Reputation and Information Aggregation
Supplemental Appendix
(for online publication)

I. Proofs for Section 4

Proof of Proposition 3. By inspection of ΔIU_0 in the proof of Lemma 4, it is easy to observe that the difference in expected instrumental utility between asking and not asking for signal-type 0 increases when p decreases.

For reputation, suppose first that, as p decreases, $\widehat{S}$ does not change. The difference in expected reputation between asking and not asking for signal-type 0, ΔR_0, reads:

$$\Pr(\widehat{S}|\omega = 0)\Pr(\omega = 0|\sigma = 0)(w - x) + \Pr(\widehat{S}|\omega = 1)\Pr(\omega = 1|\sigma = 0)(v - y).$$

As (i) only $\Pr(\omega|\sigma = 0)$ depends on p, (ii) $\Pr(\omega = 0|\sigma = 0)$ decreases as p decreases, and (iii) $w(\mu) - x < 0 < v(\mu) - y$, for a given μ, ΔR_0 increases as p decreases. Moreover, it is straightforward to observe that $\overline{\mu}$ increases as p decreases. By Lemma 3, part (i), an increase of μ induces an increase in expected reputation of signal-type 0 after asking. Thus, the difference in the overall expected payoff of signal-type 0 between asking and not asking under $\overline{\mu}$ and $\mu = 0$ increases as p decreases. Then, since the difference in expected instrumental utility is positive by A1,[1] for signal-type 0 to remain indifferent between asking and not asking as p decreases, $\widehat{\rho}$, and $\underline{\rho}$ must increase.

Consider now a change in $\widehat{S}$ as p marginally decreases. Namely, suppose that for some $k \leq n$ and each vector of advices s with $o(s) = k$, some signal-type σ switches from

[1] Hence, when signal-type 0 is indifferent between asking and not asking, the difference in expected reputation is negative.

considering $\omega = 0$ to considering $\omega = 1$ more likely. When $\sigma = 0$, were she to still decide 0, the reasoning for the case in which $\widehat{S}$ does not change would hold. By switching to $d = 1$, she improves her expected payoff after asking. When $\sigma = 1$, this means that, after s, signal-type 1 considers $\omega = 0$ and $\omega = 1$ equally likely. Then, if the prior is updated with s but not with σ, $\omega = 0$ results more likely than $\omega = 1$. Thus, given s *and* $\sigma = 0$, the decision-maker prefers to be perceived as signal-type 0 rather than pooling with signal-type 1 on $d = 0$. This observation is equivalent to Lemma 3, part (ii), as the probability of $\omega = 0$ conditional on s but not on σ is higher than 1/2 like the prior p. Hence, the switch of signal-type 1 to $d = 1$ increases the expected reputation of signal-type 0 after s. Thus, a change in $\widehat{S}$ may only increase the difference in the expected payoff of signal-type 0 between asking and not asking, and this makes $\underline{\rho}$ and $\widehat{\rho}$ increase even further. ■

Proof of Proposition 4. From Section 3.3, $\overline{\mu}$ is defined by

$$\Pr(\omega = 0|m^1)|_{\mu=\overline{\mu}} = rh + (1-r)l. \tag{SB}$$

An increase in $rh + (1-r)l$ allows an increase in $\Pr(\omega = 0|m^1)|_{\mu=\overline{\mu}}$, hence an increase in $\overline{\mu}$.

To see the effect of an increase in the competence of the decision-maker ($\gamma = qg + (1-q)b$), write $\Pr(\omega = 0|m^1)$ as

$$\frac{[\Pr(m^1|\sigma = 0)\Pr(\sigma = 0|\omega = 0) + \Pr(m^1|\sigma = 1)\Pr(\sigma = 1|\omega = 0)]\Pr(\omega = 0)}{num. + [\Pr(m^1|\sigma = 0)\Pr(\sigma = 0|\omega = 1) + \Pr(m^1|\sigma = 1)\Pr(\sigma = 1|\omega = 1)]\Pr(\omega = 1)}.$$

So we get

$$\Pr(\omega = 0|m^1)|_{\mu=\overline{\mu}} = \frac{(\overline{\mu}\gamma + 1 - \gamma)p}{(\overline{\mu}\gamma + 1 - \gamma)p + (\overline{\mu}(1-\gamma) + \gamma)(1-p)}.$$

As γ goes up, since $\overline{\mu} < 1$ by A2, $(\overline{\mu}\gamma + 1 - \gamma)$ goes down and $(\overline{\mu}(1-\gamma) + \gamma)$ goes up. (Thus, as expected, $\Pr(\omega = 0|m^1)|_{\mu=\overline{\mu}}$ goes down.) Then, to restore equality (SB), $\overline{\mu}$ must go up, so that, by $\gamma > 1/2$ (informative signals), $(\overline{\mu}\gamma + 1 - \gamma)$ increases more than $(\overline{\mu}(1-\gamma) + \gamma)$. ■

Proof of Proposition 5. Higher prior competence of the advisors affects the decision-maker's asking/not asking incentives in two ways. The most straightforward effect is a higher incentive to ask for advice due to more valuable advisors' information.

The less obvious effect is a possible discontinuous decrease in the expected reputation of signal-type 0 from asking (hence, a lower incentive to ask). It can arise because, with higher advisors' competence, there is a lower chance for signal-type 0 to separate and reveal her signal *after* asking (for instance, in the extreme case of the advisors receiving perfect signals, both signal-types will always take the same decision after asking). Suppose $\Pr(\omega = 1|\sigma = 1) > 1/2$ and consider a situation in which a certain profile of advices moves signal-type 1's belief closer to $1/2$, so that she considers $\omega = 1$ just marginally more likely than $\omega = 0$. Then, under this profile of advices, the two signal-types separate with the decision, but a marginal increase in the prior quality of the advisors will make $\omega = 0$ more likely for signal-type 1 and, thus, will make her switch to $d = 0$. This induces a discrete fall in signal-type 0's expected reputation after asking, by the same argument as in the proof of Proposition 3.

Consider $\rho = \underline{\rho}$ and suppose that a marginal improvement in the prior quality of advisors does not cause the second effect. Then such an improvement makes signal-type 0 strictly prefer to ask, which is going to destroy the advisors' truthtelling.[2] Effectively, $\underline{\rho}$ moves up, and at the initial value of $\underline{\rho}$ there is no equilibrium with information aggregation (under our equilibrium selection at the advising and decision stages). In this case, higher advisors' competence harms through provoking excessive advice-seeking.

Consider now $\rho = \widehat{\rho}$ and suppose we are exactly at the point where a marginal increase in the advisors' competence is going to cause the second effect. To keep signal-type 0 indifferent between asking and not asking, the discontinuous decrease in her incentives to ask needs to be compensated by a discrete increase in μ. However, μ cannot go above $\overline{\mu}$ without destroying advisors' truthtelling. Consequently, the second-best equilibrium as well as any hypothetical equilibrium sufficiently close to the second best (in terms of the probabilities of asking) will be destroyed. Thus, the fall in information aggregation will be discontinuous. In such a case, a higher advisors' competence harms through provoking excessive advice-avoidance. ■

[2]Better advisors' competence also widens the set of beliefs about the state for which they report truthfully. However, since we consider a marginal improvement in competence, this change will be marginal, whereas the change in the asking/not asking behavior of signal-type 0 is discrete (and actually extreme), meaning a discrete jump in $\Pr(\omega = 0|m^1)$. Hence, truthful reporting will fail.

II. Robustness

II.A. Asking a subset of advisors

Suppose the decision-maker could choose to ask any subset of advisors. Assume this choice is observed by everyone (we address secret advice-seeking below). First, our separating and partially separating equilibria of the baseline model survive. This is ensured by the off-the-path belief that asking a proper subset of advisors (rather than all advisors) implies that the decision-maker has received $\sigma = 0$, thus resulting in no truthtelling, by A2.

There can be other equilibria, but the crucial thing is that signal-type 0 cannot ask a subset of advisors different from the one approached by signal-type 1 and receive informative advice at the same time: in any such equilibrium, she will be recognized and, hence, provided with no information. Thus, all other equilibria look qualitatively similar to those of the baseline model, with the full set of advisors being substituted by a proper subset.[3] Obviously, no such equilibrium can do better than our second best.

It should be noted that asking a proper subset of advisors extends the set of ρ where some information aggregation is possible: Lowering the number of advisors that are asked reduces the incentive of signal-type 0 to ask and, thus, lowers $\underline{\rho}$. Thus, once we go below $\underline{\rho}$ of the baseline model, we can still sustain *some* information aggregation by reducing the equilibrium subset of asked advisors. However, it is rather obvious that, as ρ moves down further, information aggregation eventually deteriorates due a lower and lower number of advisors being asked.[4] Hence, our qualitative result that a too low weight on reputation is detrimental to information aggregation still holds.

II.B. Impossibility of not asking

In some real-life contexts, it could be impossible to prevent an advisor from expressing his opinion by not asking. In such cases, "not asking" essentially becomes unfeasible, and the decision-maker can only make non-verifiable statements about her signal prior

[3] Deviations to asking a subset of advisors different from the equilibrium one are ruled out by picking appropriate off-the-path beliefs. For example, we can impose that asking any non-equilibrium subset makes the advisors believe that $\sigma = 0$, which results in herding.

[4] The first statement of Proposition 2 is applicable here, as any equilibrium in which a proper subset of advisors is asked (and provides some information), is equivalent to an equilibrium in which all advisors are asked but some of them babble.

to receiving advice. Hence, instead of "asking"and "not asking", let us interpret m^1 and m^0 as such statement.

Clearly, in this modified game, our separating and partially separating equilibria exist and have the same characteristics as in the baseline model for the same values of ρ: If only signal-type 0 sends message m^0, the advisors herd due to A2, thus making m^0 equivalent to not asking.

More generally, *all* equilibria of the original game (i.e., even those in which signal-type 1 does not always ask for advice) have identical counterparts in the modified setting. The reason is that at least one of the two messages necessarily leads to herding. To see this, notice that if neither message leads to herding, then it must be that both $\Pr(m^1|\sigma=0)/\Pr(m^1|\sigma=1) \leq \overline{\mu}$ and $\Pr(m^0|\sigma=0)/\Pr(m^0|\sigma=1) \leq \overline{\mu}$.[5] However, this is impossible because $\overline{\mu} < 1$. Then, a message that triggers herding can just be replaced by "not asking".

II.C. Asking accompanied by statements

Suppose that the decision-maker can make a non-verifiable statement about her signal after asking for advice. We argue that this additional cheap talk stage does not change the results of the model.

Consider an equilibrium of the modified game in which both signal-types ask with positive probability, make different and informative statements τ' and τ'', and both statements trigger truthful reporting by the advisors (otherwise one would be clearly equivalent to not asking)

First, for truthful reporting, it must be the case that $\Pr(m^1,\tau|\sigma=0)/\Pr(m^1,\tau|\sigma=1) < 1$ for $\tau=\tau',\tau''$. Hence, signal-type 0 does not always ask. Since the two statements are different and signal-type 0 makes both less frequently than signal-type 1, by Lemma 5 from the main Appendix signal-type 0 strictly prefers and makes only one of the two statements, say τ'. Then, signal-type 1 would strictly prefer τ' to τ'' if she would consider state 0 more likely. Since sometimes she states τ'', it must be that she considers state 1 more likely. Thus, signal-type 0 is perceived as such after not asking (for signal-type 1 decides 1). Moreover, since she plays both τ' and not asking, she must be indifferent be-

[5] It is straightforward to show that the advisors' incentives after message m, i.e., the left-hand side of $\Pr(\omega=0|m) \leq rh+(1-r)l$, are determined by the ratio $\Pr(m|\sigma=0)/\Pr(m|\sigma=1)$.

tween the two. But then, since expected reputation depends only on relative probabilities, there also exists (and aggregates more information) our partially separating equilibrium, where signal-type 0 asks with probability $\Pr(m^1, \tau'|\sigma = 0)/\Pr(m^1, \tau'|\sigma = 1)$ (instead of $\Pr(m^1, \tau'|\sigma = 0)$ like here).

II.D. Secret asking

Suppose now that the decision-maker was given the additional opportunity to ask for advice without being observed by the observer. Obviously, this can happen only when the observer is not the advisors.

It is straightforward to note that the separating and partially separating equilibria of the baseline model have equivalent counterparts in the modified game: We just need to impose the advisors' belief that any asking behavior except asking publicly all advisors implies that the decision-maker has received $\sigma = 0$. Then, secret asking generates herding by advisors, and, thus, it is equivalent to not asking.

Furthermore, no equilibrium in the modified game can improve over the second-best. For such an improvement to materialize it must be that signal-type 0 asks for advice with the aggregate probability higher than $\overline{\mu}$, i.e., $\Pr(m^1_{secret}|\sigma = 0) + \Pr(m^1_{public}|\sigma = 0) > \overline{\mu}$. This means that either $\Pr(m^1_{secret}|\sigma = 0)/\Pr(m^1_{secret}|\sigma = 1) > \overline{\mu}$ or $\Pr(m^1_{public}|\sigma = 0)/\Pr(m^1_{public}|\sigma = 1) > \overline{\mu}$. This, in turn, implies that either after secret or after public advice seeking, the advisors will herd. Suppose this happens after public asking. Then, for the secret asking to trigger truthful reporting it must be that $\Pr(m^1_{secret}|\sigma = 0)/\Pr(m^1_{secret}|\sigma = 1) \leq \overline{\mu}$, which implies $\Pr(m^1_{secret}|\sigma = 0) \leq \overline{\mu}$. Hence, signal-type 0 will receive informative advice with probability weakly below the second-best one.

II.E. Publicly unobservable advice

A separate issue is observability of the advisors' messages by the external observer. This issue is irrelevant for the behavior of the advisors, as they only care about their reputation in the eyes of the decision-maker (we discuss what happens if they have other concerns in subsection II.I below). As for the decision-maker, making the advisors' messages unobservable by the external observer would generally affect her incentives. This is because the decision is affected by advice, and, therefore, the observer's inference about the decision-maker's signal is affected by information on both the decision and the advice.

Absent any information about the advice, for each signal-type of the decision-maker, the expected reputation after each decision will be independent of the received advice, and there will be just a threshold number of truthfully reported 0 signals above which $d = 0$ will be taken. This may not result in always taking the decision that corresponds to the more likely state, given the advice and the signal-type. However, this does not modify qualitatively the asking/not asking trade-off faced by the signal-types: Information will still be relatively more valuable to signal-type 1, whereas, reputation-wise, being perceived as signal-type 0 will be relatively more attractive for signal-type 0 (from the pre-asking perspective). Therefore, our equilibria would not qualitatively change.

II.F. Sequential public advice

First of all, notice that in our setup, for a given advisors' belief conditional on being asked, sequential public advice always provides the decision-maker with less information. If the advisors herd under simultaneous advice, so will they under sequential advice starting from the first speaker. At the same time, if the advisors tell the truth under simultaneous advice, they will still start herding under sequential advice once the number of messages in one direction exceeds that in the other direction by one or two (depending on the direction of messages).

Thus, if the choice of the advice scheme (sequential versus simultaneous) is part of the game, then the conclusions we reached in the discussion of asking a subset of advisors apply here as well (in particular, all baseline model equilibria survive) If, in contrast, sequentiality of advice is exogenous, our results still stay qualitatively intact: Although sequential advice is less informative, the fundamental trade-off between reputation and receiving information remains, generating the familiar types of equilibria.

II.G. Privately known advisors' types

First of all, what is crucial for our story is the distortion of the advisors' incentives when the confidence about the state rises. Although in our model this distortion arises due to reputation concerns, costly information acquisition by advisors would generate a similar effect (we elaborate more on that in subsection II.I below), even when they know their types.

Second, while unawareness of an advisor about his type may be an extreme assumption, full awareness is equally extreme. Presumably, an advisor could learn his type through experience, i.e., by assessing correctness of his signals in the past. However such learning is limited: Even for good advisors signals are never perfectly precise and, moreover, advisors may not always receive accurate ex-post information on whether their signals matched the state.

Finally, even under the assumption that the advisors know their types, herd behavior does not fully disappear. By Lemma 4 of Ottaviani and Sørensen (2001), low types still herd with positive probability whenever $\Pr(\overline{\omega}|m^1) > l$, where $\overline{\omega}$ is the more likely state conditional on being asked. Therefore, the problem of "excessive asking", though becoming less severe, remains relevant. Hence, having a (moderately) reputation-concerned decision-maker remains beneficial, similarly to the baseline model.

II.H. Privately known decision-maker's type

Assume the decision-maker knows her competence-type. There will be now four privately known competence-signal-types (call them just "types"), as each of the competence-types $\{G, B\}$ can receive either $\sigma = 0$ or $\sigma = 1$: $G0$, $G1$, $B0$, $B1$.

The first thing to notice is that Proposition 1 qualitatively holds. If $\rho = 0$ or is sufficiently small, all types will be tempted to ask. Consequently, the advisors will herd.

Intuitively, all informative equilibria will, roughly speaking, have the following feature: signal-types 0 will refrain from asking more often than signal-types 1, similarly to the baseline model. Let us focus, for simplicity, on equilibria in pure strategies. As an example, consider the following equilibrium: $G0$ does not ask for advice, while $G1$, $B0$ and $B1$ ask, and the advisors report truthfully.[6] Such an equilibrium must exist for a range of parameters. Provided that p it is neither too high nor too low relative to the

[6] Another possible equilibrium is the one in which $G0$ and $B0$ always refrain from asking, while $G1$ and $B1$ always ask and receive truthful advice. From the point of view of the advisors, the strategy of the decision-maker conveys the same information as in the separating equilibrium of the baseline model. From the point of view of the decision-maker, since asking and not asking are unable to signal the competence-type directly, the trade-off is qualitatively the same as in the baseline model, and it is solved through a similar single-crossing argument. (Of course, since the competence-types are privately known, the relevant incentive compatibility constraints will not be exactly the same as in the baseline model). Provided that $G1$ does not have a too strong belief in $\omega = 1$, she will prefer pooling with $B1$ instead of not asking and choosing $d = 1$.

Finally, when p is sufficiently close to 1/2, there can potentially exist equilibria in which G-types never ask and take the decisions corresponding to their signals (the behavior of B-types is likely to vary depending on the equilibrium).

precision of the good competence-type, g, the advisors' belief after being asked will be sufficiently close to 1/2 so that (1) holds (if the proportion of bad competence-types is high enough, then p should just not be too high).

Thus we will have the familiar trade-off between having a higher instrumental utility from asking and higher reputational payoff from not asking. Since $G0$ is most confident about the state of all types, her expected instrumental utility from asking is smaller compared to the other three. For simplicity (not crucial), we can assume that not asking followed by $d = 1$ yields an (off-the-path) belief that the decision-maker is $G1$. Then not asking always yields the belief that $\theta = G$. Then, naturally, there will be thresholds $\underline{\rho}'$ and $\overline{\rho}'$ such the that the equilibrium under consideration exists if and only if $\rho \in [\underline{\rho}', \overline{\rho}']$ Threshold $\underline{\rho}'$ will be determined by the incentive compatibility of $G0$: when $\rho < \underline{\rho}'$, the reputation concerns are so low that $G0$ will want to deviate to asking for advice. Threshold $\overline{\rho}'$ will be determined by the incentive compatibility of either $B0$ or $G1$ (the deviation incentive of $B1$ is obviously weaker than that of $B0$): when $\rho > \overline{\rho}'$, high reputation concerns will make either of these types deviate to not asking.

As p grows, type $G0$ becomes more confident about the state and, thus, less tempted to ask for advice. Therefore, a lower level of reputation concerns becomes enough for her to refrain from asking, i.e., $\underline{\rho}'$ decreases. Type $B0$ also becomes more confident that $\omega = 0$, which makes her less willing to ask. Consequently, a lower level of reputation concerns is needed to keep $B0$ asking. If $\overline{\rho}'$ is determined by the incentive compatibility of $B0$, this means that $\overline{\rho}'$ goes down. If $\overline{\rho}'$ is determined by the incentive compatibility of $G1$, it must be that $G1$ believes that $\omega = 1$ is more likely. Then, a higher p results in higher willingness to ask by $G1$, meaning an increase in $\overline{\rho}'$. It is clear, however, that at some point $\overline{\rho}'$ becomes determined by the incentive compatibility of $B0$, and, thus, eventually goes down. Hence, comparative statics with respect to p remains qualitatively similar to that in the baseline model.

II.I. Advisors' other incentives

Our setup can be modified to allow an advisor to care about the quality of decisions in addition to reputation. The optimal weight of the advisors' reputation concerns would then be as small as possible, to maximize their truthtelling incentives. However, in reality, it is hardly possible to eliminate the reputation concerns altogether. Therefore, the herd behavior would still be a problem (albeit for a smaller set of beliefs), and all our qualitative results would survive.

In addition to reputation in the eyes of the decision-maker, an advisor may care about his reputation in front of other people. If advice is public, this is inconsequential for our model. In contrast, if the advisors' messages are observed only by the decision-maker, such extra reputation concerns may help truthtelling indirectly, through the incentive to reduce the probability of wrong decisions. However, provided that *some* concerns for reputation in the eyes of the decision-maker remain, the argument in the previous paragraph applies here as well.

A key ingredient of our story is that the advisors are willing to provide information only when they feel uncertain about the state of nature. Apart from reputation concerns, there may be other reasons that generate a similar incentive. For example, assume that advisors have no reputation concerns and care about the quality of decisions, but need to incur a cost of acquiring (or transmitting) a signal. Then their incentives to acquire information will be stronger (and hence the quality of information received by the decision-maker will be higher) the more undecided they think the decision-maker is. Consequently, like in our baseline model, it will be crucial to avoid "excessive asking" by a decision-maker with the signal confirming the prior. At the same time, the temptation to ask for advice increases in the prior uncertainty. Thus, such a framework would generate the same qualitative results as the current one. We now formalize these arguments.

Costly information acquisition by advisors. Formalization

Consider an alternative setup in which the advisors have no reputation concerns and care about the quality of decisions, but need to incur a cost of acquiring a signal. Assume, for simplicity, there is only one advisor, whose ex-post payoff is

$$\begin{cases} 1 \text{ if } d = \omega; \\ 0, \text{ if } d \neq \omega. \end{cases}$$

At stage 2 the advisor has no signal. At stage 4 (if asked for advice at stage 3) by incurring a fixed cost c the advisor can acquire a binary signal $s \in \{0, 1\}$ of precision $\alpha := \Pr(s = \omega)$ for any ω. If he does not invest in information acquisition, no signal is acquired. The advisor then reports whether he invested in information and which signal he received. For the rest, the model is the same as in our baseline setup.

Note that there is absolutely no reason for the advisor to lie, so let us assume he will always reveal the truth.

Denote $\gamma := qg + (1-q)b \equiv \Pr(\sigma = \omega)$ – the ex-ante expected precision of the decision-maker's signal.

To avoid uninteresting cases, let us make the following two assumptions. First, assume

$$\Pr(\omega = s|\sigma, s) > 1/2 \text{ for any } \sigma \text{ and } s.$$

That is, like in our baseline model (assumption A1), the advice is always potentially useful for both signal-types. The decision-maker, thus, will *always* follow the advice. This means that, regardless of the state, the decision will be correct with probability α – the expected quality of the advisor's signal. Thus, if the advisor decides to invest in information acquisition, his expected payoff (after he is asked for advice) will be $\alpha - c$.

Second, assume that absent any information about the decision-maker's signal, the advisor prefers not to acquire information (this is analogous to the first part of A2). If the advisor acquires no information, then the following happens. If $\Pr(\omega = 0|\sigma = 1) > 1/2$, then both signal-types take $d = 0$, and the advisor's payoff is p. If $\Pr(\omega = 0|\sigma = 1) \leq 1/2$, then the decision-maker takes $d = \sigma$, and the advisor's payoff is γ. It is straightforward to show that $\Pr(\omega = 0|\sigma = 1) > 1/2$ is equivalent to $p > \gamma$. Thus, the necessary and sufficient condition for no information acquisition (assuming information acquisition in case of indifference) is

$$a - c < \max\{\gamma, p\}.$$

Finally, assume that learning that $\sigma = 1$ always results in information acquisition (this is analogous to the second part of A2). If the advisor learns that $\sigma = 1$ and acquires no information, the decision-maker will take $d = \overline{\omega}$, where $\overline{\omega}$ is the more likely state, given $\sigma = 1$. The advisor's payoff will then be $\Pr(\omega = \overline{\omega}|\sigma = 1) \equiv \max\{\Pr(\omega = 0|\sigma = 1), \Pr(\omega = 1|\sigma = 1)\}$. Thus, the advisor will acquire information after learning that $\sigma = 1$ (assuming information acquisition in case of indifference) if and only if

$$\alpha - c \geq \max\{\Pr(\omega = 0|\sigma = 1), \Pr(\omega = 1|\sigma = 1)\} \equiv \frac{\max\{\gamma, p\} - \gamma p}{(1-\gamma)p + \gamma(1-p)}$$

Let us first analyze the second-best solution. Like in the baseline model, signal-type 1 must always ask in the second best, as more asking by signal-type 1 never hurts the advisor's incentives. Indeed, it raises the advisor's belief that the decision-maker is less confident about the state (as signal-type 1 is always less confident), which can only encourage information acquisition.

Let μ be the probability that signal-type 0 asks for advice.

Consider first the case $p > \gamma$, that is, the case in which signal-type 1 believes that $\omega = 0$ is more likely. If the advisor does not acquire information, the decision-maker will always take $d = 0$. Thus, the advisor's expected payoff will be $\Pr(\omega = 0|m^1)$, and he will invest in information if and only if

$$\alpha - c \geq \Pr(\omega = 0|m^1).$$

Clearly, $\Pr(\omega = 0|m^1) < p$ for all $\mu < 1$, increases with μ and reaches p at $\mu = 1$. The second-best μ, i.e., $\overline{\mu}$, is determined by $\alpha - c = \Pr(\omega = 0|m^1)$; it exists thanks to our assumptions that $a - c < \max\{\gamma, p\}$ and $\alpha - c \geq \Pr(\omega = 0|\sigma = 1)$. As we decrease p, $\Pr(\omega = 0|m^1)$ declines for given μ. Hence, $\overline{\mu}$ has to go up to keep the equality satisfied.

If $p \leq \gamma$, signal-type 1 believes that $\omega = 1$ is weakly more likely. If the advisor does not acquire information, the decision-maker will always take $d = \sigma$. Thus the advisor's expected payoff will be $\Pr(\omega = \sigma|m^1)$, and he will invest in information if and only if

$$\alpha - c \geq \Pr(\omega = \sigma|m^1).$$

The right-hand side can be rewritten as $\Pr(\omega = 0|m^1)\Pr(\sigma = 0|\omega = 0, m^1) + \Pr(\omega = 1|m^1)\Pr(\sigma = 1|\omega = 1, m^1)$, which, after some algebra, yields

$$\frac{\gamma(p\mu + 1 - p)}{\gamma(p\mu + 1 - p) + (1 - \gamma)[p + (1 - p)\mu]} \equiv \frac{1}{1 + \dfrac{(1 - \gamma)}{\gamma}\dfrac{[p + (1 - p)\mu]}{1 - p + p\mu}}$$

It is easy to show that this expression is increasing in μ (since $p > 1/2$) and reaches γ when $\mu = 1$. Hence $\overline{\mu}$ is determined by $\alpha - c = \Pr(\omega = \sigma|m^1)$; it exists thanks to our assumptions that $a - c < \max\{\gamma, p\}$ and $\alpha - c \geq \Pr(\omega = 1|\sigma = 1)$ (it is easy to show that $\Pr(\omega = \sigma|m^1) \geq \Pr(\omega = 1|\sigma = 1)$). Furthermore, the expression goes up as p declines, meaning that $\overline{\mu}$ goes up with a decrease in p.

Note finally that $\Pr(\omega = 0|m^1) = \Pr(\omega = \sigma|m^1)$ at $p = \gamma$, for a fixed μ, because $\Pr(\omega = 0|m^1) = \Pr(\omega = 0, \sigma = 0|m^1) + \Pr(\omega = 0, \sigma = 1|m^1)$, $\Pr(\omega = \sigma|m^1) = \Pr(\omega = 0, \sigma = 0|m^1) + \Pr(\omega = 1, \sigma = 1|m^1)$, and $\Pr(\omega = 1, \sigma = 1|m^1) = \Pr(\omega = 0, \sigma = 1|m^1)$ at $p = \gamma$ (since $\Pr(\omega = 0|\sigma = 1) = \Pr(\omega = 1|\sigma = 1)$). This implies continuity of $\overline{\mu}(p)$ at $p = \gamma$. Thus, the entire dynamics of $\overline{\mu}$ is the same as in the baseline model.

Let us turn now to the equilibrium asking/not asking behavior. For the decision-maker's incentives, it is immaterial why exactly advisors provide truthful information; it only matters whether, for a given asking/non-asking behavior, asking triggers information provision. Hence, the trade-offs of the two signal-types remain qualitatively the same as in the baseline model.

Given the dynamics of $\overline{\mu}$, the equilibrium structure, thus, also remains the same. The existence of $\underline{\rho}$, below which no information provision occurs is straightforward: Under too low reputation concerns, signal-type 0 would deviate from any (partially or fully) separating equilibrium and ask for advice. The existence of $\widehat{\rho}$ (corresponding to $\overline{\mu}$) is straightforward as well. Given that $\overline{\mu}$ is decreasing in p, we can apply the arguments of the baseline model to show that both $\underline{\rho}$ and $\widehat{\rho}$ increase with the degree of prior uncertainty.

II.J. Relaxing A2

In the baseline model, we assumed that if an advisor learned that the decision-maker received signal 1, his own signal determined which state he considered more likely. Let us drop this restriction. Let $\overline{\omega}$ be the more likely state from the perspective of an advisor conditional on being asked but ignoring the own signal. Then the necessary and sufficient condition for truthtelling by the advisors becomes

$$\Pr(\omega = \overline{\omega}|m^1) \leq rh + (1 - r)l, \tag{TR}$$

which is a generalized condition (1) from the main text.

Consider the case when, after learning that $\sigma = 1$, an advisor believes that $\omega = 1$ is more likely regardless of his own signal, which is equivalent to $\Pr(\omega = 1|\sigma = 1) > rh + (1 - r)l$. Then, some asking by signal-type 0 becomes necessary to induce truthful revelation by advisors. In other words, any equilibrium with information aggregation will necessarily be partially separating. The following proposition is true.

Proposition 6 *Consider the case when* $\Pr(\omega = 1|\sigma = 1) > rh + (1-r)l$. *A partially separating equilibrium in which signal-type* 0 *is indifferent between asking and not asking for advice, signal-type* 1 *always asks, and the advisors report their signals truthfully, exists if and only if* $\rho \in [\underline{\widehat{\rho}}, \widehat{\rho}]$, *where* $\underline{\widehat{\rho}} \in (0,1)$, $\widehat{\rho} \in (\underline{\widehat{\rho}}, 1)$. *In this equilibrium,* μ *is strictly increasing in* ρ, *ranging from some* $\underline{\mu} > 0$ *at* $\underline{\widehat{\rho}}$ *to* $\overline{\mu}$ *at* $\widehat{\rho}$.

Proof. The argument is exactly the same as in the proof of Theorem 1, with the only difference that $\mu = 0$ is no longer compatible with truthful reporting by advisors. Notice that: (1) $\Pr(\omega = 1|m^1)$ is decreasing in μ, (2) $\Pr(\omega = 1|m^1)$ equals $1 - p < rh + (1-r)l$ for $\mu = 1$ and $\Pr(\omega = 1|\sigma = 1) > rh + (1-r)l$ for $\mu = 0$. Hence, there exists a value of μ, denoted by $\underline{\mu}$, such that $\Pr(\omega = 1|m^1) = rh + (1-r)l$. This is the lowest value of μ compatible with (TR). The value of ρ making signal-type 0 indifferent between asking and not asking for $\mu = \underline{\mu}$ is denoted by $\underline{\widehat{\rho}}$. Since $\overline{\mu}$ is determined by $\Pr(\omega = 0|m^1) = rh + (1-r)l$, $\underline{\mu} < \overline{\mu}$, which implies $\underline{\widehat{\rho}} < \widehat{\rho}$. ■

It turns out that threshold $\underline{\widehat{\rho}}$ has the same properties as $\underline{\rho}$ from the baseline model:

Proposition 7 *Consider the case when* $\Pr(\omega = 1|\sigma = 1) > rh + (1-r)l$. *The following is true:*

i) Under our equilibrium selection at the advising and decision stages, there exist no equilibria with information aggregation for $\rho < \underline{\widehat{\rho}}$.

ii) The value of $\underline{\widehat{\rho}}$ *rises with an increase in the prior uncertainty about the state of nature.*

iii) $\underline{\widehat{\rho}}$ *converges to* $\underline{\rho}$ *when* $\Pr(\omega = 1|\sigma = 1)$ *approaches* $rh + (1-r)l$.

Proof. For part (i), the logic is the same as in the proof of Proposition 1. The expected reputational loss from asking is the highest possible under partial separation with $\nu \equiv \mu = \underline{\mu}$, under the constraint that the advisors report truthfully, i.e., that $\nu \geq \underline{\mu}$. Moreover, for $\rho < \underline{\widehat{\rho}}$, by Proposition 6, there is no partially separating equilibrium, because signal-type 0 would strictly prefer to ask. The two things combined bring to the same conclusion as Proposition 1.

For part (ii), notice that, similarly to $\overline{\mu}$, $\underline{\mu}$ increases as p goes down, to ensure that (TR) binds. Furthermore, at $\underline{\widehat{\rho}}$, signal-type 0 is indifferent between asking and not asking in the equilibrium under consideration. Thus, the same argument as for $\widehat{\rho}$ in the proof of Proposition 3 applies for $\widehat{\underline{\rho}}$.

Finally, as $\Pr(\omega = 1|\sigma = 1)$ approaches $hr + l(1 - r)$, $\underline{\mu}$ clearly approaches 0. Consequently, our partially separating equilibrium at $\underline{\mu}$ approaches the purely separating equilibrium, and $\widehat{\underline{\rho}}$ converges to $\underline{\rho}$. ■

Thus, all our qualitative results survive if we drop the second part of A2.