\
\implies m'(B) &\leq -\frac{1}{2}[v(z(B)) + w(z(B))] + \frac{1}{2}[v(z(B)) - w(z(B))] = -w(z(B)) < 0.
\end{aligned}
$$

3) From Lemma A.3 the hazard rate v is convex, so for any real x, $v(x) + w(x) = v(x) + v(-x) \geq 2v(0) > 0$, and $m(B)$ tends to $-\infty$ as B tends to ∞.

From 1)-3) follows that there exists $B_m : m(B) \geq 0$ iff $B \leq B_m$. Also

$$
\left(\frac{B_m}{2} - \sigma^2\right)[v(z(B_m)) + w(z(B_m))] = b_1 + b_2. \tag{10}
$$

As $z(B_m)$ satisfies equation (4), we have:

$$
\left(\frac{B_m}{2} - \sigma^2\right)[v(z(B_m)) - w(z(B_m))] + z(B_m) = b_1 - b_2. \tag{11}
$$

From the previous discussion and (10) we have $B_m \geq 2\sigma^2$. Also, (10) and the inequality $v(x) + w(x) = v(x) + v(-x) \geq 2v(0) = \frac{2}{\sqrt{\pi}\sigma}$ give an upper bound on B_m:

$$
\left(\frac{B_m}{2} - \sigma^2\right)\frac{2}{\sqrt{\pi}\sigma} \leq b_1 + b_2.
$$

Subtracting (11) from (10), we get a lower bound on B_m:

$$
2b_2 = (B_m - 2\sigma^2)w(z(B_m)) - z(B_m) \leq (B_m - 2\sigma^2)w(0) = (B_m - 2\sigma^2)\frac{1}{\sqrt{\pi}\sigma}.
$$

Finally, we calculate the expected bias of the chosen offer, its variance, and players' utilities:

$$
\begin{aligned}
\bar{b}(k_1^m, k_2^m, L) &= -2\sigma^2 f(z^*) + k_2^m F(z^*) + k_1^m(1 - F(z^*)) \\
&= -2\sigma^2 f(z^*) + b_2 F(z^*) + \rho f(z^*) + b_1(1 - F(z^*)) + \rho f(z^*) \\
&= b_1(1 - F(z^*)) + b_2 F(z^*) - B f(z^*); \\
Var(k_1^m, k_2^m, L) &= \sigma^2 - 4\sigma^4 f^2(z^*) - 2\sigma^2 z^* f(z^*)(2F(z^*) - 1) + (z^*)^2 F(z^*)(1 - F(z^*)).
\end{aligned}
$$

□

Corollary A.1. *In the min equilibrium,*

$$\begin{aligned} V(k_1^m, k_2^m, L) &= -\sigma^2 - b_1^2(1 - F(z^*)) - b_2^2 F(z^*) + B(b_1 + b_2) f(z^*) + \Delta(z^*) \\ U_1(k_1^m, k_2^m, L) &= -\sigma^2 - (b_1 - b_2)^2 F(z^*) + B(1 - F(z^*)) - B(b_1 - b_2) f(z^*) + \Delta(z^*) \\ U_2(k_1^m, k_2^m, L) &= -\sigma^2 - (b_1 - b_2)^2 (1 - F(z^*)) + BF(z^*) + B(b_1 - b_2) f(z^*) + \Delta(z^*), \end{aligned}$$

where $\Delta(z^*) := \left(\sigma^4 - \frac{B^2}{4}\right) \frac{f^2(z^*)}{F(z^*)(1-F(z^*))}$.

Proof. Immediate from applying Lemma A.1 to the markups given by Theorem A.1. □

Proof of Proposition 6. From Corollary A.1 above and Corollary S.3 in the supplementary appendix,

$$V_{min} - V_{max} = (2F(z^*) - 1)(b_1^2 - b_2^2) + 2B(b_1 + b_2) f(z^*) \geq 0,$$

with equality if and only if either $b_1 + b_2 = 0$ or both $B = 0$ and $b_1 = b_2$. □

Proof of Proposition 7. The principal's utilities V_{min} and V_{sim} from the min equilibrium and simple delegation to expert 2, respectively, are

$$\begin{aligned} V_{min} &= -\sigma^2 - b^2(1 - F) - r^2 b^2 F + \sigma^4 vw \\ V_{sim} &= -\sigma^2 - r^2 b^2 \\ \implies V_{min} - V_{sim} &= \sigma^4 vw - b^2(1 - r^2)(1 - F), \end{aligned}$$

where f, F, v, w take argument $z^* > 0$ defined implicitly by (4) specialized to $(b_1, b_2) = (b, rb)$:

$$z^* = b(1 - r) + \sigma^2(v(z^*) - w(z^*)). \tag{12}$$

It is immediate from inspection that if $r = 1$, $V_{min} > V_{sim}$, so for the rest of the proof, consider $r \in [-1, 1)$. Solving for b in (12), we have

$$V_{min} - V_{sim} = \sigma^4 vw - \frac{1 + r}{1 - r}(1 - F)[z - \sigma^2(v - w)]^2 =: A(z).$$

As $A(0) > 0$, we wish to show that (i) if $r \in [-1, 0]$, $A(z) > 0$ for all $z > 0$, and (ii) if $r \in (0, 1)$, there exists $\bar{z}$ such that $A(z) < 0$ for all $z > \bar{z}$. Note that $A(z) = 0$ iff

$$1/F(z) = \frac{1 + r}{1 - r} g(z)^2, \tag{13}$$

where $g(z) := \frac{z-\sigma^2(v-w)}{v\sigma^2}$. Note that $g(z) > 0$, using the fact that $g_2(z) := z - \sigma^2(v-w)$ satisfies $g_2(0) = 0$ and $g_2'(z) = 1 - \sigma^2(v' - w') > 0$ by Lemma A.3.

For case (i), the LHS of (13) is bounded below by 1 while the RHS is at most $g(z)^2$. Hence, it suffices to show that $g(z) < 1$, or equivalently $z < G(z) := \sigma^2[2v(z) - w(z)]$, for all $z > 0$. We have $v'(z) - v'(-z) = v'(z) + w'(z) = (v(z) + w(z))[v(z) - w(z) - z/(2\sigma^2)]$. Since v is strictly convex, for $z > 0$, $v'(z) - v'(-z) > 0$, and hence $v(z) > w(z) + z/(2\sigma^2)$. Therefore, $G(z) - z = \sigma^2[2v(z) - w(z)] - z > \sigma^2[2(w(z) + z/(2\sigma^2)) - w(z)] - z = \sigma^2 w(z) > 0$. We conclude that $G(z) > z$ and thus $g(z) < 1$ for all $z > 0$, as desired.

For case (ii), Lemma A.3 implies that $g(z) \to 1$ as $z \to +\infty$, so the RHS of (13) tends to $\frac{1+r}{1-r} > 1$, while the LHS tends to 1, and hence $\bar{z}$ exists.

For the final claim of the proposition, note that by using a change of variables $\tilde{z} := z/\sigma$ and $\tilde{b} = b/\sigma$, and letting ϕ and Φ denote the pdf and CDF of $N(0,1)$, respectively, we have $\tilde{z} = \tilde{b}(1-r) + \frac{\phi(\tilde{z}/\sqrt{2})}{\sqrt{2}}\left(\frac{1}{1-\Phi(\tilde{z}/\sqrt{2})} - \frac{1}{\Phi(\tilde{z}/\sqrt{2})}\right)$ in which σ only appears in $\tilde{z}$ and $\tilde{b}$, and $\frac{V_{min}-V_{sim}}{\sigma^2} = \frac{\frac{1}{2}\phi(\tilde{z}/\sqrt{2})^2}{(1-\Phi(\tilde{z}/\sqrt{2}))\Phi(\tilde{z}/\sqrt{2})} - \tilde{b}^2(1-r^2)(1-\Phi(\tilde{z}/\sqrt{2}))$. Fixing b and taking $\sigma \to +\infty$, we have $\tilde{b} \to 0$ and $\tilde{z} \to 0$, and it is easy to see that this implies $\lim_{\sigma \to 0} \frac{V_{min}-V_{sim}}{\sigma^2} > 0$. Hence, the min equilibrium is preferred for sufficiently large σ. □

Proof of Proposition 8. The principal's payoffs in the min equilibrium and under simple delegation to a super expert are, respectively,

$$V_{min} = -(\bar{b})^2 - Var = -\left(b - \frac{B}{2\sqrt{\pi}\sigma}\right)^2 - \left(1 - \frac{1}{\pi}\right)\sigma^2$$
$$V_{sim,super} = -\frac{\sigma^2}{2} - b^2.$$

Since V_{min} is a concave quadratic in B, either the set of B such that $V_{min} > V_{sim,super}$ is empty or this set is a nonempty interval of the form $(\underline{B}, \overline{B})$ where $\underline{B}, \overline{B} \in \mathbb{R}$. Note that V_{min} attains a maximum value of $-\left(1 - \frac{1}{\pi}\right)\sigma^2$ when $B = 2b\sqrt{\pi}\sigma$, where the first term vanishes. This maximum value exceeds $V_{sim,super}$ iff $-\left(1 - \frac{1}{\pi}\right)\sigma^2 > -\frac{\sigma^2}{2} - b^2$ which is equivalent to $b > \sigma\sqrt{\frac{1}{2} - \frac{1}{\pi}}$ as given in the proposition statement. To see that when this inequality holds we have $\underline{B} > 0$, note that by inspection, $V_{min} < V_{sim,super}$ when $B = 0$. □

Proof of Proposition 9. Min equilibrium exists only if $B \leq B_m$, and for opposite biases, the interval in Theorem A.1 implies $B_m = 2\sigma^2$; note that B_m is larger for symmetric (positive) biases, so min equilibrium exists for both cases when $B \leq 2\sigma^2$. From Theorem A.1 and

Lemma A.1,

$$V_{symm}(b) = -\sigma^2 - b^2 + 2Bbf(0) + (4\sigma^4 - B^2)f^2(0)$$
$$V_{opp}(b) = -\sigma^2 - b^2 + \left(\sigma^4 - \frac{B^2}{4}\right)\frac{f^2(z^*_{opp})}{F(z^*_{opp})(1-F(z^*_{opp}))} \geq -\sigma^2 - b^2 = V_{sim}(b),$$

where z^*_{opp} is the min equilibrium markup difference for oppositely biased experts. As $2Bbf(0) \geq 0$ and $\frac{f^2(z)}{F(z)(1-F(z))}$ reaches its maximum of $4f^2(0)$ at $z=0$, $V_{symm}(b) \geq V_{opp}(b)$. □

A.3 Proofs for Section 4

Proof of Proposition 2. If $b_1 = b_2 = b > 0$, then the upper and lower bounds on B_m from Theorem A.1 coincide, so $B_m = 2\sigma^2 + 2\sqrt{\pi}\sigma b$. The experts' markups are $k_1^m = k_2^m = k_m = b + \frac{\rho}{\sigma\sqrt{\pi}}$ and $z^* = 0$. □

Proof of Proposition 3. If $b_1 = b_2 = b > 0$, then the upper and lower bounds on B_M from Theorem S.1 in the supplementary appendix coincide, so $B_M = 2\sigma^2 - 2\sqrt{\pi}\sigma b$. The experts' markups are $k_1^M = k_2^M = k_M = b - \frac{\rho}{\sigma\sqrt{\pi}}$ and $z^* = 0$. □

Proof of Proposition 4. The comparison of the min and max equilibria follows as a special case of Proposition 6. The comparison of the min equilibrium to simple delegation follows as a special case of Proposition 7 by setting $r = 1$. □

Given $B > 0$, let σ^* denote the unique positive solution to the equation $(4\pi - 4)\sigma^4 + 2\sqrt{\pi}bB\sigma = B^2$.

Lemma A.4. *Consider $b_1 = b_2 = b > 0$. The principal's payoff is related to the noise level σ as follows:*

- *If $0 < B < \frac{2(\pi-1)\pi}{(\pi-2)^2}b^2$, there exists a non-empty interval $\sigma \in [\frac{\sqrt{\pi b^2+2B}-\sqrt{\pi}b}{2}, \sigma^*)$, where min equilibrium exists and the principal's expected payoff is increasing in σ; when $\sigma > \sigma^*$, the principal's expected payoff is decreasing in σ.*

- *If $B = 0$ or $B \geq \frac{2(\pi-1)\pi}{(\pi-2)^2}b^2$, the principal's expected payoff is always decreasing in σ.*

Proof. From Proposition 2, the symmetric min equilibrium exists if $B \leq 2\sigma^2 + 2\sqrt{\pi}b\sigma$ or, equivalently, $\sigma \geq \frac{\sqrt{\pi b^2+2B}-\sqrt{\pi}b}{2}$. The principal's expected payoff in the min equilibrium is equal to $V = -(b - \frac{B}{2\sqrt{\pi}\sigma})^2 - (1-\frac{1}{\pi})\sigma^2$. Then $V'(\sigma) = \frac{B}{\sqrt{\pi}\sigma^2}(\frac{B}{2\sqrt{\pi}\sigma} - b) - 2(1-\frac{1}{\pi})\sigma = -\frac{1}{2\pi\sigma^3}[(4\pi-4)\sigma^4 + 2\sqrt{\pi}bB\sigma - B^2] > 0$ if and only if $\sigma < \sigma^*$.

If $B = 0$, $V'(\sigma) < 0$. Otherwise, denote $\sigma_0 = \frac{\sqrt{\pi b^2+2B}-\sqrt{\pi}b}{2} > 0$.

The interval $[\sigma_0, \sigma^*)$ is non-empty if and only if $0 > (4\pi - 4)\sigma_0^4 + 2\sqrt{\pi}bB\sigma_0 - B^2 = (4\pi - 4)\sigma_0^4 - 2B\sigma_0^2 + B[2\sigma_0^2 + 2\sqrt{\pi}b\sigma_0 - B] = (4\pi - 4)\sigma_0^4 - 2B\sigma_0^2$ or, equivalently, $0 > (2\pi - 2)\sigma_0^2 - B = (2\pi - 2)\sigma_0^2 - (2\sigma_0^2 + 2\sqrt{\pi}b\sigma_0) = (2\pi - 4)\sigma_0^2 - 2\sqrt{\pi}b\sigma_0$. The latter holds if and only if $\sigma_0 < \frac{\sqrt{\pi}b}{\pi - 2}$ or, equivalently, $B = 2\sigma_0^2 + 2\sqrt{\pi}b\sigma_0 < \frac{2(\pi-1)\pi}{(\pi-2)^2}b^2$. □

Lemma A.5. *Suppose that bonus B is paid by the principal. The principal's optimal choice of B depends on the bias-to-noise ratio as follows:*

- *If $\frac{b}{\sigma} \leq \sqrt{\pi}$, then the principal pays no bonus: $B = 0$.*
- *If $\frac{b}{\sigma} > \sqrt{\pi}$, then the principal chooses bonus $B = 2\sqrt{\pi}\sigma^2 \left[\frac{b}{\sigma} - \sqrt{\pi}\right]$. The increase in the principal's expected payoff, relative to when the bonus is restricted to be 0, is equal to $(b - \sqrt{\pi}\sigma)^2$.*

Proof. By Proposition 6, only min equilibrium should be considered. Therefore, we seek to maximize the quadratic function $V(k_m, k_m, L) - B = -\left(b - \frac{B}{2\sqrt{\pi}\sigma}\right)^2 - \sigma^2 + \frac{\sigma^2}{\pi} - B$ on the interval $B \in [0, 2\sigma^2 + 2\sqrt{\pi}\sigma b]$.

a) First consider $\frac{b}{\sigma} \leq \sqrt{\pi}$. In this case the maximum is achieved at $B = 0$ and is equal to $V(k_m, k_m, L|B = 0) = -b^2 - \left(1 - \frac{1}{\pi}\right)\sigma^2$. Hence, if $\frac{b}{\sigma} < \sqrt{\pi}$, then the principal pays no bonus.

b) Now consider $\frac{b}{\sigma} \geq \sqrt{\pi}$. In this case the principal achieves maximum in the min equilibrium at $B_R = 2\sqrt{\pi}\sigma\,(b - \sqrt{\pi}\sigma)$ (min equilibrium exists for this point) and is equal to $V(k_m, k_m, L|B = B_R) = (\pi - 1 + \frac{1}{\pi})\sigma^2 - 2\sqrt{\pi}\sigma b$.

Her gains compared to $B = 0$ (if she is legally restricted from paying bonuses) are equal to:

$$V(k_m, k_m, L|B = B_R) - V(k_m, k_m, L|B = 0) = \left(b - \sqrt{\pi}\sigma\right)^2.$$

□

Proof of Proposition 5. For nonmonotonicity in b, note that the principal's payoff in the min equilibrium is $-\bar{b}^2 - Var = -\left(b - \frac{B}{2\sqrt{\pi}\sigma}\right)^2 - \left(1 - \frac{1}{\pi}\right)\sigma^2$, which is clearly increasing in b for $b \in \left[0, \frac{B}{2\sqrt{\pi}\sigma}\right]$. Nonmonotonicity in σ is established by Lemma A.4, and nonmonotonicity in B is established by Lemma A.5. □

A.4 Proofs for Section 6

Before proving Proposition 10, we first state and prove two lemmas.

Lemma A.6. *Set $B = 0$, and suppose the principal commits to (p_m, p_M, p_1, p_2). For any b_1, b_2, in any equilibrium of the game between experts, if $p_{-i} \neq 1$, the expected bias of expert i's action conditional on expert i being chosen is b_i.*

Proof. We prove that, more generally, fixing any (not necessarily optimal) markup k_j of the other expert, the first order condition for expert i implies $\bar{b}_i = b_i$, where $\bar{b}_i$ is the expected bias of the action conditional on expert i being chosen.

The expected payoff for expert 1 given k_1, k_2 and the principal's strategy is

$$\begin{aligned} u_1(k_1, k_2) = {} & p_1(-\sigma^2 - (k_1 - b_1)^2) + p_2(-\sigma^2 - (k_2 - b_1)^2) \\ & + p_m \left[-\sigma^2 + 2\sigma^2(k_1 + k_2 - 2b_1)f(z) - (k_2 - b_1)^2 F(z) \right. \\ & \quad \left. -(k_1 - b_1)^2(1 - F(z))\right] \\ & + p_M \left[-\sigma^2 - 2\sigma^2(k_1 + k_2 - 2b_1)f(z) - (k_1 - b_1)^2 F(z) \right. \\ & \quad \left. -(k_2 - b_1)^2(1 - F(z))\right]. \end{aligned}$$

The first order condition for expert 1 is

$$\begin{aligned} 0 = \frac{\partial}{\partial k_1} u_1(k_1, k_2) = {} & -2k_1 p_1 + 2b_1 p_1 + p_m \left[2\sigma^2 f(z) - (k_1 + k_2 - 2b_1) z f(z) - (k_2 - b_1)^2 f(z) \right. \\ & \left. +(k_1 - b_1)^2 f(z) - 2(k_1 - b_1)(1 - F(z))\right] \\ & + p_M \left[-2\sigma^2 f(z) + (k_1 + k_2 - 2b_1) z f(z) - 2(k_1 - b_1)F(z) \right. \\ & \left. -(k_1 - b_1)^2 f(z) + (k_2 - b_1)^2 f(z)\right] \\ \implies k_1 = {} & b_1 + \sigma^2 \frac{(p_m - p_M)f(z)}{p_1 + p_m(1 - F(z)) + p_M F(z)}, \end{aligned} \tag{14}$$

which is well-defined whenever $p_2 \neq 1$.

Note that conditioning on $\theta = 0$ implies $\bar{b}_1 = \mathbb{E}[s_1 + k_1 | \text{expert 1 chosen}]$. We have

$$\begin{aligned} \bar{b}_1 \cdot \Pr(\text{expert 1 chosen}) = {} & p_1 \mathbb{E}[s_1 + k_1] + p_m \Pr(s_1 + k_1 \leq s_2 + k_2) \mathbb{E}[s_1 + k_1 | s_1 + k_1 \leq s_2 + k_2] \\ & + p_M \Pr(s_1 + k_1 \geq s_2 + k_2) \mathbb{E}[s_1 + k_1 | s_1 + k_1 \geq s_2 + k_2] \\ = {} & p_1 k_1 + p_m(1 - F(z))(k_1 + \mathbb{E}[s_1 | s_1 \leq s_2 + k_2 - k_1]) \\ & + p_M F(z)(k_1 + \mathbb{E}[s_1 | s_1 \geq s_2 + k_2 - k_1]) \\ = {} & p_1 k_1 + p_m(1 - F(z)) \left(k_1 - \sigma^2 \frac{f(z)}{1 - F(z)}\right) \\ & + p_M F(z) \left(k_1 + \sigma^2 \frac{f(z)}{F(z)}\right). \end{aligned}$$

Now $\Pr(\text{expert 1 chosen}) = p_1 + p_m(1-F(z)) + p_M F(z)$, so the equation above implies

$$\begin{aligned}\bar{b}_1 &= \frac{p_1 k_1 + p_m(1-F(z))\left(k_1 - \sigma^2 \frac{f(z)}{1-F(z)}\right) + p_M F\left(k_1 + \sigma^2 \frac{f(z)}{F(z)}\right)}{p_1 + p_m(1-F(z)) + p_M F(z)} \\ &= k_1 - \sigma^2 \frac{(p_m - p_M) f(z)}{p_1 + p_m(1-F(z)) + p_M F(z)} \\ &= b_1,\end{aligned}$$

where the last line follows from (14). A symmetric argument for expert 2 shows that $k_2 = b_2 + \sigma^2 \frac{(p_m - p_M) f(z)}{p_2 + p_m F(z) + p_M(1-F(z))}$ is expert 2's unique best response. □

Lemma A.7. *Suppose the principal commits to (p_m, p_M, p_1, p_2). For any pair of markups $(k_1, k_2) \in \mathbb{R}^2$, the variance of the chosen offer is bounded below by $\sigma^2(1 - 1/\pi)$.*

Proof. To simplify calculations, normalize $\theta = 0$. Recall from Lemma A.2 the random variables $\xi(k_1, k_2) := \min(\epsilon_1 + k_1, \epsilon_2 + k_2)$ and $\eta(k_1, k_2) := \max(\epsilon_1 + k_1, \epsilon_2 + k_2)$, where $\epsilon_1, \epsilon_2 \sim N(0, \sigma^2)$ are independent. The expectation of the chosen offer is $\mathbb{E}[a] = p_m \mathbb{E}\xi(k_1, k_2) + p_M \mathbb{E}\eta(k_1, k_2) + p_1 k_1 + p_2 k_2$. The expectation of the square of the chosen offer is

$$\mathbb{E}[a^2] = p_m \mathbb{E}\xi^2(k_1, k_2) + p_M \mathbb{E}\eta^2(k_1, k_2) + p_1(\sigma^2 + k_1^2) + p_2(\sigma^2 + k_2^2).$$

The variance $A(p_m, p_M, p_1, p_2) = \mathbb{E}[a^2] - (\mathbb{E}[a])^2$ is a concave function of (p_m, p_M, p_1, p_2).[19] It follows that its minimum is attained at one of the vertices of the simplex $\{(p_m, p_M, p_1, p_2) \in [0,1]^4 : p_m + p_M + p_1 + p_2 = 1\}$. Clearly, $A(0,0,1,0) = A(0,0,0,1) = \sigma^2 > \sigma^2(1 - 1/\pi)$. Moreover, from the variance expression in Lemma A.1,

$$A(1,0,0,0) = \sigma^2(1 - 4\sigma^2 f^2(z)) - 2\sigma^2 z f(z)(2F(z) - 1) + z^2 F(z)(1 - F(z)).$$

Note that the expression above is an even function of z, and when evaluated at $z = 0$, it simplifies to $\sigma^2(1 - 4\sigma^2 f(0)^2)) = \sigma^2(1 - 1/\pi)$. Hence, to establish that $\sigma^2(1 - 1/\pi)$ is a lower bound for all $z \in \mathbb{R}$, it suffices to show that the function is nondecreasing for $z \geq 0$. Differentiating with respect to z yields $2zF(z)(1 - F(z)) - 2\sigma^2 f(z)(2F(z) - 1)$. This is nonnegative if and only if $z \geq \sigma^2 \frac{f(z)(2F(z)-1)}{F(z)(1-F(z))} = \sigma^2(v(z) - w(z))$, where $v(z)$ and $w(z)$ are defined as in Lemma A.3. Since there is equality at $z = 0$, and by Lemma A.3, $\sigma^2(v' - w') < 1$, we have $z > \sigma^2(v(z) - w(z))$ for all $z > 0$, as desired. We conclude that $A(1,0,0,0) \geq \sigma^2(1 - 1/\pi)$, and by symmetry, it follows that $A(0,1,0,0) \geq \sigma^2(1 - 1/\pi)$.

[19] In particular, after plugging in the expressions for $\mathbb{E}\xi(k_1, k_2)$, $\mathbb{E}\xi^2(k_1, k_2)$, $\mathbb{E}\eta(k_1, k_2)$, and $\mathbb{E}\eta^2(k_1, k_2)$ from Lemma A.1, it can be shown that the Hessian matrix for $A(p_m, p_M, p_1, p_2)$ is negative semi-definite as its leading principal minors are $(-2(k_1[1 - F(z)] + k_2 F(z) - 2\sigma^2 f(z))^2, 0, 0, 0)$.

Hence, we have shown that all four vertices are bounded below by $\sigma^2(1-1/\pi)$, completing the proof. □

Proof of Proposition 10. By taking expectations over the chosen expert, Lemma A.6 implies that the ex ante expected bias of the chosen offer is b. It follows that the principal's expected utility from committing to (p_m, p_M, p_1, p_2) is $-b^2 - A(p_m, p_M, p_1, p_2)$, where A is the variance of the chosen offer as in Lemma A.7. By Lemma A.7, $A(p_m, p_M, p_1, p_2) \geq \sigma^2(1-1/\pi)$, and as shown in the proof of Lemma A.7, the lower bound $\sigma^2(1-1/\pi)$ is achieved when the principal commits to $(p_m, p_M, p_1, p_2) = (1, 0, 0, 0)$ and $k_1 = k_2$ (so that $z = 0$). By committing to $(p_m, p_M, p_1, p_2) = (1, 0, 0, 0)$, the principal induces equilibrium markups $k_1 = k_2$, so this policy achieves the lower bound. This establishes the optimality of committing to the policy as in the proposition statement.

For the second claim, suppose $B = 0$, $b_1 = b > 0$ and $b_2 = rb$, where $r \in [-1, 0)$. By similar arguments to those in the proof of Theorem A.1, there is a unique equilibrium of the game between experts; let $V(p)$ denote the principal's utility. It suffices to show that $V'(1) < 0$ for sufficiently large b. From (14) and the analogous expression for expert 2, the markups are

$$\begin{aligned} k_1(p) &= b + \sigma^2 \frac{(2p-1)f(z(p))}{W(p)} \\ k_2(p) &= rb + \sigma^2 \frac{(2p-1)f(z(p))}{1-W(p)} \\ \Longrightarrow \quad z(p) &= b(1-r) + \sigma^2 \frac{(2p-1)f(z(p))(1-2W(p))}{W(p)(1-W(p))}, \end{aligned} \tag{15}$$

where $W(p) := p(1-F(z(p))) + (1-p)F(z(p))$. Differentiating with respect to p, evaluating at $p = 1$ and solving for $z'(1)$ yields

$$z'(1) = \frac{2\sigma^2 f(2F-1)}{2F^2(1-F)^2 + zfF(1-F)(2F-1) - 2\sigma^2 f^2[1-2F+2F^2]}.$$

The principal's utility is

$$V(p) = -\sigma^2 + 2\sigma^2(2p-1)f(k_1+k_2) - k_1^2 W - k_2^2(1-W).$$

By differentiating with respect to p, evaluating at $p = 1$, substituting in the above expression for $z'(1)$ and using (15) to eliminate b, we obtain

$$
\begin{aligned}
V'(1) &= \frac{g_1}{g_2}, \text{ where} \\
g_1 &:= 2(1+r)z^2(1-F)^3F^3(2F-1) + (1+r)z(z^2-4\sigma^2)fF^2(1-F)^2(2F-1)^2 \\
&\quad + 8\sigma^6 f^4[-1+3(1+r)F-6(1+r)F^2+4(1+r)F^3] \\
&\quad + 2z\sigma^4 f^3(2F-1)[-1+7(1+r)F-18(1+r)F^2+12(1+r)F^3] \\
&\quad - 2\sigma^2 f^2(1-F)F\left[-(1+r)z^2+2r\sigma^2+2(1+r)(4z^2-3\sigma^2)F \right.\\
&\qquad \left. -6(1+r)(3z^2-2\sigma^2)F^2+4(1+r)(3z^2-2\sigma^2)F^3\right], \\
g_2 &:= (1-r)(1-F)F[2(1-F)^2F^2+zfF(1-F)(2F-1)-2\sigma^2f^2(1-2F+2F^2)].
\end{aligned}
$$

By similar reasoning as in the proof of Proposition 7, it suffices to show that $V'(1) < 0$ for sufficiently large z. In particular, we show that for sufficiently large z, $g_2 > 0$ and $g_1 < 0$.

First, to show that $g_2 > 0$ for large z, note that

$$
g_2/[(1-r)F(1-F)^3] = \frac{2(1-F)^2F^2+zfF(1-F)(2F-1)-2\sigma^2f^2(1-2F+2F^2)}{(1-F)^2}.
$$

By repeated use of L'Hopital's rule and application of Lemma A.3, it is easy to see that the above has limit 1, and hence $g_2 > 0$ for sufficiently large z.

To show that $g_1 < 0$ for sufficiently large z, we show that $g_1/(zf^3) \to 2r\sigma^4 < 0$. After eliminating terms that vanish, $g_1/(zf^3)$ has the same limit as $\tilde{g}_1/(zf^3)$, where $\tilde{g}_1 := \gamma_1 + \gamma_2 + \gamma_3$, and where in turn

$$
\begin{aligned}
\gamma_1 &:= (1+r)z(z^2-4\sigma^2)fF^2(1-F)^2(2F-1)^2 \\
\gamma_2 &:= 2z\sigma^4f^3(2F-1)[-1+7(1+r)F-18(1+r)F^2+12(1+r)F^3] \\
\gamma_3 &:= -2\sigma^2f^2(1-F)F\left[-(1+r)z^2+2r\sigma^2+2(1+r)(4z^2-3\sigma^2)F \right.\\
&\qquad \left. -6(1+r)(3z^2-2\sigma^2)F^2+4(1+r)(3z^2-2\sigma^2)F^3\right].
\end{aligned}
$$

Using from Lemma A.3 that $\lim_{z\to\infty} z(1-F)/f = 2\sigma^2$, we obtain the limits $\gamma_1/(zf^3) \to 4\sigma^4(1+r)$, $\gamma_2/(zf^3) \to 2r\sigma^4$ and $\gamma_3/(zf^3) \to -4\sigma^4(1+r)$, and hence $g_1/(zf^3) \to 2r\sigma^4 < 0$, as desired. □

Proof of Proposition 11. The expected payoff of expert i is

$$U_i(k_i, k_j, L) = \int\limits_{k_i - k_j}^{\infty} \left[B - \left(k_i - b - \frac{t}{2}\right)^2 - \frac{\sigma^2}{2}\right] f(t)\, dt$$

$$= \left[B - \sigma^2 - (k_i - b)^2\right](1 - F(k_i - k_j)) + \left[2\sigma^2(k_i - b) - \frac{1}{2}\sigma^2(k_i - k_j)\right] f(k_i - k_j).$$

First, we calculate marginal utilities:

$$U_i'(k_i) = -2(k_i - b)[1 - F(k_i - k_j)] + \left[\frac{1}{2}\sigma^2 + 2\rho + \frac{1}{4}(k_i + k_j - 2b)^2\right] f(k_i - k_j).$$

Here, setting $U_i'(k) = 0$ gives two critical points:
$k = b + \sqrt{\pi}\sigma - \sqrt{\left(\pi - \frac{5}{2}\right)\sigma^2 + B}$ and $k = b + \sqrt{\pi}\sigma + \sqrt{\left(\pi - \frac{5}{2}\right)\sigma^2 + B}$.

The second derivative is:

$$U_i''(k_i) = -2[1 - F(k_i - k_j)]$$
$$+ \left[2(k_i - b) + \frac{1}{2}(k_i + k_j - 2b) + \left(\frac{B}{2\sigma^2} - \frac{5}{4}\right)(k_i - k_j) - \frac{1}{8\sigma^2}(k_i - k_j)(k_i + k_j - 2b)^2\right] f(k_i - k_j).$$

We get that only $k^* = b + \sqrt{\pi}\sigma - \sqrt{\left(\pi - \frac{5}{2}\right)\sigma^2 + B}$ is a local maximum of experts' utility functions.

Optimality for the principal holds if and only if $k^* \geq 0$ or, equivalently, $B \leq b^2 + 2\sqrt{\pi}b\sigma + \frac{5}{2}\sigma^2$.

Calculating, we get that $U_i(k^*, k^*, L) = \frac{(\pi-1)\sigma}{\sqrt{\pi}}\sqrt{\left(\pi - \frac{5}{2}\right)\sigma^2 + B} - (\pi - \frac{7}{4})\sigma^2$. As we noted earlier, a necessary condition for equilibrium is $U_i(k^*, k^*, L) \geq 0$ or, equivalently, $B \geq (\frac{5}{2} - \frac{3\pi(8\pi-11)}{16(\pi-1)^2})\sigma^2$.

Therefore, min equilibrium may exist only if $B \in \left[\left(\frac{5}{2} - \frac{3\pi(8\pi-11)}{16(\pi-1)^2}\right)\sigma^2, \frac{5}{2}\sigma^2 + 2\sqrt{\pi}b\sigma + b^2\right]$.

To finish the proof, we show that if B lies on this interval, then $k = k^*$ is a global maximum of $U_1(k, k^*, L)$.

Denote $g(k) = -2(k - b) + \left[\frac{\sigma^2}{2} + 2\rho + \frac{1}{4}(k + k^* - 2b)^2\right] v(k - k^*)$.[20]

Then $U_1'(k) = -2(k - b)[1 - F(k - k^*)] + \left[\frac{\sigma^2}{2} + 2\rho + \frac{1}{4}(k + k^* - 2b)^2\right] f(k - k^*) = [1 - F(k - k^*)]g(k)$ and $sign(U_1'(k)) = sign(g(k))$.

[20]Recall that $v(k - k^*) = \frac{f(k-k^*)}{1-F(k-k^*)}$).

The first and second derivatives of g are

$$g'(k) = -2 + \left[\frac{\sigma^2}{2} + 2\rho + \frac{1}{4}(k+k^*-2b)^2\right] v'(k-k^*) + \frac{1}{2}(k+k^*-2b)v(k-k^*)$$
$$g''(k) = \left[\frac{\sigma^2}{2} + 2\rho - B + \frac{1}{4}(k+k^*-2b)^2\right] v''(k-k^*) + (k+k^*-2b)v'(k-k^*) + \frac{1}{2}v(k-k^*).$$

Consider two cases.

1. $B \in \left[\left(\frac{5}{2} - \frac{3\pi(8\pi-11)}{16(\pi-1)^2}\right)\sigma^2, \frac{5}{2}\sigma^2\right]$. Here $k^* \geq b$.

a) On the interval $k < b$, $U_1'(k) > 0$ and hence there is no point of maximum there.

b) On the interval $k \geq b$ we also have that $k + k^* - 2b \geq 0$. As all v, v' and v'' are strictly positive functions, $g''(k) > 0$. As $g'(k^*) < 0$ and $g'(+\infty) > 0$, hence there exists $k^{**} > k^*$: for $k < k^{**}$ $g(k)$ is decreasing; for $k > k^{**}$, $g(k)$ is increasing. As $g(b) > 0$, $g(k^*) = 0$, $g(k^{**}) < 0$ and $g(+\infty) > 0$, then there exists $k_0 > k^{**} : g(k_0) = 0$. In summary, $g(k)$ is negative only on (k^*, k_0). Consequently, $U_1(k)$ is increasing on $[b, k^*)$, decreasing on (k^*, k_0), increasing for $k > k_0$. Hence, to show that k^* is a maximum on the interval $k \geq b$ it is sufficient to verify that $U_1(k^*) \geq U_1(+\infty) = 0$, which has already been done.

2. $B \in \left[\frac{5}{2}\sigma^2, \frac{5}{2}\sigma^2 + 2\sqrt{\pi}b\sigma + b^2\right]$. Here $k^* \leq b$.

a) On the interval $k < k^*$: $U_1'(k) > 2(b-k^*)[1-F(k^*-k^*)] + \left[\frac{\sigma^2}{2} + 2\rho + \frac{1}{4}(k^*+k^*-2b)^2\right] f(k-k^*) = b - k^* - \frac{b-k^*}{f(0)} f(k-k^*)$ (as k^* is a solution of $\frac{5}{2}\sigma^2 - B + (k^*-b)^2 = \frac{k^*-b}{f(0)}$). Therefore $U_1'(k) > b - k^* - \frac{b-k^*}{f(0)} f(k-k^*) \geq b - k^* - \frac{b-k^*}{f(0)} f(0) = 0$.

b) On the interval $k \in (k^*, b]$: $\frac{U_1'(k)}{f(k-k^*)} = -2(k-b)\frac{1-F(k-k^*)}{f(k-k^*)} + \left[\frac{\sigma^2}{2} + 2\rho + \frac{1}{4}(k+k^*-2b)^2\right] < 2(b-k^*)\frac{1-F(k^*-k^*)}{f(k^*-k^*)} + \left[\frac{\sigma^2}{2} + 2\rho + \frac{1}{4}(k^*+k^*-2b)^2\right] = \frac{b-k^*}{f(0)} - \frac{b-k^*}{f(0)} = 0$, hence $U_1'(k) < 0$.

c) On the interval $k \in \left(b, 2b - k^* + 2\sqrt{B - \frac{5}{2}\sigma^2}\right]$ we also have $\frac{\sigma^2}{2} + 2\rho + \frac{1}{4}(k+k^*-2b)^2 \leq 0$. Then $U_1'(k) = -2(k-b)[1-F(k-k^*)] + \left[\frac{\sigma^2}{2} + 2\rho + \frac{1}{4}(k+k^*-2b)^2\right] f(k-k^*) < 0$.

d) On the interval $k > 2b - k^* + 2\sqrt{B - \frac{5}{2}\sigma^2}$ we also have $k + k^* - 2b > 0$. Hence, on this interval $g''(k) > 0$. Also notice that $g(2b - k^* + 2\sqrt{B - \frac{5}{2}\sigma^2}) < 0$. Then two cases are possible: (i) $g'(2b - k^* + 2\sqrt{B - \frac{5}{2}\sigma^2}) \geq 0$. Then on the whole interval $g'(k) > 0$ and $g(k)$ is increasing. As $g(2b - k^* + 2\sqrt{B - \frac{5}{2}\sigma^2}) < 0$ and $g(+\infty) > 0$, there exists k_0: $g(k) < 0$ for $k < k_0$ and $g(k) > 0$ for $k > k_0$. Now $U_1(k)$ is decreasing for $k < k_0$ and increasing for $k > k_0$. Hence, to show that k^* is a global maximum it is enough to check that $U_1(k^*) \geq U_1(+\infty) = 0$, which has already been done.

(ii) $g'(2b - k^* + 2\sqrt{B - \frac{5}{2}\sigma^2}) < 0$. Then there exists $k^{**} > 2b - k^* + 2\sqrt{B - \frac{5}{2}\sigma^2}$: for $k < k^{**}$ $g(k)$ is decreasing; for $k > k^{**}$ $g(k)$ is increasing. As $g(+\infty) > 0$, there exists k_0: $g(k) < 0$ for $k < k_0$ and $g(k) > 0$ for $k > k_0$. Then $U_1(k)$ is decreasing for $k < k_0$ and increasing for $k > k_0$ and k^* is a global maximum as $U_1(k^*) \geq U_1(+\infty) = 0$.

We now verify that $k_m = k^* > k_m^{bas.}$:

$$k_m = b + \left(\sqrt{\pi} - \sqrt{\pi + \frac{B}{\sigma^2} - \frac{5}{2}}\right)\sigma > k_m^{bas.} = b + \left(1 - \frac{B}{2\sigma^2}\right)\frac{\sigma}{\sqrt{\pi}}$$

$$\iff \sqrt{\pi} - \sqrt{\pi + \frac{B}{\sigma^2} - \frac{5}{2}} > \left(1 - \frac{B}{2\sigma^2}\right)\frac{1}{\sqrt{\pi}} \iff \sqrt{\pi} + \frac{1}{\sqrt{\pi}}\left(\frac{B}{2\sigma^2} - 1\right) > \sqrt{\pi + \frac{B}{\sigma^2} - \frac{5}{2}}$$

$$\iff \pi + \frac{B}{\sigma^2} - 2 + \frac{1}{\pi}\left(\frac{B}{2\sigma^2} - 1\right)^2 > \pi + \frac{B}{\sigma^2} - \frac{5}{2} \iff \frac{1}{2} + \left(\frac{B}{2\sigma^2} - 1\right)^2 > 0.$$

□

Proof of Proposition 12. Conjecturing common markups $k > 0$, after observing offers (a_1, a_2), the principal infers the signals were $(s_1, s_2) = (a_1 - k, a_2 - k)$. The principal's posterior expectation of θ is thus $\frac{s_1+s_2}{2}$ as in the baseline model, and the principal prefers the lower offer. As in the baseline model, imposing the FOC for expert 1 at k_1 conjecturing k_2 yields

$$\begin{aligned} k_1 &= b + \mathbb{E}[\theta | s_2 > s_1 + k_1 - k_2, s_1] - s_1 \\ &= b + \mathbb{E}[\mathbb{E}[\theta | s_1, s_2] | s_2 > s_1 + k_1 - k_2, s_1] - s_1 \\ &= b + \mathbb{E}[(s_1 + s_2)/2 | s_2 > s_1 + k_1 - k_2, s_1] - s_1 \\ &= b + \frac{1}{2}\mathbb{E}[s_2 | s_2 > s_1 + k_1 - k_2, s_1] - s_1/2. \end{aligned}$$

Now $s_2 | s_1 \sim N(s_1, 2\sigma^2(1-p))$,[21] and thus $\mathbb{E}[s_2 | s_2 > s_1 + k_1 - k_2, s_1] = s_1 + \hat{\sigma}\frac{\phi((k_1-k_2)/\hat{\sigma})}{1-\Phi((k_1-k_2)/\hat{\sigma})}$, where $\hat{\sigma} := \sqrt{2\sigma^2(1-p)}$, and thus the FOC for expert 1 reduces to $k_1 = b + \frac{\hat{\sigma}}{2}\frac{\phi((k_1-k_2)/\hat{\sigma})}{1-\Phi((k_1-k_2)/\hat{\sigma})}$. Repeating the process for expert 2, subtracting the equations, and simplifying, we get that an equation that implies $k_1 = k_2$. Substituting this back into the FOC for expert 1 yields the unique candidate equilibrium markup

$$k^* = b + \frac{\hat{\sigma}}{2}\frac{\phi(0)}{1 - \Phi(0)} = b + \frac{\hat{\sigma}}{2}\frac{\frac{1}{\sqrt{2\pi}}}{\frac{1}{2}} = b + \sigma\sqrt{(1-p)/\pi}.$$

As $k^* > b \geq 0$, the sum of markups is positive, so choosing the lower offer is indeed a best response for the principal. By similar calculations to those used in the proof of Theorem A.1, it can be verified that there is a unique solution to the FOC for expert i given any conjecture about k_j, and the SOC is satisfied. Hence, (k^*, k^*, L) is an equilibrium, and it is unique within the class of min equilibria.

[21] To see this, note that $s_2 = \theta + \epsilon_2 = s_1 - \epsilon_1 + \epsilon_2$. As $(-\epsilon_1)$ and ϵ_2 have correlation $-p$ and each has variance σ^2, $-\epsilon_1 + \epsilon_2 \sim N(0, 2\sigma_2 + 2\sigma^2(-p))$.

We now analyze the variance of the chosen offer (conditional on $\theta = 0$). Let $Y = \min\{s_1, s_2\}$; by standard formulas (Nadarajah and Kotz, 2008), $\mathbb{E}[Y] = -\sigma\sqrt{(1-p)/\pi}$ and $\mathbb{E}[Y^2] = \sigma^2$, and thus the variance of the chosen offer is $\text{Var}(k^* + Y) = \text{Var}(Y) = \sigma^2(1-(1-p)/\pi) \leq \sigma^2$, with equality iff $p = 1$. Since the expected bias is b in either the min equilibrium or in simple delegation, by the same decomposition of payoffs as in Section 4.3, the principal prefers the min equilibrium. □

Proof of Proposition 13. It remains true that when $B = 0$, fixing any markups k_{-i} of the other experts, expert i's markup is such that $s_i + k_i = b + \mathbb{E}[\theta|\text{expert } i \text{ is chosen}]$. Hence, conjecturing that $k_j = k^*$ for all $j \neq i$, we have

$$
\begin{aligned}
s_i + k_i &= b + \mathbb{E}[\theta | s_i \leq s_j + k^* - k_i \quad \forall j \neq i] \\
&= b + \mathbb{E}[\mathbb{E}[\theta|s_1, \ldots, s_N]|s_i \leq s_j + k^* - k_i \quad \forall j \neq i] \\
&= b + \mathbb{E}\left[\frac{s_1 + \cdots + s_N}{N}|s_i \leq s_j + k^* - k_i \quad \forall j \neq i\right] \\
&= b + \frac{s_i}{N} + \frac{N-1}{N}\mathbb{E}[s_j|s_i \leq s_j + k^* - k_i \quad \text{for fixed } j \neq i] \\
&= b + \frac{s_i}{N} + \frac{N-1}{N}\left(s_i + 2\sigma^2\frac{f(k_i - k^*)}{1 - F(k_i - k^*)}\right) \\
\iff \quad k_i &= b + \frac{N-1}{N}2\sigma^2 v(k_i - k^*),
\end{aligned}
$$

recalling that $v = f/(1-F)$. Now from Lemma A.3, $0 < v' < 1/(2\sigma^2)$, so the LHS exceeds the RHS when k_i is sufficiently high, and there is at most one crossing point. And since $v > 0$, the RHS exceeds the LHS when k_i is sufficiently low. Hence, there is a unique solution, and it is easy to verify that it is indeed a best response. For k^* defined in the proposition statement, $k_i = k^*$ is a best response for each i when $k_j = k^*$ for all $j \neq i$. And since $b > 0$, the expected bias of the lowest offer is positive, and all other offers are more biased upward by definition, so given that she can only use ordinal information, it is a best response for the principal to choose the expert with the lowest offer. It follows that there is a unique symmetric min equilibrium, and it is characterized by a common markup $k^* = b + \frac{N-1}{N}2\sigma^2 v(0) = b + \frac{N-1}{N}\frac{2\sigma}{\sqrt{\pi}}$ as in the proposition statement.

An upper bound on the variance of the minimum of N standard normal random variables is 1 (see Boucheron et al. (2012)), and hence (after scaling) the variance of the chosen action is at most σ^2. Since the expected bias of the chosen action is b in both the symmetric min equilibrium and under simple delegation, the principal prefers the former. □

Proof of Proposition 14. First, we show that if both experts play constant markup strategies with $k_i > b_i$, it is a best response for the principal to always choose the lower offer. Note that conditional on observing offers a_1 and a_2, the principal's posterior belief about θ is symmetric with mean $\frac{s_1+s_2}{2}$, where $s_i = a_i - k_i$, $i = 1, 2$, are the implied signals of the experts under their conjectured strategies. To see this, note that by the symmetry of the noise distributions around 0 (and symmetry across experts), given state θ', the signal pair (s_1, s_2) has density $g(s_1 - \theta')g(s_2 - \theta') = g(\theta' - s_1)g(\theta' - s_2) = g(s_2 - \theta'')g(s_1 - \theta'')$, where $\theta'' = s_1 + s_2 - \theta'$, which is the density of signal pair (s_1, s_2) given state θ''. As θ' and θ'' are symmetric across $\frac{s_1+s_2}{2}$ and θ' is arbitrary, given the diffuse prior, the principal's posterior must have mean $\frac{s_1+s_2}{2}$. Now $k_i > b_i$ for $i = 1, 2$ implies that $k_1 + k_2 > b_1 + b_2 \geq 0$. Since ℓ is symmetric about 0 and strictly concave, the principal's expected payoff is maximized by choosing the lower of the two offers, as in the baseline model, as this offer is the closest to the principal's expectation of the state.

Hence, to complete the proof, it suffices to show that there exists a (Bayesian Nash) equilibrium in constant markup strategies of the game induced by the principal's strategy of choosing the lower offer, and the markups satisfy with $k_i > b_i$. We first verify that if the principal chooses the lower offer and expert 2 applies any constant markup k_2, expert 1's best response is a constant markup $k_1 > b_1$; an analogous argument then holds for expert 2's best response. Due to the diffuse prior and the stationarity of the principal's strategy and expert 2's strategy, we can analyze expert 1's best response problem by fixing $\theta = 0$. Let G denote the CDF corresponding to the density g. Expert 1's expected payoff from choosing a markup of k_1 is

$$U_1(k_1, k_2, L) = \int_{-\infty}^{\infty} [1 - G(\epsilon_1 + k_1 - k_2)]\ell(\epsilon_1 + k_1 - b_1)g(\epsilon_1)d\epsilon_1 + \int_{-\infty}^{\infty} \int_{-\infty}^{\epsilon_1+k_1-k_2} \ell(\epsilon_2 + k_2 - b_1)g(\epsilon_2)d\epsilon_2 g(\epsilon_1)d\epsilon_1.$$

As expected payoffs are continuous in markups and the space of feasible markups $[-M, M]$ is compact, an equilibrium in constant markups exists for the game between experts. Let (k_1, k_2) be the pair of markups for any such equilibrium. We claim that $k_1 \leq b_1$ is not a best response for expert 1, so $k_1 > b_1$. Differentiating w.r.t. k_1 yields

$$\begin{aligned}\frac{d}{dk_1}U_1(k_1, k_2, L) &= \int_{-\infty}^{\infty} [1 - G(\epsilon_1 + k_1 - k_2)]\ell'(\epsilon_1 + k_1 - b_1)g(\epsilon_1)d\epsilon_1 \\ &= \int_{0}^{\infty} \{[1 - G(\epsilon_1 + k_1 - k_2)]\ell'(\epsilon_1 + k_1 - b_1) \\ &\quad +[1 - G(-\epsilon_1 + k_1 - k_2)]\ell'(-\epsilon_1 + k_1 - b_1)\} \, g(\epsilon_1)d\epsilon_1,\end{aligned}$$

where we have used the symmetry of g to rewrite the integral. Now by the concavity and symmetry of ℓ, for all $\epsilon_1 > 0$ and $k_1 \leq b_1$, $\ell'(-\epsilon_1 + k_1 - b_1) > 0$ and $\ell'(-\epsilon_1 + k_1 - b_1) \geq -\ell'(\epsilon_1 + k_1 - b_1)$. Clearly, if $\ell'(\epsilon_1 + k_1 - b_1) \geq 0$ (which holds for ϵ_1 in a positive-measure set), then the integrand above is strictly positive. And if $\ell'(\epsilon_1 + k_1 - b_1) < 0$, we have

$$\begin{aligned}
&\{[1 - G(\epsilon_1 + k_1 - k_2)]\ell'(\epsilon_1 + k_1 - b_1) + [1 - G(-\epsilon_1 + k_1 - k_2)]\ell'(-\epsilon_1 + k_1 - b_1)\} \\
&\quad > \{[1 - G(-\epsilon_1 + k_1 - k_2)]\ell'(\epsilon_1 + k_1 - b_1) + [1 - G(-\epsilon_1 + k_1 - k_2)]\ell'(-\epsilon_1 + k_1 - b_1)\} \\
&\quad = [1 - G(-\epsilon_1 + k_1 - k_2)]\,\{\ell'(\epsilon_1 + k_1 - b_1) + \ell'(-\epsilon_1 + k_1 - b_1)\} \\
&\quad \geq 0.
\end{aligned}$$

Hence, $\frac{d}{dk_1}U_1(k_1, k_2, L) > 0$ whenever $k_1 \leq b_1$; and by continuity, there exists $\delta > 0$ such that $\frac{d}{dk_1}U_1(k_1, k_2, L) > 0$ whenever $k_1 \leq b_1 + \delta$. It follows that if k_1 is a best response, then $k_1 > b_1$, and by an analogous argument, $k_2 > b_2$. We conclude that there exists a min equilibrium of the game and it satisfies $k_i > b_i$, $i = 1, 2$. □

References

Alonso, R. and N. Matouschek (2008). Optimal delegation. *The Review of Economic Studies 75*, 259–293.

Ambrus, A. and A. Kolb (2020). On defining ex ante payoffs in games with diffuse prior. *forthcoming in Economic Theory*.

Ambrus, A. and S. E. Lu (2014). Almost fully revealing cheap talk with imperfectly informed senders. *Games and Economic Behavior 88*, 174–189.

Ambrus, A. and S. Takahashi (2008). Multi-sender cheap talk with restricted state spaces. *Theoretical Economics 3*, 1–27.

Argenziano, R., S. Severinov, and F. Squintani (2016). Strategic information acquisition and transmission. *American Economic Journal: Microeconomics 8*(3), 119–55.

Austen-Smith, D. (1993). Interested experts and policy advice: multiple referrals under open rule. *Games and Economic Behavior 5*, 3–43.

Battaglini, M. (2002). Multiple referrals and multidimensional cheap talk. *Econometrica 70*, 1379–1401.

Battaglini, M. (2004). Policy advice with imperfectly informed experts. *Advances in Theoretical Economics 4*, 1–32.

Bernhardt, D., M. Campello, and E. Kutsoati (2006). Who herds? *Journal of Financial Economics 80*, 657–675.

Bhattacharya, S. and A. Mukherjee (2013). Strategic information revelation when experts compete to influence. *The RAND Journal of Economics 44*, 522–544.

Boucheron, S., M. Thomas, et al. (2012). Concentration inequalities for order statistics. *Electronic Communications in Probability 17*.

Callander, S. (2011). Searching for good policies. *American Political Science Review 105*, 643–662.

Calvert, R. L. (1985). Robustness of the multidimensional voting model: Candidate motivations, uncertainty, and convergence. *American Journal of Political Science 29*, 69–95.

Chen, Q. and W. Jiang (2006). Analysts' weighting of private and public information. *Review of Financial Studies 19*, 319–355.

Crawford, V. and J. Sobel (1982). Strategic information transmission. *Econometrica: Journal of the Econometric Society 50*, 1431–1451.

Deimen, I. and D. Szalay (2018). Delegated expertise, authority, and communication.

Dessein, W. (2002). Authority and communication in organizations. *The Review of Economic Studies 69*, 811–838.

Dewatripont, M. and J. Tirole (1999). Advocates. *Journal of Political Economy 107*, 1–39.

Di Pei, H. (2015). Communication with endogenous information acquisition. *Journal of Economic Theory 160*, 132–149.

Downs, A. (1957). *An economic theory of democracy.* New York: Harper & Row.

Ederer, F., R. Holden, and M. A. Meyer (2014). Gaming and strategic opacity in incentive provision. Working paper.

Filipovich, D. (2008). Cheap talk on the circle. *Report, CEE, El Colegio de Mexico*.

Friedman, D. (1991). A simple testable model of double auction markets. *Journal of Economic Behavior and Organization 15*, 47–70.

Gentzkow, M. and E. Kamenica (2017). Competition in persuasion. *The Review of Economic Studies 84*, 300–322.

Gerardi, D., R. McLean, and A. Postlewaite (2009). Aggregation of expert opinions. *Games and Economic Behavior 65*, 339–371.

Gibbons, R. (1988). Learning in equilibrium models of arbitration. *The American Economic Review*, 896–912.

Gilligan, T. and K. Krehbiel (1989). Asymmetric information and legislative rules with a heterogeneous committee. *American Journal of Political Science 33*, 459–490.

Gratton, G. (2014). Pandering and electoral competition. *Games and Economic Behavior 84*, 163–179.

Heidhues, P. and J. Lagerlöf (2003). Hiding information in electoral competition. *Games and Economic Behavior 42*, 48–74.

Holmström, B. (1977). *On incentives and control in organizations*. Stanford University.

Hotelling, H. (1929). Stability in competition. *The Economic Journal 39*, 41–57.

Iezzoni, L., S. Rao, C. DesRoches, C. Vogeli, and E. Campbell (2012). Survey shows that at least some physicians are not always open or honest with patients. *Health Affairs 31*, 383–391.

Kamenica, E. (2008). Contextual inference in markets: On the informational content of product lines. *The American Economic Review 98*, 2127–2149.

Kartik, N., F. Squintani, and K. Tinn (2015). Information revelation and pandering in elections. Technical report, mimeo, Columbia University.

Klemperer, P. (1999). Auction theory: a guide to the literature. *Journal of Economic Surveys 13*, 227–286.

Krishna, V. and J. Morgan (2001a). Asymmetric information and legislative rules: Some ammendments. *American Political Science Review 95*, 435–452.

Krishna, V. and J. Morgan (2001b). A model of expertise. *Quarterly Journal of Economics 116*, 747–775.

Krishna, V. and J. Morgan (2008). Contracting for information under imperfect commitment. *The RAND Journal of Economics 39*, 905–925.

Laslier, J.-F. and K. Van der Straeten (2004). Electoral competition under imperfect information. *Economic Theory 24*, 419–446.

Li, H. and W. Suen (2004). Delegating decisions to experts. *Journal of Political Economy 112*, 311–335.

Li, H. and W. Suen (2009). Viewpoint: Decision-making in committees. *Canadian Journal of Economics 42*, 359–392.

Loertscher, S. et al. (2012). Location choice and information transmission. Technical report, mimeo, University of Melbourne.

Martinelli, C. (2001). Elections with privately informed parties and voters. *Public Choice 108*, 147–167.

Martinelli, C. and A. Matsui (2002). Policy reversals and electoral competition with privately informed parties. *Journal of Public Economic Theory 4*, 39–61.

Melumad, N. and T. Shibano (1991). Communication in settings with no transfers. *The RAND Journal of Economics 22*, 173–198.

Migrow, D. and F. Squintani (2018). Multi-agent information acquisition and sharing.

Milgrom, P. and J. Roberts (1986). Relying on the information of interested parties. *The RAND Journal of Economics 17*, 18–32.

Morris, S. and H. Shin (2002). The social value of private information. *The American Economic Review 92*, 1521–1534.

Morris, S. and H. Shin (2003). Global games: Theory and applications. Econometrics Society Monographs.

Myatt, D. and C. Wallace (2014). Central bank communication design in a Lucas-Phelps economy. *Journal of Monetary Economics 63*, 64–79.

Mylovanov, T. and A. Zapechelnyuk (2013). Optimal arbitration. *International Economic Review 54*(3), 769–785.

Nadarajah, S. and S. Kotz (2008). Exact distribution of the max/min of two gaussian random variables. *IEEE Transactions on very large scale integration (VLSI) systems 16*(2), 210–212.

Ottaviani, M. and P. Sørensen (2006). Professional advice. *Journal of Economic Theory 126*, 120–142.

Pesendorfer, W. and A. Wolinsky (2003). Second opinions and price competition: Inefficiency in the market for expert advice. *The Review of Economic Studies 70*, 417–437.

Pétry, F. and B. Collette (2009). Measuring how political parties keep their promises: A positive perspective from political science. Volume 15 of *Studies in Public Choice*, pp. 65–80.

Prendergast, C. (1993). A theory of "yes men". *American Economic Review 83*, 757–770.

Sampford, M. (1953). Some inequalities on Mill's ratio and related functions. *The Annals of Mathematical Statistics 24*, 130–132.

Schultz, C. (1996). Polarization and inefficient policies. *The Review of Economic Studies 63*, 331–344.

Shin, H. (1998). Adversarial and inquisitorial procedures in arbitration. *The RAND Journal of Economics 29*, 378–405.

Treitel, G. (1999). *The law of contract* (10 ed.). Sweet & Maxwell.

Wittman, D. (1983). Candidate motivation: A synthesis of alternative theories. *American Political Science Review 77*, 142–157.

Wolinsky, A. (2002). Eliciting information from multiple experts. *Games and Economic Behavior 41*, 141–160.

Zitzewitz, E. (2001). Measuring herding and exaggeration by equity analysts and other opinion sellers. Technical report, mimeo, Stanford University.