The (Human) Sampler's Curses

By Mark Thordal-Le Quement*

We present a cheap talk model in which a receiver (*R*) *sequentially consults multiple experts who are either unbiased or wish to maximize R's action, bias being unobservable. Consultation is costly and R cannot commit to future consultation behavior. We find that individual expert informativeness negatively relates to consultation extensiveness and expert trustworthiness due to biased experts' incentive to discourage further consultation by mimicking unbiased experts. We identify three* (*sampler's*) *curses: R may lose from an increase in the number or in the trustworthiness of experts as well as from a decrease in consultation costs.* (*JEL* D82, D83)

In a complex world, decision makers rely on experts for most of their information. A crucial issue is that experts can often be legitimately suspected of pursuing an own agenda and that adequate information about these outside motives is typically lacking. An important instance of this problem is the case of online product reviews. American novelist J. Franzen, discussing Amazon's book selling business, reports a claim that about one-third of posted reviews are fakes.[1] Given the anonymity of reviews, writing fake positive (negative) reviews of own (competitors') products appears like an easy way to game the system. Another instance of the problem arises in the market for credence goods where, as noted in Ely and Välimäki (2003), "*the seller first diagnoses the client's needs and then chooses a product to sell.*" Doctors, mechanics, as well as management and legal consultants are examples. These may be tempted to recommend the most expensive (and profitable) procedure rather than the most adequate one. Though a share of experts doubtlessly succumbs to this temptation, many are committed to unbiased advice.

A helpful strategy in the above situations is to consult multiple experts in the hope of encountering a truthful person at some point. This process is typically sequential. After hearing a first expert, one asks for a second opinion if still uncertain, then potentially for a third, etc. The sequential nature of the process can have different underlying motives. Information may be complex and take time to decode

* School of Economics, University of East Anglia, Norwich, United Kingdom, NR4 7TJ (e-mail: M.Le-Quement@uea.ac.uk). I am indebted to Kohei Kawamura for many stimulating discussions at early stages of this project. I thank Karl Schlag and Joel Sobel for helpful comments and suggestions on an early version focused on simultaneous consultation. Last but not least, I thank two anonymous referees for many helpful comments and suggestions.

† Go to http://dx.doi.org/10.1257/mic.20150009 to visit the article page for additional materials and author disclosure statement or to comment in the online discussion forum.

[1] "*Maybe the internet experiment in consumer reviewing will result in such flagrant corruption* (*already one third of product reviews are said to be bogus*) *that people will clamor for the return of professional reviewers*" p. 274, in *The Kraus Project*, 2013, J. Franzen, Fourth Estate, London.

(e.g., scientific reports, specialized product reviews) or experts may be located in different locations (e.g., doctors, lawyers). Even the process of *googling* is essentially a sequential consultation problem. One types a set of keywords and the search engine presents an ordered list of results that can only be examined sequentially. Consultation may also be sequential (as opposed to simultaneous) by choice, because it spares wasteful consultations whenever uncertainty is resolved early. Why order three reports if one is likely to have a clear picture after reading the first two? A final dimension of sequential consultation is the absence of a commitment ability on the part of the decision maker regarding future consultation behavior (whether or not she will ask for another opinion).

This paper evaluates key dimensions of the sequential consultation problem. A decision maker (R) faces a set of experts who all know the state ω. Each of these either shares R's preferences or wishes to maximize her action, expert preferences being unobservable. R consults experts sequentially and is unable to commit to her consultation behavior. We ask three questions: Does R necessarily gain from facing a larger pool of experts, from experts being on average more trustworthy and from consultation being less costly? These questions are relevant because the rise of the internet seems to be driving precisely such changes in parameter values thanks to the multiplication of sources of expertise (blogs, independent news websites, product reviews). We provide negative answers to all three questions. The set of instances identified and their underlying mechanisms constitute what we call the *human sampler's curses*. We speak of the *human sampler*'s curses for two reasons. First, because repeated consultation is a sampling problem in which exogenous signals are replaced by strategic experts. Second, because the decision maker is *human, all too human* in the sense that she cannot commit not to ask for another opinion whenever this is advantageous. We speak of *curses* because our insights are of a negative, sobering nature.

The sampler's curses originate in two main trade-offs. The first trade-off is that more equilibrium consultations imply a lower quality of individual communication. Accordingly, a larger pool of experts or a lower consultation cost, by leading to more equilibrium consultation and worsened individual reporting, can decrease R's expected payoff. The second trade-off is that higher average trustworthiness implies coarser individual reporting. For a fixed number of equilibrium consultations an increase in average trustworthiness can thereby cause a decrease in R's expected payoff. The source of both trade-offs lies in biased experts' incentive to discourage future consultation. This incentive increases when more experts are consulted in equilibrium or when the chance of R encountering an unbiased expert increases.

Biased senders' incentives are best seen in the two experts case. An equilibrium with multiple sequential consultations takes the form of a partitional equilibrium with N intervals, interval 1 (N) being the lowest (highest), an unbiased expert truthfully reveals the interval in which the state ω is located by sending m_i if ω is in interval i. A biased expert systematically claims that ω is in the highest interval. R in turn asks for a second opinion if the first expert consulted claims that ω is in the highest interval. By consulting again, R will receive $m_i \neq m_N$ if ω is located in interval $i < N$ and the second expert is unbiased, in which case she correctly learns that ω is located i. If ω is located in a low interval, this behavior of R, however, creates an incentive for a biased expert to deviate to sending m_{N-1} so as to discourage R from asking a second opinion. For such a downward deviation by a biased expert not to be profitable it should be costly, implying that the second highest interval $(N - 1)$ should cause low beliefs, which requires the latter to be large. By the incentive conditions of unbiased senders, this in turn implies that the third highest interval $(N - 2)$ should be large, etc.

In the case of a unique equilibrium consultation, there exists a qualitatively different *semi-revealing* scenario involving a continuum of equilibrium messages and described by a unique threshold $\theta \in (0, 1)$. An unbiased expert sends $m = \omega$ if $\omega < \theta$ and sends $m = \theta$ for any $\omega \geq \theta$. A biased expert always sends $m = \theta$. R's beliefs are monotonically increasing and continuous in m for $m \in [0, \theta]$. The distinctive feature of this scenario is that unbiased experts communicate perfectly below θ, making individual communication very informative. The scenario is incentive compatible because R does not consult again after $m = \theta$, implying that a biased sender has no incentive to deviate to $m' = \theta - \epsilon$ (for ϵ positive but small) so as to discourage a further consultation.

Literature Review.—A main strand in the cheap talk literature (e.g., Crawford and Sobel 1982; Morgan and Stocken 2003) studies one-to-one interaction and shows that preference misalignment causes noisy communication. Papers on the multiple sender case identify externalities between senders. In Krishna and Morgan (2001), if two senders have different biases and at least one of them is not an "extremist," consulting two senders is always beneficial and may induce full revelation. Battaglini (2002) studies a setup featuring multiple perfectly informed experts and a multidimensional state. Full revelation of information in all states of nature is generically possible, even when the conflict of interest is arbitrarily large. What really matters is the local behavior of senders' indifference curves at the ideal point of R rather than the proximity of players' ideal point. Battaglini (2004) studies a setup with noisy sender signals and identifies a trade-off between information aggregation and information extraction: Multiple consultation reduces idiosyncratic noise but comes at the cost of less precise communication by each sender. The key is that noisy information means that each sender's message affects R's decision on all dimensions, which in turn implies that a sender can always profitably bias R's decision by deviating in a putative truthtelling equilibrium. Kawamura (2013) studies a sampling problem with uncertainty about the preferences of senders and similarly finds a trade-off between the quality of communication and sample size. The intuition is that as the sample size increases, each sender's influence on R decreases, thus generating an incentive for each sender to exaggerate his reporting. In the limit, only binary communication is feasible. Morgan and Stocken (2008) study a polling problem and assume a binary message space, thus discarding the issue of how the fineness of the equilibrium partitioning changes with respondents' strategic incentives. R can always learn more by polling more agents and information aggregates in the limit. Li (2010) examines a simple cheap talk game in which a decision maker faces two experts whose privately known bias can be null, negative, or positive. If deciding to hear both experts, R may use a simultaneous, sequential, or hierarchical communication protocol (we omit the latter, which we find less immediately relevant). Sequential communication is modeled differently than in our model: The first expert sends a message that is observed by the second expert who then sends his message, after which the decision maker finally observes both messages simultaneously. A first result is that all two-expert mechanisms do better than the one-expert mechanism. The communication strategy of experts under multiple simultaneous consultation replicates that used given a single consultation, so that there is no trade-off between individual informativeness and the number of experts consulted. Adding more experts simply increases the likelihood of different biases and increases the probability of receiving undistorted information from one of the experts. A second finding is that simultaneous communication is mostly better than sequential communication because allowing the second expert to condition his messaging on the first expert's message reduces the informativeness of the second sender's communication.

Information elicitation from multiple experts has also been studied from a mechanism design perspective allowing R to commit to sophisticated decision rules. In Gerardi, McLean, and Postlewaite (2009), a mediator uses the correlation among signals to threaten individual experts with punishment if their report does not match that of others. In Wolinsky (2002), truthtelling is achieved through an ex post inefficient decision rule generating pivotal scenarios in which lying is costly.

Sarvary and Parker (1997) and Sarvary (2002) examine how the value of information is influenced by the presence of multiple senders. In the first paper, the value of the information sold by one information provider may increase when a competitor enters the market. In the second paper, featuring sequential purchase of information, the second information seller may be able to charge a higher price than the first information seller.

Finally, another literature involving multiple consultations (typically across time but sometimes also across individuals) is the reputational concerns literature (see Sobel 1985; Ottaviani and Sørensen 2001; Morris 2001; Ely and Välimäki 2003; Avery and Meyer 2012). The broad idea of these models is that experts have incentives to bias their reports so as to ensure a good reputation, the latter being intrinsically valued or instrumental to generating future opportunities to be consulted.

The paper contributes to the existing literature in two main ways. It considers the little studied sequential consultation problem and proposes a simple partitional communication scenario that is analytically tractable. In so doing, it identifies a new strategic incentive to affect future consultation and demonstrates that this can generate the basis for nontrivial comparative statics. A substantial caveat to the above remarks is that novelty comes at the price of extreme (perhaps excessive) simplification in remaining assumptions.

We proceed as follows. Section I presents a simple example with three states aimed at demonstrating the key mechanism underlying our results. Section II introduces the general model featuring a continuous state space. Section III provides an analysis of the general model for the two experts case. Section IV extends results to the n-experts case. Section V discusses other main extensions. Section VI concludes. Unless explicitly remarked upon, proofs are relegated to Appendixes. Part I of the online Appendix examines in detail the case of simultaneous consultation. Part II of the online Appendix contains the joint proof of Lemmas 6 and 8.

I. A Simple Example

A receiver R faces two experts. The state of the world is given by a variable ω drawn from a commonly known uniform distribution on $\{-1, 0, 1\}$. We call the states L, M, and H, where $L = -1$, $M = 0$, and $H = 1$. Each sender privately knows the state, while R only knows the ex ante distribution. Given a choice of action a and a state ω, R's utility is $u(a, \omega) = -(a - \omega)^2$. It follows that R's optimal action is the expected value of the state given available information. The preference type of individual senders is privately known to these and drawn from $\{U, B\}$ according to the distribution $\{\beta, 1 - \beta\}$, each sender's type being independently drawn. If a sender's type is U, his utility function coincides with that of R, while if his type is B, his utility function is given by $\hat{u}(a, \omega) = a$. A sender thus either wants to induce correct beliefs or to maximize beliefs.

R may consult senders at a cost of c per consultation. If consulted, a sender does not know his position in the consultation sequence. This implies that if he is consulted second, he does not observe the message sent by the preceding expert. We consider two communication strategy profiles. Strategy profile σ_{sr} specifies that an unbiased sender sends $m = \omega$, while a biased sender always sends $m = H$. Strategy profile σ_p specifies that an unbiased sender sends $m = L$ if $\omega = L$ and otherwise sends H, while a biased sender always sends $m = H$. Communication by unbiased senders is thus less informative in the second profile. We consider two consultation strategies of R. The first (ϕ_1) consists in consulting only once. The second (ϕ_2) specifies that after consulting once, R consults again if she may learn new information from asking for a second opinion. In equilibria featuring the two communication strategy profiles defined above, this means that on the equilibrium path R consults again only after message H. After stopping consultation, R picks the expected payoff maximizing action given her beliefs.

In what follows, we show that committing to consulting only once can lead to more informative individual communication and more learning than retaining the option to ask for a second opinion. In all scenarios considered, we assume that R believes that the state is -1 after an out of equilibrium message profile. We speak of R's belief to describe her expected value of the state.

Single Consultation.—Consider a putative equilibrium featuring σ_{sr} and ϕ_1, assuming that R is able to commit to ϕ_1. Beliefs of R after L and M are given by -1 or 0, respectively. After H, R's belief is

$$\mu_H^{sr}(\beta) = -\frac{\frac{1}{3}(1-\beta)}{\frac{1}{3} + \frac{2}{3}(1-\beta)} + \frac{\frac{1}{3}}{\frac{1}{3} + \frac{2}{3}(1-\beta)}.$$

Note that this equilibrium exists if and only if $\mu_H^{sr} \geq 0$, because otherwise a biased sender profitably deviates to M, which triggers the higher belief 0. It is easily checked that $\mu_H^{sr} \geq 0$ for any β, so that the equilibrium always exists. Setting $c = 0$, this guarantees R an expected payoff of

$$\Pi_1(\beta) = -\frac{1}{3}(1-\beta)\left(\mu_H^{sr}(\beta)+1\right)^2 - \frac{1}{3}(1-\beta)\left(\mu_H^{sr}(\beta)\right)^2 - \frac{1}{3}\left(\mu_H^{sr}(\beta)-1\right)^2.$$

Note that conditional on R committing to ϕ_1, there also always exists an equilibrium featuring σ_p. In such an equilibrium, deviating to L is never an attractive option for a biased sender because it yields belief -1. Note furthermore that the first equilibrium is more Blackwell-informative (see Blackwell 1951) than the second, so that R always favors the first over the second equilibrium.

Repeated Consultation.—We shall now examine two (putative) equilibria in which R asks for a second opinion whenever this is informative, i.e., uses ϕ_2. Consider first a scenario in which senders use σ_{sr}. R consults again if and only if the first consultation yielded H. Indeed, after any other message, she knows that she encountered an unbiased sender and will not learn more from a second consultation. Beliefs after an equilibrium message profile containing at least one L or M message are given by -1 and 0, respectively. After two H messages, R's belief is

$$\mu_{HH}^{sr}(\beta) = -\frac{\frac{1}{3}(1-\beta)^2}{\frac{1}{3}+\frac{2}{3}(1-\beta)^2} + \frac{\frac{1}{3}}{\frac{1}{3}+\frac{2}{3}(1-\beta)^2}.$$

We now check deviation incentives of senders. If we can ensure that at $\omega = -1$ a biased sender has no profitable deviation to M, then a biased sender never has a profitable deviation. Message M triggers belief 0 for sure as R does not consult again. Message H either leads to μ_{HH}^{sr} or to -1. R's final belief will be μ_{HH}^{sr} in two events (A and B). In A, the expert is the second expert consulted. In B, the expert is the first expert consulted, R consults again and meets a biased expert. The probability of $(A \cup B)$ can be shown to be

$$\frac{2\beta(1-\beta)}{1-(1-\beta)^2}.$$

A biased sender thus has no deviation incentive if and only if

$$(1) \qquad \frac{2\beta(1-\beta)}{1-(1-\beta)^2}\,\mu_{HH}^{sr}(\beta) + \left(1-\frac{2\beta(1-\beta)}{1-(1-\beta)^2}\right)(-1) \geq 0.$$

It can be shown that (1) is satisfied if and only if $\beta < 0.5$. It follows that (σ_{sr}, ϕ_2) is an equilibrium if and only if $\beta < 0.5$.

Consider now a putative equilibrium featuring the communication strategy σ_p and the consultation strategy ϕ_2. After two H-messages, R's belief is μ_{HH}^{p}. This equilibrium always exists because a biased sender trivially never deviates to message L

which yields the minimal belief -1. Assuming $c = 0$, the expected payoff of R in this scenario is

$$\Pi_2(\beta) = -\frac{1}{3}(1-\beta)^2\left(\mu^p_{HH}(\beta)+1\right)^2 - \frac{1}{3}\left(\mu^p_{HH}(\beta)-0\right)^2 - \frac{1}{3}\left(\mu^p_{HH}(\beta)-1\right)^2.$$

Comparing Single and Repeated Consultation.—We may now conclude. It can be shown that for β large enough, $\Pi_1(\beta) > \Pi_2(\beta)$, implying that R learns more from a highly informative single consultation than from two moderately informative consultations.

II. Model

A receiver R faces a set of M experts. The state of the world is given by a variable ω drawn from a uniform distribution F on $[0, 1]$. Each of the senders (or experts) is perfectly informed of the state, while R only knows F. Given a choice of action a and a state ω, R's utility is $u(a, \omega) = 1 - (a - \omega)^2$. It follows that R's optimal action, given information I, is to choose $a = E[\omega \mid I]$. The preference type of individual senders is privately known to these and drawn from $\{U, B\}$ according to the distribution $P(U) = \beta$, each sender's preference type being drawn independently. If a sender's type is U, his utility function coincides with that of R, while if his type is B, his utility function is given by $\hat{u}(a, \omega) = a$. A sender of type B thus simply aims at maximizing R's action.

R may consult senders at a cost of c per consultation. Consultation is sequential and sequentially rational. R randomly picks a first sender, receives a message, then consults a second sender if she wishes to, and potentially a third, etc. Once she stops consultation, R makes a decision. A strategy of R specifies a consultation strategy as well as a decision rule. A consultation strategy is a complete specification of R's consultation behavior at every information set. A decision rule determines the action chosen by R at the end of the consultation phase given messages observed.

A sender does not observe whether any other sender has been consulted before him, which entails that he also does not observe other senders' reports. A set of messages $[0, 1]$ is available to each sender, the emission of a message being costless. A strategy for a sender S_i specifies a distribution over $[0, 1]$ at each information set of S_i. A pure communication strategy of S_i is a mapping $\{U, B\} \times [0, 1] \to [0, 1]$. We focus on strategy profiles, such that all senders use the same pure communication strategy (so-called symmetric profiles).

We examine Perfect Bayesian equilibria and simply call them equilibria. An equilibrium of the game is given by a communication strategy profile and a receiver strategy, such that none of the parties has an advantageous deviation. R thus has no incentive to stop consulting earlier or later than specified by her strategy, and always simply chooses the ex post optimal action given beliefs. Furthermore, an individual sender has no incentive to deviate from his prescribed communication strategy. We speak of R's *gross* expected payoff in a given equilibrium if ignoring consultation costs. Once these costs are included, we instead speak of R's *net* expected payoff or simply of her expected payoff.

III. The Two Experts Case

We now analyze the case of $M = 2$. We shall examine three types of communication strategy profiles that are ordered in terms of the informativeness of individual experts' communication. As we shall see, the first two types of communication strategies are not compatible with sequential consultation. The first is only incentive compatible for senders under repeated consultation, but discourages a second consultation. The second incentivizes a second consultation, but is not incentive compatible for senders given repeated consultation. The third type of communication strategy is in contrast compatible with sequential consultation.

A. Fully Revealing Communication

A first scenario is fully revealing communication: Each sender sends $m = \omega, \forall \omega$. Consider a putative equilibrium in which R consults sequentially two experts. In such a scenario, R, however, stops consultation after the first expert given that she learns nothing from the second consultation. Note that there exists no equilibrium involving fully revealing communication and a single consultation. In such a putative equilibrium a biased sender profitably deviates to sending the highest message $m = 1$.

B. Semi-Revealing Communication

We now introduce a communication strategy originally proposed by Morgan and Stocken (2003) for the case of a single consultation of an expert with unknown bias. We generalize it to the case of more than one consultation.

DEFINITION 1: *A semi-revealing strategy profile is characterized by a threshold* $\theta \in [0, 1]$, *such that* (*i*) *an unbiased expert sends* $m = \omega$ *if* $\omega < \theta$ *and sends* $m = \theta$ *if* $\omega \geq \theta$, *while* (*ii*) *a biased expert always sends* $m = \theta$.

In an equilibrium involving the above strategy, the highest message $m = \theta$ is suspicious (and is the only one to be so). The latter message renders a second consultation potentially informative. By asking for a second opinion after $m = \theta$, R may, with positive probability, receive $m' < \theta$, in which case she will know that she now met an unbiased expert, and will adopt belief m'. If in contrast R first receives a message $m < \theta$, she recognizes that she cannot learn more about the state than what she already knows. By consulting again, she will either receive m again or $m' = \theta$, both cases leading to no change in her beliefs, which continue to assign probability one to $\omega = m$. As a matter of fact, after $m < \theta$, all R learns through a second consultation is whether the second expert is unbiased (if $m' = m$) or biased (if $m' = \theta$). We now introduce the following sequential consultation strategy of R.

DEFINITION 2: *Assume that experts use the semi-revealing strategy profile* θ. *The sequential consultation strategy* τ_2 *specifies that R consults for sure once and* *consults a second expert if and only if she receives the high message $m = \theta$ in the first consultation. The sequential consultation strategy τ_1 specifies that R consults exactly once for sure.*

We shall now define conditions under which given the consultation strategy $\tau_n \in \{\tau_1, \tau_2\}$, the semi-revealing strategy profile θ is incentive compatible. Define first the following function:

$$B(\theta, \beta, n) = \frac{\left(1-(1-\beta)^n\right)(1-\theta)E(\omega \mid \omega \geq \theta)}{\left(1-(1-\beta)^n\right)(1-\theta)+(1-\beta)^n} + \frac{(1-\beta)^n E(\omega)}{\left(1-(1-\beta)^n\right)(1-\theta)+(1-\beta)^n}.$$

The function $B(\theta, \beta, n)$ denotes the expected value of the state conditional on (i) senders being known to use the semi-revealing strategy profile θ and (ii) R using τ_n and having over the course of n sequential consultations received n times message $m = \theta$. It can be checked that for given β, n, a unique $\theta \in (0, 1)$ satisfies

$$B(\theta, \beta, n) = \theta. \tag{2}$$

It is given by

$$\theta^*(\beta, n) = \frac{1}{\sqrt{(1-\beta)^n}+1}.$$

Uniqueness of a solution to (2) follows because $B(\theta, \beta, n)$ is a concave function of θ and $B(0, \beta, n) > 0$ as well as $B(1, \beta, n) < 1$. We may now state the following:

LEMMA 1: (*i*) *Assuming that R uses $\tau_n \in \{\tau_1, \tau_2\}$, θ is incentive compatible for unbiased senders if and only if $\theta = \theta^*(\beta, n)$*:

(ii) *Assuming that R uses τ_1, $\theta^*(\beta, 1)$ is incentive compatible for biased senders*.

(iii) *Assuming that R uses τ_2, $\theta^*(\beta, 2)$ is not incentive compatible for biased senders*.

Point (i) states a necessary condition for (θ, τ_n) to constitute an equilibrium strategy profile. Given n, the belief of R after receiving n times message $m = \theta$ should be equal to θ. It follows that the set of final beliefs induced in an equilibrium (θ, τ_n) is the interval $[0, \theta]$. Point (iii) implies that there exists no equilibrium featuring a semi-revealing strategy profile and τ_2. The argument is as follows. In the only possible candidate semi-revealing profile $\theta^*(\beta, 2)$, given $\omega = 0$, a biased sender deviates to $\theta^*(\beta, 2) - \epsilon$, for ϵ positive and small. He thereby preempts a potentially damaging next consultation and ensures beliefs arbitrarily close to the maximum equilibrium belief $\theta^*(\beta, 2)$.

Point (ii) implies that there may exist an equilibrium of the form $(\theta^*(\beta, 1), \tau_1)$. We now examine incentives of R in such a putative equilibrium. We introduce some preliminary notation for the marginal value of a consultation in equilibrium for R. Given β and senders using $\theta^*(\beta, 1)$, let $v_{sr}^1(\beta)$ denote the increase in gross expected payoff achieved through the first consultation under the assumption that no consultation is subsequently done. Clearly, given a sequentially rational consultation strategy in which R consults a second time if advantageous, the value of the first consultation is weakly larger than $v_{sr}^1(\beta)$. Similarly, given β and senders using $\theta^*(\beta, 1)$, let $v_{sr}^2(\beta)$ denote the gain in gross expected payoff achieved through a second and final consultation, evaluated conditional on having received $m = \theta^*(\beta, 1)$ in the first consultation.

LEMMA 2: *Assume that senders use* $\theta^*(\beta, 1)$:

(i) *For c sufficiently low,* τ_1 *is not incentive compatible.*

(ii) *Let* $c = v_{sr}^2(\beta) + \epsilon$. *For* β *large enough and* ϵ *sufficiently small,* τ_1 *is incentive compatible.*

Point (i) trivially follows from the fact that after receiving $m = \theta^*(\beta, 1)$ in the first consultation, R may potentially learn more from consulting again and will thus deviate to doing so if c is sufficiently low. Point (ii) offers sufficient conditions for the existence of an equilibrium featuring the strategy profile $(\theta^*(\beta, 1), \tau_1)$. The key behind point (ii) is that for β large $v_{sr}^1(\beta) > v_{sr}^2(\beta)$ because $\lim_{\beta\to 1} v_{sr}^1(\beta) = \frac{1}{12}$, while $\lim_{\beta\to 1} v_{sr}^2(\beta) = 0$. It follows that for β large enough, one can pick a $c \in (0, 1)$ s.t.

$$v_{sr}^1(\beta) \geq c \geq v_{sr}^2(\beta),$$

thus making τ_1 incentive compatible given $\theta^*(\beta, 1)$.

C. Partitional Communication

Lemmas 1 and 2 imply that given an equilibrium consultation plan specifying multiple consultations the communication strategy of senders must be some third type of strategy. We now introduce the class of partitional communication strategies and show that it is compatible with repeated consultation.

DEFINITION 3: *An N-intervals strategy profile is characterized by thresholds* $t_0 = 0 < t_1 < .. < t_{N-1} < t_N = 1$, *such that* (i) *an unbiased sender sends* m_i *if* $\omega \in [t_{i-1}, t_i)$, $\forall\ i \in \{1, .., N\}$ *and sends* m_N *if* $\omega = 1$, *while* (ii) *a biased sender always sends* m_N.

In what follows, we denote an N-intervals strategy profile by $(t_1, .., t_{N-1})$, thus omitting t_0 and t_N whose values are exogenously fixed at 0 and 1, respectively. In such a profile, an unbiased expert correctly reveals the interval to which ω belongs, while a biased expert claims that ω belongs to the highest interval. As in the case of an equilibrium featuring a semi-revealing profile θ, the highest message is the only one justifying further consultation. By asking for a second opinion after m_N, R may receive $m_i \neq m_N$ in which case she will know that she now met an unbiased expert and will revise her beliefs to $E[\omega \mid \omega \in [t_{i-1}, t_i)]$. If instead R receives $m_i \neq m_N$ in the first consultation, she recognizes that another consultation will not affect her beliefs. Whether the second consultation yields m_i again or m_N, she will indeed continue to assign probability one to $\omega \in [t_{i-1}, t_i)$. We now introduce a sequential consultation strategy that is a counterpart of τ_2.

DEFINITION 4: *Assume that senders use an N-intervals strategy profile. ϕ_2 specifies that R consults once for sure and consults a second expert if and only if she receives the high message m_N in the first consultation. The sequential consultation strategy ϕ_1 specifies that R consults exactly once for sure.*

In what follows, we study putative equilibria involving a partitional strategy profile and ϕ_1 or ϕ_2, with a primary focus on ϕ_2. Before going into formal detail, we provide a general intuition for the strategic incentives faced by senders given that R uses ϕ_2. A biased expert, conditional on $\omega = 0$ and being consulted first, would prefer the receiver not ask for a second opinion as she might meet an unbiased expert who will send the low message m_1. This creates an incentive for the biased expert to deviate to reporting m_{N-1} instead of m_N. Indeed, by not sending m_N, he signals himself as a truthful expert and preempts a second consultation. A biased sender will only refrain from deviating to m_{N-1} if the latter is very costly because m_{N-1} triggers very low beliefs. To make m_{N-1} very unattractive, $[t_{N-2}, t_{N-1})$ should be large, which in turn implies that $[t_{N-3}, t_{N-2})$ also should be large to ensure that unbiased senders are indifferent between m_{N-1} and m_{N-2}, etc. Large intervals to the left of t_{N-1}, however, imply that an unbiased expert communicates in a very coarse fashion. Ensuring that a biased expert does not deviate thus requires that unbiased senders communicate noisily by a logic of contagion.

The next two lemmas provide a formal characterization of sender incentives. We define the conditions under which an N-intervals strategy profile is incentive compatible for unbiased senders and biased senders, respectively, given that R follows the sequential sampling strategy ϕ_n, for $n \in \{1, 2\}$. We call such a profile U-IC and B-IC, respectively.

LEMMA 3: *Assume that R uses $\phi_n \in \{\phi_1, \phi_2\}$:*

(i) The N-intervals strategy profile $(t_1, .., t_{N-1})$ is U-IC if and only if

$$(3) \qquad B(t_{N-1}, \beta, n) - t_{N-1} = t_{N-1} - \left(\frac{t_{N-1} + t_{N-2}}{2}\right)$$

and

$$t_{i+1} - t_i = t_i - t_{i-1}, \quad \forall i \in \{1, .., N-2\}. \tag{4}$$

(ii) *For every $N \geq 2$, there exists a unique U-IC N-intervals strategy profile.*

Note for later reference that (3) and (4) together imply

$$\frac{B(t_{N-1}, \beta, n)}{t_{N-1}} = \frac{2(N-1)+1}{2(N-1)}. \tag{5}$$

Equalities (3) and (4) state the standard indifference condition on thresholds defining intervals. At every interior threshold t_i, an unbiased sender must be indifferent between m_i and m_{i+1}. For $t_1, .., t_{N-2}$, indifference requires

$$t_i - E[\omega \mid \omega \in [t_{i-1}, t_i)] = E[\omega \mid \omega \in [t_i, t_{i+1})] - t_i,$$

which simplifies to (4). The condition stated for t_{N-1} in (3) is slightly more complex because the belief triggered by m_N is not given by $E[\omega \mid \omega \in [t_{N-1}, 1]]$, but instead by $B(t_{N-1}, \beta, n)$. Recall indeed that a homogeneous profile of n m_N-messages implies $E[\omega \mid \omega \in [t_{N-1}, 1]]$ only if it was sent by n unbiased senders, which cannot be ascertained.

Point (ii) implies that at most one N-intervals strategy profile can constitute an equilibrium for any given $N \geq 2$. We henceforth denote by $\{t_r(\beta, N, n)\}_{r=1}^{N-1}$ the unique threshold profile that is incentive compatible for unbiased senders for given β, N, n. Point (ii) also demonstrates that restrictions on the fineness N of equilibrium partitions do not originate in unbiased senders given that there exists a U-IC N-intervals profile for any $N \geq 2$.

Finally, it is easily shown that given β and $n \in \{1, 2\}$, for N tending to infinity the unique U-IC N-intervals strategy profile tends to the semi-revealing strategy profile $\theta^*(\beta, n)$. To see this, note two facts. First, as N increases $B(t_{N-1}(\beta, N, n), \beta, n)$ tends toward $t_{N-1}(\beta, N, n)$, so that in the limit $t_{N-1}(\beta, N, n)$ satisfies the same condition as θ in (2). Second, as N increases, intervals to the left of t_{N-1} become infinitely many and infinitely small.

We now examine incentives of biased senders. Note that if R consults only once in equilibrium (i.e., uses ϕ_1), the U-IC N-intervals profile is always incentive compatible for biased senders because message m_N trivially maximizes R's expected beliefs. If R instead uses ϕ_2, whether m_N maximizes R's expected beliefs is nontrivial because m_N triggers further consultation and may lead to a low message if the state is low and the next sender is unbiased.

LEMMA 4: *Assume that R uses $\phi_n \in \{\phi_1, \phi_2\}$. The U-IC N-intervals strategy profile $(t_1, .., t_{N-1})$ is B-IC if and only if*

$$\frac{B(t_{N-1}, \beta, n)}{t_{N-1}} \geq \left(\frac{3}{2} - \frac{1}{2}\frac{n\beta(1-\beta)^{n-1}}{1-(1-\beta)^n}\right). \tag{6}$$

The above condition ensures that a biased expert prefers to send m_N rather than deviating to m_{N-1} if $\omega = 0$, which is in turn sufficient to ensure that he has no profitable deviation for any $\omega > 0$. The condition is trivially satisfied for $n = 1$ given that the right-hand side reduces to 1 while the left-hand side is always greater than 1. Consider now the case of $n = 2$, i.e., where R asks for a second opinion after m_N. Let $\omega = 0$. Sending m_{N-1} triggers for sure belief $\mu(m_{N-1})$ satisfying

$$\mu(m_{N-1}) = \frac{t_{N-1} + t_{N-2}}{2} = \left(\frac{4t_{N-1} - 2B(t_{N-1}, \beta, 2)}{2}\right), \tag{7}$$

as R will not consult again after m_{N-1}. On the other hand, the final belief of R after m_N is uncertain. One possible scenario (Scenario 1) is that R consults again and meets an unbiased expert, thus adopting $\mu(m_1)$ satisfying

$$\mu(m_1) = \frac{t_1}{2} = \frac{2\big(B(t_{N-1}, \beta, 2) - t_{N-1}\big)}{2}.$$

R's final belief will instead be $\mu(m_N) = B(t_{N-1}, \beta, 2)$ if the expert is the second to be consulted (Scenario 2.a) or if R consults again and meets a biased expert (Scenario 2.b), the summed probability of Scenarios 2.a and 2.b being

$$\frac{2\beta(1-\beta)}{1-(1-\beta)^2}$$

The expected belief triggered by m_N is thus given by

$$\left(\frac{2\beta(1-\beta)}{1-(1-\beta)^2}\right)\mu(m_N) + \left(1 - \frac{2\beta(1-\beta)}{1-(1-\beta)^2}\right)\mu(m_1). \tag{8}$$

Summarizing, the U-IC N-intervals strategy profile $(t_1, .., t_{N-1})$ is incentive compatible for a biased sender given ϕ_n if and only if (8) is larger than (7), which is equivalent to (6).

We may now gather the conditions ensuring incentive compatibility of a given N-intervals profile for unbiased as well as biased senders. Applying simultaneously (5) and (6), we may state the following:

LEMMA 5:

(i) *Given β and ϕ_2, there exists an incentive compatible N-intervals profile if and only if $N \leq \frac{2}{\beta}$. If there exists one, it is furthermore unique.*

(ii) *Given β and ϕ_1, there exists a unique incentive compatible N-intervals profile for any $N \geq 2$.*

Point (i) is obtained by combining (5) and (6), yielding

$$\frac{2(N-1)+1}{2(N-1)} \geq \frac{3}{2} - \frac{1}{4}\,\frac{1-(1-\beta)^2}{2\beta(1-\beta)},$$

which simplifies to $\frac{2}{\beta} \geq N$. Point (ii) is immediate given that (6) is always true for $n = 1$.

Given that R uses ϕ_2, for any β there is thus a finite maximal partition fineness $\bar{N}(\beta, 2) = \langle \frac{2}{\beta} \rangle \geq 2$, where we let $\langle \frac{2}{\beta} \rangle$ denote the highest integer smaller than $\frac{2}{\beta}$. The upper bound $\bar{N}(\beta, 2)$ furthermore decreases in β. For any integer $X \geq 2$, we may also state that $\bar{N}(\beta, 2) = X$ if and only if $\beta \in \left(\frac{2}{X+1}, \frac{2}{X}\right]$. Equilibria with more than two intervals thus exist only for β small enough, while a two-intervals equilibrium exists for any β. High individual trustworthiness thus implies coarse reporting. Note that $\bar{N}(\beta, 2)$ very quickly decreases as β increases. For $\beta \in \left(\frac{1}{2}, \frac{2}{3}\right]$, we have $\bar{N}(\beta, 2) = 3$, while for $\beta > \frac{2}{3}$, $\bar{N}(\beta, 2) = 2$.

Having now stated incentive conditions for senders formally, we briefly refine the intuition behind sender incentives. First, note that a high N implies very small partition intervals, which in turn means that the belief $E[\omega \mid m_{N-1}]$ triggered by the second highest message m_{N-1} is high and close to that triggered by two messages m_N. For a biased sender, the deviation payoff attached to m_{N-1} thus increases monotonically in N, which means that keeping fixed the expected payoff attached to m_N, finer equilibria are less likely to be incentive compatible for biased senders. Second, assuming $\omega = 0$ a higher β or n implies a higher chance that R meets an unbiased truthteller if she consults again and thereby adopts belief $E[\omega \mid m_1]$. Keeping fixed the payoff attached to m_{N-1}, sending m_N thus becomes increasingly unattractive as β or n increases. It thus follows that as β or n increases equilibria break down in a top-down order. Very fine equilibria first disappear, then fine ones, then relatively coarse ones, etc., until only the two-intervals equilibrium is left. The latter never breaks down because in this equilibrium sending m_2 is always by definition more attractive than sending m_1 for a biased sender given $\omega = 0$.

D. Preliminary Comparative Statics of R's Welfare

We now proceed to the welfare analysis. We first state some preliminary results relating to the comparative statics of gross expected payoffs and subsequently use these to derive the sampler's curses.

Define $V_{sr}(\beta, 1)$ as the gross expected payoff of R in an equilibrium featuring a unique consultation and the semi-revealing profile $\theta^*(\beta, 1)$. Let $V_p(\beta, N, n)$ denote the gross expected payoff of R in an equilibrium featuring the sequential sampling strategy ϕ_n and the N-intervals strategy profile $\{t_r(\beta, N, n)\}_{r=1}^{N-1}$.

LEMMA 6:

(i) $V_{sr}(\beta, 1)$ *is continuous and increasing in* β.

(ii) *$V_p(\beta, N, n)$ is continuous and increasing in β as well as increasing in N and n.*

(iii) $\lim_{N\to\infty} V_p(\beta, N, 1) = V_{sr}(\beta, 1)$.

To understand why $V_{sr}(\beta, 1)$ is increasing in β, note that increasing β increases the threshold $\theta^*(\beta, 1)$ as well as the probability that the consulted expert is unbiased, both effects being beneficial. The comparative statics of $V_p(\beta, N, n)$ are also simple and intuitive. It increases with expert quality β, the fineness N of communication and the (maximal) number n of experts consulted. Points (ii) and (iii) together imply that conditional on R's strategy specifying a unique consultation, she favors the equilibrium featuring the semi-revealing profile over any equilibrium featuring a partitional profile.

Note that given two partitional equilibria featuring ϕ_2 together with $\{t_r(\beta, N, 2)\}_{r=1}^{N-1}$ or $\{t_r(\beta, N+1, 2)\}_{r=1}^{N}$, respectively, the second (and finer) equilibrium dominates the first not only in terms of R's gross expected payoff but also in terms of her net expected payoff. Indeed, condition (5) implies that $t_N(\beta, N+1, 2) > t_{N-1}(\beta, N, 2)$, which means that R is less likely to consult a second time in the second equilibrium. Consider now expert utilities. An unbiased expert has the same preferences as R except he does not consider consultation costs. His payoff ranking of equilibria is accordingly the same as R's (net and gross) payoff ranking of equilibria. On the other hand, a biased expert's expected payoff from a given equilibrium depends only on the expected value of R's beliefs. This being by definition the same in any equilibrium $\left(\text{and equal to } \frac{1}{2}\right)$, biased experts are ex ante indifferent across all equilibria. It follows that given a fixed $\phi_n \in \{\phi_1, \phi_2\}$, R's (gross and net) payoff ranking of partitional equilibria (and the semi-revealing equilibrium if $\phi_n = \phi_1$) is also a Pareto ranking.

E. The Sampler's Curses

We now turn to the key welfare-related comparative statics of our model. We ask whether R necessarily gains from an increase in the number M of available experts, a decrease in consultation costs c or an increase in expert trustworthiness β. We shall see that the answer to all three questions is negative, thereby identifying a set of *sampler's curses*. A note of caution on the nature of our exercise is here warranted: For each set of parameters β, n, we exclusively consider the R-optimal equilibrium within the limited class of equilibria that we study. We thus focus on the semi-revealing communication scenario in the case of only one equilibrium consultation and on N-intervals strategy profiles in the case of multiple equilibrium consultations.

We first identify a case in which R loses from an increase in the number of experts available.

PROPOSITION 1 (The Curse of Manyness): *There is a $\beta^* \in \left(\frac{1}{2}, 1\right)$, such that given $\beta \geq \beta^*$ and c low enough, R's expected payoff is larger if facing only one expert than if facing two experts.*

The above result bases on the simple trade-off faced by R between the quality of individual reporting and the quantity of experts consulted. For c low enough, when facing two experts, R cannot commit to not consulting again after the first consultation if an extra consultation is informative. It follows that there exists no equilibrium featuring a single consultation and the semi-revealing strategy profile (see Lemma 2). Instead, there only exist simple finite N-intervals equilibria in which partitioning bounds the amount of information that can be retrieved by R even if she consults an arbitrary number of experts. When in contrast only one expert is available, R is de facto committed to a unique consultation, thereby allowing semi-revealing communication to constitute an equilibrium outcome. The informativeness of such an equilibrium scenario furthermore converges smoothly to 1 as β tends to 1. It follows that for β high enough and c low enough, R prefers to face a single expert rather than two.

We now identify a case in which R loses from a decrease in the cost of consultation. Recall that $v_{sr}^2(\beta)$ denotes the gross value of a second and final consultation conditional on having received $m = \theta^*(\beta, 1)$, assuming that senders use the semi-revealing profile $\theta^*(\beta, 1)$.

PROPOSITION 2 (The Curse of Inexpensiveness): *Let there be two experts. For β sufficiently high, there is an $\epsilon^* > 0$, such that for $\epsilon \leq \epsilon^*$ R's expected payoff is larger given $c = v_{sr}^2(\beta) + \epsilon$ than given $c' = 0$.*

The above result builds on the same mechanism as Proposition 1. An excessively low c breaks the equilibrium in which R consults once and senders use the semi-revealing profile $\theta^*(\beta, 1)$ because R asks for a second opinion after $m = \theta^*(\beta, 1)$. Semi-revealing communication being very informative, a strictly positive (though not too high) c serves as a beneficial commitment device that enables semi-revealing communication.

The following proposition identifies a case in which R loses from an increase in β.

PROPOSITION 3 (The Curse of Trustworthiness): *Let there be two experts. For every $N \geq 2$, for c sufficiently small there is an $\epsilon^* > 0$, such that for $\epsilon \leq \epsilon^*$ R's expected payoff is smaller under $\beta' = \frac{2}{N} + \epsilon$ than under $\beta'' = \frac{2}{N} - \epsilon$.*

The above result bases on a simple trade-off. Given that R uses ϕ_2, a higher β increases a biased expert's incentive to deviate to m_{N-1} conditional on $\omega = 0$ and thereby negatively affects the achievable fineness of individual reporting. At the threshold value of $\beta = \frac{2}{N}$, an infinitesimal increase in β leads to a discontinuous downwards jump in the maximal achievable number of intervals while affecting only infinitesimally the likelihood of meeting a biased expert. Such an increase in β is thus unambiguously disadvantageous.

IV. More than Two Experts

We now examine the general case of $M \geq 2$ senders. Given an N-intervals strategy profile, the consultation strategy ϕ_n specifies that R stops consulting as soon as she receives a message $m_i \neq m_N$ and continues consulting for a maximum of n rounds as long as receiving m_N. By the arguments given in Lemmas 3 and 4, we may state the following:

LEMMA 7: *Assume that R uses ϕ_n, for $n \geq 1$. There exists an incentive compatible N-intervals strategy profile if and only if*

$$\frac{1}{2(N-1)} + 1 \geq \frac{3}{2} - \frac{1}{4}\,\frac{n\beta(1-\beta)^{n-1}}{1-(1-\beta)^n}.$$

If there exists one, it is unique.

Given ϕ_n and β, we denote by $\{t_r(\beta, N, n)\}_{r=1}^{N-1}$ the unique incentive compatible N-intervals strategy profile. On the basis of the above, we may state the following three comparative statics properties of the equilibrium set w.r.t. N, β, and n. The comparative statics result for N reads as follows. For any β and $n \geq 2$, there is a finite upper bound $\bar{N}(\beta, n) \geq 2$ on the feasible number of intervals given ϕ_n. The bound $\bar{N}(\beta, n)$ furthermore decreases in β and n. A two-intervals profile is always incentive compatible no matter n and β. The comparative statics result for β reads as follows. For any given $N > 2$ and $n \geq 2$ there is a finite upper bound $\bar{\beta}(N, n)$, such that given ϕ_n, there exists an incentive compatible N-intervals profile if and only if $\beta \leq \bar{\beta}(N, n)$. The bound $\bar{\beta}(N, n)$ furthermore decreases in N and n. The comparative statics result for n reads as follows. For given β and $N > 2$, there is a finite upper bound $\bar{n}(\beta, N) \geq 1$ such that there exists an incentive compatible N-intervals strategy profile if and only if ϕ_n satisfies $n \leq \bar{n}(\beta, N)$. The bound $\bar{n}(\beta, N)$ furthermore decreases in β and N.

Extending our definition of the gross expected payoff $V_p(\beta, N, n)$ to $n \geq 2$, we may now state the following:

LEMMA 8: *For any $n \geq 1$, $V_p(\beta, N, n)$ is continuous and increasing in β and increasing in N and n.*

It can furthermore be shown that

$$\lim_{n\to\infty} V_p(\beta, N, n) = 1 + \frac{1-3N}{6N^2}.$$

The above limit expression is increasing in N and converging to 1 as N tends to infinity. When consulting infinitely many experts, bias is thus immaterial while the number of intervals is the only determinant of payoffs.

Note that as in the two experts case, given two equilibria featuring ϕ_n and either $\{t_r(\beta, N, n)\}_{r=1}^{N-1}$ or $\{t_r(\beta, N+1, n)\}_{r=1}^{N}$ the second (and finer) equilibrium also dominates the first in terms of R's gross expected payoffs. Indeed, it follows from (5), which carries over to the n-experts case, that $t_N(\beta, N+1, n) > t_{N-1}(\beta, N, n)$. This implies that R is less likely to consult an mth time for any $m \in \{2, .., n\}$ in the second equilibrium. The second equilibrium thus saves consultation costs in expectation. Given the nature of expert utilities (see a discussion for the two experts case), it follows that for any fixed ϕ_n R's (net and gross) payoff ranking of partitional equilibria is also a Pareto ranking.

We now present counterparts of Propositions 1, 2, and 3 for the n-persons case. We start with a counterpart of Proposition 1, identifying cases in which R loses from an increase in the number of experts available.

PROPOSITION 4 (The Curse of Manyness):

(i) *For every $M > 1$, there is a $\beta^* \in \left(\frac{1}{2}, 1\right)$, such that given $\beta \geq \beta^*$ and c low enough, R's expected payoff is larger if facing one expert than if facing M experts.*

(ii) *For every $M > 2$, given $\beta \in \left(\beta, \frac{2}{3}\right]$ and c low enough R's expected payoff is larger if facing two experts than if facing M experts.*

The mechanism behind point (i) is the same as that behind Proposition 1. As to point (ii), note that if $\beta \in \left(\frac{1}{2}, \frac{2}{3}\right]$, $N = 3$ is the finest achievable partitioning given ϕ_2 while the finest achievable partitioning drops to $N = 2$ for $\phi_3, \ldots, \phi_M$. Furthermore, for $\beta \in \left(\frac{1}{2}, \frac{2}{3}\right]$, $\lim_{n \to \infty} V_p(\beta, 2, n) < V_p(\beta, 3, 2)$, which implies that R is better off in an equilibrium featuring a 3-intervals profile and the sequential consultation strategy ϕ_2 rather than in an equilibrium featuring a 2-intervals profile and a sequential consultation strategy ϕ_n for any $n \geq 3$.

We now present a counterpart for the n-persons case of Proposition 2. We introduce some preliminary notation for the marginal value of a consultation in equilibrium for R. For fixed β and M and given experts using $\theta^*(\beta, 1)$, let $\tilde{v}^2_{sr}(\beta, M)$ denote the critical value of c above which there exists no sequentially rational consultation plan specifying consulting again after a first consultation yielding $m = \theta^*(\beta, 1)$.

PROPOSITION 5 (The Curse of Inexpensiveness): *Let there be $M \geq 2$ experts. For β sufficiently high there is an $\epsilon^* > 0$, such that for $\epsilon \leq \epsilon^*$ R's expected payoff is larger under $c = \tilde{v}^2_{sr}(\beta, M) + \epsilon$ than under $c' = 0$.*

The underlying mechanism behind the above being the same as in Proposition 2, we do not repeat the intuition. Propositions 4 and 5 admittedly do not provide as strong generalizations of Propositions 1 and 2 as one might wish or expect. In the case of Proposition 4, one would want a result comparing the cases of M and $M+1$ experts for any M. Achieving more general results would require the ability to provide welfare comparisons of pairs of partitional equilibria of the form

$$\left(\phi_n, \{t_r(\beta, N, n)\}_{r=1}^{N-1}\right), \left(\phi_{n+1}, \{t_r(\beta, N-1, n+1)\}_{r=1}^{N-2}\right).$$

One would focus on values of β, n, such that $\overline{N}(\beta, n+1) = \overline{N}(\beta, n) - 1$, meaning that given β, n, increasing consultation by one unit induces a decrease of one unit in the maximal feasible intervals number. One would then identify cases in which transiting from ϕ_n to ϕ_{n+1} decreases R's gross payoff because one more consulted sender does not make up for the caused loss in individual informativeness.

In seeking to obtain such a result, the problem encountered is that there is no simple formula for the difference between (or the ratio of) $V_p(\beta, N-1, n+1)$ and $V_p(\beta, N, n)$, let alone for the difference (or the ratio of) $V_p(\beta, \overline{N}(\beta, n)-1, n+1)$ and $V_p(\beta, \overline{N}(\beta, n), n)$. We leave a further examination of this issue for further work.

We now present a counterpart of Proposition 3. The underlying mechanism is the same as in Proposition 3 and thus not repeated.

PROPOSITION 6 (The Curse of Trustworthiness): *Let there be $M \geq 1$ experts. For every $N \geq 2$, for c sufficiently small there is an $\epsilon^* > 0$, such that for $\epsilon \leq \epsilon^*$ R's expected payoff is smaller under $\beta = \overline{\beta}(N, M) + \epsilon$ than under $\beta' = \overline{\beta}(N, M) - \epsilon$.*

V. Extensions

A. Asymmetric Equilibria and Babbling

It may appear problematic that we do not consider the possibility of babbling by a subset of experts given that excessive consultation often hurts R. A first response to this critique is that equilibria featuring babbling can be eliminated if we assume an expert cost to becoming informed or to communicating. An expert would not bear the cost of becoming informed and/or of producing a report (and even writing a bogus report is demanding) if he knows that his report will be ignored. The only way to reconcile costs on the sender side with babbling would be to assume that senders receive a monetary transfer. But why would a sender known to babble and who is never consulted be paid and who would pay him?

A second argument builds on the neologism proofness criterion (see Farrell 1993). The criterion assumes availability of a rich exogenous language endowed with a literal meaning and involves asking whether given neologisms are credible and thereby break the equilibrium. Let $M = 2$ and let experts be called S_1 and S_2. Consider an equilibrium in which S_1 uses the semi-revealing strategy $\theta^*(\beta, 1)$, S_2 babbles, and R only consults S_1. Suppose that R (voluntarily or inadvertently) encounters S_2 after having consulted S_1 and having received $m = \theta^*(\beta, 1)$. The statement "ω *is strictly below* $\frac{2}{3}\theta^*(\beta, 1)$" constitutes a credible neologism. Indeed, supposing it is believed and assuming that R already consulted S_1 and received $\theta^*(\beta, 1)$, S_2 would only want to state this if true. To see this, note that threshold $t^* = \frac{2}{3}\theta^*(\beta, 1)$ satisfies

$$\theta^*(\beta, 1) - t^* = t^* - \frac{t^*}{2},$$

implying that an unbiased S_2 strictly prefers inducing belief $\frac{t^*}{2}$ to the equilibrium belief $\theta^*(\beta, 1)$ if and only if $\omega < t^*$. Anticipating the possibility of credible communication by S_2, for c low enough R would thus want to consult S_2 after having consulted S_1 and received $m = \theta^*(\beta, 1)$. The equilibrium is thus not neologism proof according to a slightly extended definition of the criterion that includes endogenous consultation choice. A major caveat is that the neologism proofness criterion would also destroy equilibria with symmetric partitional communication and no babbling. Indeed, given an equilibrium partitional profile $\{t_r\}_{r=1}^{N-1}$, any neologism of the form "$\omega < \frac{1}{3} t_1$," would be credible.[2]

I conjecture that an adapted version of announcement proofness (see Matthews, Okuno-Fujiwara, and Postlewaite 1991) would eliminate equilibria featuring babbling by a subset of $M - n$ experts while at the same time preserving equilibria featuring a symmetric informative partitional profile, given c low enough. The modification of the criterion entails considering endogenous consultation (as above) as well as suitably restricting the set of available announcement strategies to ones that are isomorphic (not necessarily exactly identical) to the partitional strategy profile used by non-babbling agents. In any equilibrium featuring babbling by a subset of experts, I conjecture that some permitted announcement strategy (together with corresponding announcements) is credible for a babbling expert conditional on being consulted. Given c low enough R, anticipating this, deviates to consulting a babbling expert after a nth consultation yielding m_N. In contrast, in an equilibrium involving no babbling, I conjecture that no permitted announcement strategy (and corresponding messages) is credible.

B. Asymmetric Equilibria and Observable Expert Position

In many situations, experts are likely to know their position in the consultation sequence. Search engines provide a fairly predictable ordering of results for any given search. On a given US economic issue, *Wall Street Journal* articles appear above *New York Times* articles, which themselves appear before *Washington Post* articles, etc. Most people examine the highest ranked result first, then the second, etc. Some experts might also be marginally more credible, cheaper, or better accessible and thus tend to be consulted earlier than others. Claiming that experts know their position can refer to two possible cases. One possibility is that this is hardwired into the game: Experts simply observe their position in the consultation sequence. Another possibility is that experts do not observe their position but know it in equilibrium because R's strategy specifies a specific deterministic consultation sequence. A key difference is that in the first case, a deviation from the equilibrium consultation order is detected by experts. We shall focus primarily on the first case.

Could R benefit from experts observing their position in the consultation sequence rather than ignoring it? The answer is negative if we consider symmetric N-intervals communication strategy profiles. Conditional on sending m_N, early (late) experts in the observable position scenario anticipate a higher (lower) expected number of future consultations than experts in the unobservable position scenario. It follows that early biased experts in the observable position scenario face increased misreporting incentives while late biased experts face decreased misreporting incentives. Restricting ourselves to symmetric equilibria, incentives exhibit a lowest common

[2] If $M = 1$ (a unique sender), the semi-revealing profile $\theta^*(\beta, 1)$ appears to be neologism proof. This is noteworthy given that in the Crawford and Sobel (1982) model, no influential equilibrium is neologism proof.

denominator property: A symmetric equilibrium featuring very refined individual communication requires that even early senders face no deviation incentives, which is impossible to achieve.

If we allow for asymmetric strategies, it is less clear that the observable position regime is dominated. One would then expect optimal equilibria to feature coarse communication by early senders and increasingly refined communication as one moves along the consultation sequence. Suppose there are two experts $\{S_1, S_2\}$ and that S_1 is consulted first. A simple scenario is partitional communication by both with S_2's partition featuring more intervals. Let S_j use the N^j-intervals strategy $\{t^j_r\}_{r=1}^{N^j}$, assuming $N^2 > N^1$. If biased, S_j sends m_{N^j}. If unbiased, S_j sends m_r given $\omega \in [t^j_{r-1}, t^j_r)$ (including the upper bound for $r = N^j$).

One might first consider equilibria in which R only asks for a second opinion after a suspicious high message m_{N^1}, as in the main analysis. Checking incentives here is, however, very complex. If S_2 is much more informative than S_1, there are now new potential deviation incentives for S_1 and R as compared to symmetric partitional equilibria. An unbiased S_1 might deviate to sending m_N for ω low in order to encourage a second consultation, in the hope that R meets an unbiased S_2 and learns very accurate information. For the same reason R may wish to consult again after a low message.

One might instead consider equilibria in which R always asks for a second opinion, this being sequentially rational. Consultation behavior is now unaffected by the message received in round 1 and S_1's incentives are as a consequence much simpler. We now examine such a scenario. Given the profile $\{t^1_r\}_{r=1}^{N^1-1}, \{t^2_r\}_{r=1}^{N^2-1}$, let
AQ 1 $\tilde{B}(t^1_{N^1-1}, t^2_{N^2-1}, \beta, 2)$ be the expected value of ω upon receiving m_{N^1} followed by m_{N^2}. Necessary (NB: not sufficient) conditions for the above to be incentive compatible for senders are as follows. For each $j \in \{1, 2\}$,

$$\tilde{B}(t^1_{N^1-1}, t^2_{N^2-1}, \beta, 2) - t^j_{N^j-1} = t^j_{N^j-1} - \left(\frac{t^j_{N^j-1} + t^j_{N^j-2}}{2}\right)$$

and

$$t^j_{i+1} - t^j_i = t^j_i - t^j_{i-1}, \quad \forall i \in \{1, \dots, N^j - 2\}.$$

The partitional strategy of each sender S_j thus features $N^j - 1$ equally sized intervals to the left of $t^j_{N^j-1}$. Note that these conditions echo those given in Lemma 3.

We now consider a class of partitional profiles parametrized by N^1 and the integer x. In this class, x describes the relative fineness of S_2's communication with respect to S_1's communication. Given N^1 and $x \geq 2$, set $N^2 = (2x + 1)(N^1 - 1) + x + 1$. In other words, each interval to the left of $t^1_{N^1-1}$ in S_1's partitional strategy gives rise to $2x + 1$ nested intervals in the partitional strategy of S_2. The latter sender can thus potentially refine the information provided by S_1. Set furthermore

$$t^2_{N^2-1} = t^1_{N^1-1} + \frac{t^1_{N^1-1}}{(2(N^1-1))} \frac{x}{(x + \frac{1}{2})}.$$

We focus on the limit strategy for $x \to \infty$. As x tends to infinity, N^2 tends to infinity and it can be shown that $t^2_{N^2-1}$ tends to $\tilde{B}(t^1_{N^1-1}, t^2_{N^2-1}, \beta, 2)$. In the limit, S_2 thus uses a semi-revealing strategy with threshold $t^1_{N^1-1} + \frac{t^1_{N^1-1}}{2(N_1-1)}$. To find the value of $t^1_{N^1-1}$ featured in such a limit equilibrium given N_1, we simply need to find a t that solves

$$t + \frac{t}{2(N_1-1)} = B\left(t, t + \frac{t}{2(N_1-1)}, \beta, 2\right).$$

The solution of the above is unique and has a fairly simple closed form which is omitted here. An interesting feature of this class of equilibria is that it features gradual learning even conditional on R meeting only unbiased senders and the state being low. If $\omega \leq t^1_{N^1-1}$, an unbiased second sender refines the information provided by a first unbiased sender though both hold the same information and both have the same preferences. The only reason behind gradual learning is to make sure that R consults again even after a low message in the first consultation. Which is in turn necessary to ensure that S_1's strategy is incentive compatible, as explained earlier.

Asymmetric equilibria might also exist given unobservable sender positions. Let there be two senders and assume that sender identities are still observable (S_1 and S_2), positions in the consultation sequence being, however, unobservable to senders. Consider an asymmetric partitional equilibrium in which the first consulted sender is less informative than the second consulted sender, a second consultation taking place only after a high message. This appears impossible to support as an equilibrium outcome. If the first consulted sender (say S_1) is the least informative one, R would indeed deviate to consulting S_2 first. Under unobservable positions, equilibria featuring a conditional second consultation would thus have to feature a more informative first sender. The paradox of such equilibria is that given the equilibrium consultation strategy of R, the second consulted sender could in principle be much more informative given that he is the last one (implying no incentive to discourage a next consultation). The coarseness of his communication only serves to make R's consultation plan incentive compatible.

C. Other Forms of Expert Bias

Our modeling of expert bias is very particular. Staying close to our setup, three crucial aspects that are eluded are the following. First, bias could also be negative. Second, assuming only positive bias experts may have different ex ante probabilities of being biased. Finally, bias could be moderate as opposed to radical. We briefly discuss these possibilities in what follows.

One might consider a model with symmetric experts whose bias is either zero, maximally positive or maximally negative with respective probabilities $\{1 - \alpha_h - \alpha_l, \alpha_h, \alpha_l\}$. Consider the following so-called *two-sided* partitional communication strategy given by $\{t_r\}_{r=1}^{N-1}$. An unbiased sender sends the message m_i corresponding to the ith interval. A negatively (positively) biased expert always sends m_1 (m_N). The consultation strategy of R specifies stopping consultation as soon as receiving some message in $\{m_2, .., m_{N-1}\}$ and continuing otherwise for a maximum of

n rounds. Equilibrium final beliefs are now much more numerous than in our current model with one-sided bias. In the one-sided bias model, the set of equilibrium final beliefs of R contains n different elements. It features one final belief for any equilibrium message sequence ending with a message from $\{m_1, .., m_{N-1}\}$ and one final belief for the message sequence containing n times m_N. In the two-sided bias model, under the assumed consultation strategy, the set of equilibrium final beliefs of R contains $2n - 1$ elements. There is still one final belief for each equilibrium message sequence ending with a message from $\{m_2, .., m_{N-1}\}$. There are two final beliefs corresponding to unilateral sequences of n times, respectively, m_1 and m_N. In addition, there are now also $n - 1$ final beliefs corresponding to equilibrium message sequences containing exclusively messages m_1 and m_N and at least one of each. Given the above, the belief distribution attached to m_1 or m_N by a given sender might be more difficult to compute than in the one-sided bias model, thus rendering the analysis of sender incentives more complex. New questions arise within this model. Are negative and positive biases in any sense complementary? Given a total likelihood of bias $\tilde{\alpha} = \alpha_h + \alpha_l$, does R prefer α_h and α_l to be symmetric or asymmetric?

A second variation of our model is to go back to our original setup but assume that experts have different likelihoods of being (positively) biased. There is thus a high and a low credibility expert (S_h and S_l) featuring, respectively, $\beta_h > \beta_l$. Suppose that the experts' position is unobserved by the latter. Consider a partitional communication equilibrium in which S_h is consulted first and S_l is consulted only if S_h sends the highest message. This scenario appears more attractive than the reversed scenario in which S_l is consulted first for two reasons. Consulting S_h first allows to minimize the likelihood of having to consult again. Furthermore, a biased first expert's incentive to deviate to m_{N-1} is minimized by letting the second consulted expert be the low credibility one. This allows to maximize the fineness of the first sender's partitional communication, which in turn allows the second sender's communication to be very fine without incentivizing R to reverse the consultation order (see discussion in previous subsection).

A third interesting extension would be to allow for moderate as opposed to radical bias. Consider a model in which a given expert may be either unbiased, maximally positively biased or moderately positively biased with ex ante probabilities $\{1 - \gamma_l - \gamma_h, \gamma_h, \gamma_l\}$. In the second case, the expert's utility function is given by a while in the last case it is given by $-(a - (\omega + b))^2$, for some positive and not too large b. Recall that in our main setup, the radical bias of biased senders implies that we can focus on simple equilibria in which only the highest message is suspicious. These equilibria may not exist anymore and give way to equilibria in which multiple messages are suspicious. One interesting question would be whether for given $\tilde{\gamma} = \gamma_l + \gamma_h > 0$, R prefers the case of $\gamma_l = 0$ or that of $\gamma_h = 0$. Another question is whether for a given total probability $\tilde{\gamma}$ of biased experts, R favors symmetric values of γ_l and γ_h or asymmetric ones.

D. Simultaneous Consultation

In many contexts, multiple experts are consulted simultaneously rather than sequentially. It is for example common practice for the editor of an academic journal to simultaneously order reports on a given paper from multiple referees. The Amazon product page features a *reviews aggregator* summarizing grades awarded by anonymous reviewers.

While we give a detailed analysis of simultaneous consultation in online Appendix I, we here offer a preview of results. We show that semi-revealing communication and multiple simultaneous consultations are compatible in contrast to the sequential case. Given this positive finding, whether sequential or simultaneous consultation is optimal for R is unclear. While the first protocol saves consultation costs by offering the option to terminate consultation early, the second protocol yields more informative individual reporting. We explicitly compare the two protocols for $M = 2$ and show that for intermediate β and c high enough, sequential consultation is optimal. The intuition is as follows. A relatively high c implies that it is important to avoid ordering superfluous reports. An intermediate β implies that individual (semi-revealing) reports under simultaneous consultation are not radically more informative than individual (partitional) reports under sequential consultation, so that the informational benefit of the former protocol does not compensate for its cost-inefficiency.

Assume that R is exogenously committed to simultaneously consulting two experts. Given this, there exists an equilibrium with full revelation by both senders. As explained in online Appendix I, it is, however, not robust to the presence of noise in the messaging process. Given the assumed consultation strategy, there also exists a unique incentive compatible semi-revealing profile given by $\theta = \theta^*(\beta, 2)$. The key behind the possibility of semi-revealing communication is that a biased sender cannot affect the number of senders consulted and thus does not have an incentive (as in the sequential case) to deviate to $\theta^*(\beta, 2) - \epsilon$ to preempt further consultation. The semi-revealing scenario is furthermore robust to the presence of noise in the messaging process.

We wish to compare simultaneous and sequential consultation protocols in terms of R's welfare. Denote by $V_{sr}(\beta, 2)$ the gross expected payoff of R given that she simultaneously consults two experts using $\theta^*(\beta, 2)$. Under sequential consultation, we assume that R uses ϕ_2 while senders use the partitional strategy $\{t_r(\beta, \bar{N}, 2)\}_{r=1}^{\bar{N}-1}$, where $\bar{N}$ is shorthand for $\bar{N}(\beta, 2)$ which is the maximal achievable intervals number given β and ϕ_2. For each β, we shall consider values of c, such that given $\{t_r(\beta, \bar{N}, 2)\}_{r=1}^{\bar{N}-1}$, the sequential consultation strategy ϕ_2 is indeed incentive compatible.

In order to compare the two protocols, we compare the net expected payoff of R under each, and thus examine the inequality

$$\begin{aligned} V_{sr}(\beta, 2) - 2c \leq\ & V_p(\beta, \bar{N}, 2) \\ & - \left[1 + t_{\bar{N}-1}(\beta, \bar{N}, 2)(1 - \beta) + 1 - t_{\bar{N}-1}(\beta, \bar{N}, 2)\right] c, \end{aligned}$$

which simplifies to

$$V_{sr}(\beta, 2) - V_p(\beta, \bar{N}, 2) \leq \beta\, t_{\bar{N}-1}(\beta, \bar{N}, 2)\, c.$$

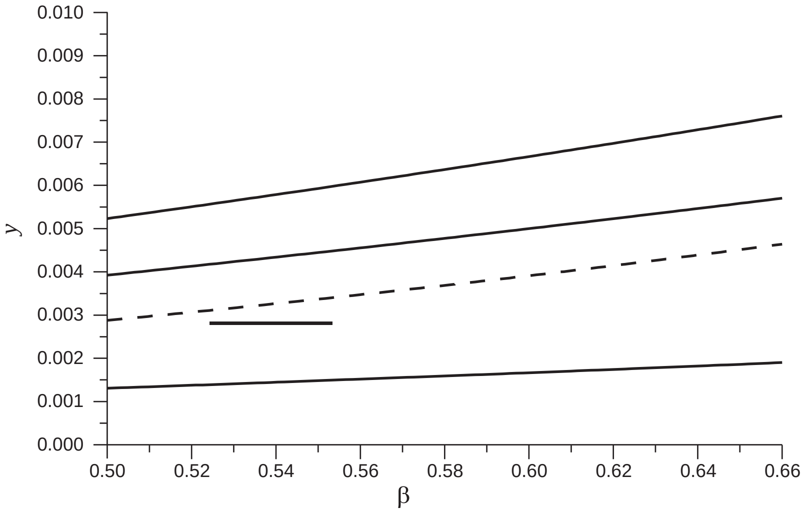

Figure 1

In Figure 1, the dashed curve corresponds to $V_{sr}(\beta, 2) - V_p(\beta, \bar{N}, 2)$, while solid curves correspond to $\beta t_{\bar{N}}(\beta, \bar{N}, 2)\, c$ for $c = 0.005, 0.01, 0.015, 0.02$. We consider $\beta \in \left(\frac{1}{2}, \frac{2}{3}\right]$, implying $\bar{N}(\beta, 2) = 3$. The figure reveals that for $c \geq 0.015$, sequential consultation dominates simultaneous consultation.

VI. Conclusion

Our main finding is that when information is elicited from multiple experts sequentially (and sequentially rationally), more extensive consultation comes at the price of less individually informative consultations. This trade-off lies at the core of the three sampler's curses identified: Access to more experts, a lower cost of consultation, as well as higher expert trustworthiness, may all hurt *R*. Further research should assess whether our findings survive in more general environments and also ought to draw implications for the optimal design of online review mechanisms. Incentives of unbiased and biased experts to provide potentially costly reviews also ought to be taken into consideration.

Appendix

A. *Proof of Lemma* 1

Step 1: This proves the *if* part of the statement contained in point (i) (i.e., sufficiency). Let *R* use τ_n. We first show that if θ satisfies (2), it is incentive compatible for unbiased senders. There are two cases to consider; either $\omega < \theta^*(\beta, n)$ or $\omega \geq \theta^*(\beta, n)$. In the first case, an unbiased sender knows that by sending a message ω, the final belief of *R* will be given by ω for sure. This is true as any equilibrium message profile containing at least one message $\omega < \theta^*(\beta, n)$ induces *R* to assign probability 1 to ω. An unbiased sender thus has no strict incentive to deviate for $\omega < \theta^*(\beta, n)$. On the other hand, if $\omega \geq \theta^*(\beta, n)$, sending $\theta^*(\beta, n)$ will lead to the final belief $\theta^*(\beta, n)$. Indeed, in equilibrium, R will receive for sure only message $\theta^*(\beta, n)$. Recall finally that any message profile to which Bayes' rule cannot be applied gives rise to belief $\omega = 0$. If $\omega \geq \theta^*(\beta, n)$ it is thus clear that an unbiased sender has no incentive to send a message $\omega' \neq \theta^*(\beta, n)$.

Step 2: This proves the *only if* part of the statement in point (i) (i.e., necessity). We now show that if θ does not satisfy (2), it is not incentive compatible for unbiased senders. Assume that the threshold θ does not satisfy (2). Recall that $B(\theta, \beta, n)$ is strictly larger than θ for $\theta < \theta^*(\beta, n)$ and strictly smaller than θ for $\theta > \theta^*(\beta, n)$. Let $\theta < \theta^*(\beta, n)$. It follows that the highest equilibrium message $m = \theta$ is s.t. $B(\theta, \beta, n) > \theta$. If $\omega = \theta + \epsilon$, for ϵ positive and sufficiently small, an unbiased sender can thus profitably deviate to sending message $\theta - \epsilon$. This is trivial if $n = 1$. If $n = 2$, there are two possibilities. Either R is the second expert to be consulted and the first sender sent θ, in which case the message profile $(\theta, \theta - \epsilon)$ is an equilibrium profile and gives rise to belief $\theta - \epsilon$. Or R is the first to be consulted, so that $\theta - \epsilon$ will discourage a further consultation and R's final belief will be $\theta - \epsilon$. Let $\theta > \theta^*(\beta, n)$. It follows that the highest equilibrium message $m = \theta$ is s.t. $B(\theta, \beta, n) < \theta$. If $\omega \geq \theta$, an unbiased sender can thus profitably deviate to sending message $\theta - \epsilon$, for strictly ϵ positive but sufficiently small. If $n = 1$, this is immediate. If $n = 2$, the argument is the same as that given above for $n = 2$.

Step 3: This proves point (ii). Note simply that if R consults only once for sure, then message $\theta^*(\beta, 1)$ triggers the highest possible expected belief among all available messages. Recall that any message to which Bayes' rule cannot be applied gives rise to belief $\omega = 0$.

Step 4: This proves point (iii). Suppose that $\omega = 0$. A biased sender can advantageously deviate to a message $\theta^*(\beta, 2) - \epsilon$, for ϵ positive but arbitrarily small. By doing so, he ensures that R adopts belief $\theta^*(\beta, 2) - \epsilon$ with probability 1. Indeed, either this is R's first consultation and she will not consult again or this is her second consultation and she first encountered a biased expert who sent $\theta^*(\beta, 2)$, in which case the message profile $(\theta^*(\beta, 2), \theta^*(\beta, 2) - \epsilon)$ leads to belief $\theta^*(\beta, 2) - \epsilon$. If, instead, the biased sender follows his equilibrium strategy and sends the high message $\theta^*(\beta, 2)$, he recognizes that R will adopt either belief $\theta^*(\beta, 2)$ or 0, both with strictly positive probability. Clearly, in expectation, he is better off sending $\theta^*(\beta, 2) - \epsilon$, for ϵ sufficiently small. ■

B. *Proof of Lemma* 2

Point (i) is proved in the main text. This proves point (ii). Assume that senders use the semi-revealing strategy profile $\theta^*(\beta, 1)$. Note that $\lim_{\beta\to 1} \theta^*(\beta, 1) = 1$. It follows immediately that

$$\lim_{\beta\to 1} v^1_{sr}(\beta) = 1 - \left(1 - \int_0^1 \left(\frac{1}{2} - \omega\right)^2 d\omega\right) = \frac{1}{12}.$$

In contrast, $\lim_{\beta \to 1} v^2_{sr}(\beta) = 0$. This holds true because the expected payoff of R, conditional on having received $m = \theta^*(\beta, 1)$ in the first consultation and stopping consultation, tends to 1 for β tending to 1. Formally, the latter expected payoff is given by

$$1 - \frac{(1-\beta)}{1+(1-\beta)} \frac{1}{\theta^*(\beta,1)} \int_0^{\theta^*(\beta,1)} (\theta^*(\beta,1) - \omega)^2 \, d\omega$$
$$- \frac{1}{1+(1-\beta)} \frac{1}{1-\theta^*(\beta,1)} \int_{\theta^*(\beta,1)}^1 (\theta^*(\beta,1)-\omega)^2 \, d\omega.$$

The corresponding closed form can be shown to converge to 1 as β tends to 1. The fact that this expected payoff tends to 1 immediately implies that the marginal value of an extra consultation tends to 0 conditional on a first consultation yielding $m = \theta^*(\beta, 1)$.

It follows that for β large enough, $v^1_{sr}(\beta) > v^2_{sr}(\beta)$. We may conclude that for β large enough, setting $c = v^2_{sr}(\beta) + \epsilon$ and picking ϵ sufficiently small, there exists an equilibrium featuring a unique consultation and the semi-revealing communication strategy profile $\theta^*(\beta, 1)$. Indeed, the marginal value of the first (second) consultation is higher (lower) than its cost. R can thus commit to consulting once and only once. ■

C. *Proof of Lemma* 3

Step 1: Recall that $B(t_{N-1}, \beta, n)$ is the expected value of the state conditional on senders using the N-intervals strategy profile $(t_1, .., t_{N-1})$ and having received n times the highest message m_N. In order for an unbiased sender to be indifferent between m_{N-1} and m_N at $\omega = t_{N-1}$, it must hold true that

$$(9) \qquad B(t_{N-1}, \beta, n) - t_{N-1} = t_{N-1} - \left(\frac{t_{N-1} + t_{N-2}}{2}\right) \Leftrightarrow t_{N-1} - t_{N-2} = 2(B(t_{N-1}, \beta, n) - t_{N-1}).$$

Clearly, in order for an unbiased sender to be indifferent between m_{i-1} and m_i at every t_i, for $i \in \{1, .., N-2\}$, we furthermore need $t_i - t_{i-1} = \Delta$, for $i \in \{1, .., N-2\}$, for some constant $\Delta > 0$. Given (9), it must be that $\Delta = 2(B(t_{N-1}, \beta, n) - t_{N-1})$.

Step 2: The ratio $\frac{t_{N-1}}{\Delta}$ must be an integer for there to exist an integer number of intervals of size Δ between t_{N-1} and 0. In order for there to be $N-1$ intervals to the left of the threshold t_{N-1}, it must hold true that $\frac{t_{N-1}}{\Delta} = N-1$. Using the fact that $\Delta = 2(B(t_{N-1}, \beta, n) - t_{N-1})$ this is equivalent to requiring

$$(10) \qquad \frac{B(t_{N-1}, \beta, n)}{t_{N-1}} = \frac{2(N-1)+1}{2(N-1)}.$$

Step 3: For any $n \geq 1$, $\beta \in (0, 1)$, and $N \geq 2$, there is a unique $t_{N-1} \in (0, \theta^*(\beta, n))$ that satisfies (10). First note that $\frac{2(2-1)+1}{2(2-1)} = \frac{3}{2}$ while $\frac{2(N-1)+1}{2(N-1)}$ is decreasing in N and converging to 1 as N tends to infinity. As to the behavior of $\frac{B(t_{N-1}, \beta, n)}{t_{N-1}}$, note first that

$$\frac{\partial \left(\frac{B(\theta, \beta, n)}{\theta}\right)}{\partial \theta}$$

$$= -\frac{\left(\theta^2 - \theta^2(1-\beta)^n - 2\theta + 2\theta(1-\beta)^n + 1\right)}{2\theta^2\left(\theta(1-\beta)^n - \theta + 1\right)^2} < 0, \quad \forall \theta, \beta, n.$$

Furthermore, given that $B(0, \beta, n) = B(1, \beta, n) = \frac{1}{2}$, it follows that $\lim_{\theta \to 0} \frac{B(\theta, \beta, n)}{\theta} = +\infty$. Finally, recall that by definition,

$$\frac{B(\theta^*(\beta, n), \beta, n)}{\theta^*(\beta, n)} = 1.$$

We may thus conclude that for any n, β, and $N \geq 2$, there is a unique $t_{N-1} \in (0, \theta^*(\beta, n))$ that satisfies (10). ■

D. *Proof of Lemma* 4

Step 1: Let $\omega = 0$. Assume a putative equilibrium featuring the unique U-IC N-intervals strategy profile $(t_1, \ldots, t_{N-1})$ and ϕ_n. The pool from which senders are drawn contains a total of M individuals. We wish to compute the distribution that a consulted sender assigns to his possible location l in the sampling sequence. Given n, the ex ante probability of being the first to be consulted is $\frac{1}{M}$. The ex ante probability of being the second picked is $\left(\frac{M-1}{M}\right)\frac{1}{M-1}(1-\beta)$. Summarizing, the ex ante likelihood of being the k th sender consulted is $\frac{1}{M}(1-\beta)^{k-1}$. The probability of being the *kth* sender consulted *conditional on being consulted* is thus

$$P(l = k \mid \omega = 0) = \frac{\frac{1}{M}(1-\beta)^{k-1}}{\sum_{i=1}^{i=n} \frac{1}{M}(1-\beta)^{i-1}} = \frac{\beta(1-\beta)^{k-1}}{1-(1-\beta)^n}.$$

The denominator in the middle expression denotes the total probability of being consulted (either first or second, etc.). We now want to examine the incentives of a biased sender. Assume that you are the kth expert consulted and $\omega = 0$. Sending m_{N-1} triggers for sure belief $\mu(m_{N-1})$ given by

$$\mu(m_{N-1}) = \frac{t_{N-1} + t_{N-2}}{2} = \left(\frac{4t_{N-1} - 2B(t_{N-1}, \beta, n)}{2}\right),$$

while m_N triggers either belief $\mu(m_1)$ given by

$$\mu(m_1) = \frac{t_1 + t_0}{2} = \frac{2\big(B(t_{N-1}, \beta, n) - t_{N-1}\big)}{2},$$

if R meets an unbiased sender at any point in the future, or $\mu(m_N) = B(t_{N-1}, \beta, n)$ if R meets no unbiased sender at any point until she stops consulting. Given that you are the kth expert consulted and that you send m_N, the likelihood that R meets no unbiased sender at any point until she stops consulting is $(1 - \beta)^{n-k}$. So the expected belief triggered by sending m_N is given by

$$(1 - \beta)^{n-k} \mu(m_N) + (1 - (1 - \beta)^{n-k}) \mu(m_1).$$

Let $E(\mu \mid m_i, \omega)$ denote the expected belief of R, as computed by an expert not knowing his position in the consultation sequence, if he sends message m_i and the state is ω. Clearly,

$$E(\mu \mid m_N, 0) = \sum_{k=1}^{n} P(l = k \mid \omega = 0)\Big((1 - \beta)^{n-k} \mu(m_N) + \big(1 - (1 - \beta)^{n-k}\big)\mu(m_1)\Big).$$

Note that

$$P(l = k \mid \omega = 0)(1 - \beta)^{n-k} = \frac{\beta(1 - \beta)^{n-1}}{1 - (1 - \beta)^n}$$

so that

$$\sum_{k=1} P(l = k \mid \omega = 0)(1 - \beta)^{n-k} = \frac{n\beta(1 - \beta)^{n-1}}{1 - (1 - \beta)^n}.$$

We may thus write

$$E(\mu \mid m_N, 0) = \frac{n\beta(1 - \beta)^{n-1}}{1 - (1 - \beta)^n}\mu(m_N) + \left(1 - \frac{n\beta(1 - \beta)^{n-1}}{1 - (1 - \beta)^n}\right)\mu(m_1).$$

Finally, the expected belief triggered by message m_{N-1} is simply

$$E(\mu \mid m_{N-1}, 0) = \mu(m_{N-1}) = \left(\frac{4t_{N-1} - 2B(t_{N-1}, \beta, n)}{2}\right).$$

Indeed, sending m_{N-1} leads R to immediately stop consulting.

Step 2: Assume that R uses the strategy ϕ_n. The U-IC N-intervals strategy profile $(t_1, .., t_{N-1})$ is thus B-IC, given $\omega = 0$, if and only if

$$\left(1 - \frac{n\beta(1-\beta)^{n-1}}{1-(1-\beta)^n}\right)\left(\frac{2\big(B(t_{N-1},\beta,n) - t_{N-1}\big)}{2}\right)$$
$$+ \frac{n\beta(1-\beta)^{n-1}}{\left(1-(1-\beta)^n\right)} B(t_{N-1},\beta,n)$$
$$\geq \left(\frac{4t_{N-1} - 2B(t_{N-1},\beta,n)}{2}\right),$$

which, after reorganizing, reduces to

$$(11) \qquad \frac{B(t_{N-1},\beta,n)}{t_{N-1}} \geq \left(\frac{3}{2} - \frac{1}{2}\frac{n\beta(1-\beta)^{n-1}}{1-(1-\beta)^n}\right).$$

Step 3: Assume that $\omega \in (0, t_1)$. Then clearly, the expected payoff to a biased sender of sending m_N is the same as in the case of $\omega = 0$. Indeed, if R encounters an unbiased sender she receives m_1. On the other hand, the payoff attached to m_{N-1} is the same as in the case of $\omega = 0$. It follows that the U-IC N-intervals strategy profile $(t_1, .., t_{N-1})$ is B-IC, given $\omega \in (0, t_1)$, if and only if (11) holds.

Assume instead that $\omega \in (t_i, t_{i+1})$ for $i \in \{1, .., N-2\}$. Then the payoff to a biased sender of sending m_N is larger than in the case of $\omega = 0$. Indeed, if R encounters an unbiased sender she receives m_{i+1}, which triggers action $\frac{t_i + t_{i+1}}{2} > \frac{t_1}{2}$. On the other hand, the payoff attached to m_{N-1} is the same as in the case of $\omega = 0$. It follows that if the U-IC N-intervals strategy profile $(t_1, .., t_{N-1})$ is B-IC, given $\omega = 0$, then it is B-IC, given $\omega \in (t_i, t_{i+1})$, for $i \in \{1, .., N-2\}$.

We may conclude that the U-IC N-intervals strategy profile $(t_1, .., t_{N-1})$ is B-IC if and only if (11) holds. ■

E. *Proof of Proposition* 1

If only one sender is available, R can obtain $V_{sr}(\beta, 1)$, which is a continuous and monotonously increasing function of β tending to one for β tending to one. On the other hand, for any $\beta \in (0, 1)$ and finite N, $V_p(\beta, N, 2)$ is a continuous and increasing function of β and $\lim_{\beta\to 1} V_p(\beta, N, 2) < 1$. Finally, we know that $\overline{N}(\beta, 2) = 2$ for $\beta \geq \frac{2}{3}$. ■

F. *Proof of Proposition* 2

Step 1: We first consider the semi-revealing scenario. We know from Lemma 2 that setting $c = v_{sr}^2(\beta) + \epsilon$, for β sufficiently large and ϵ sufficiently small there exists an equilibrium featuring a unique consultation and the semi-revealing communication strategy profile $\theta^*(\beta, 1)$. Indeed, the marginal value of the first (second) consultation is higher (lower) than its cost. R can thus commit to consulting once and only once. Second, for any $c < v_{sr}^2(\beta)$ (and thus for $c = 0$), there exists no equilibrium featuring such a communication strategy because R deviates to consulting again after a first consultation yielding the high message $m = \theta^*(\beta, 1)$.

Step 2: We now consider the partitional scenario. For $\beta \geq \frac{2}{3}$, $\bar{N}(\beta, 2) = 2$. Thus, if there exists an equilibrium featuring a partitional communication strategy and multiple consultations, this involves two intervals. For $c = 0$, assuming a two-intervals communication strategy profile, the sequential consultation strategy ϕ_2 is furthermore trivially incentive compatible.

Step 3: For $c = v_{sr}^2(\beta) + \epsilon$, R can thus obtain expected utility $V_{sr}(\beta, 1) - v_{sr}^2(\beta) - \epsilon$, while for $c = 0$, she obtains expected utility $V_p(\beta, 2, 2)$. Note that

$$V_{sr}(\beta, 1) - v_{sr}^2(\beta) - \epsilon > V_p(\beta, 2, 2) \quad \Leftrightarrow \tag{12}$$

$$V_{sr}(\beta, 1) - V_p(\beta, 2, 2) > v_{sr}^2(\beta) + \epsilon.$$

Now, note that $\lim_{\beta \to 1} t_1(\beta, 2, 2) = \frac{1}{2}$, implying in turn that $\lim_{\beta \to 1} V_p(\beta, 2, 2) < 1$. Also, $\lim_{\beta \to 1} V_{sr}(\beta, 1) = 1$ given that $\lim_{\beta \to 1} \theta^*(\beta, 1) = 1$. Finally, $\lim_{\beta \to 1} v_{sr}^2(\beta) = 0$. It follows that given β large enough, inequality (12) holds for ϵ small enough. ■

G. *Proof of Proposition* 3

Step 1: Assuming that R follows the sequential consultation strategy ϕ_2, there exists an incentive compatible N-intervals strategy profile if and only if $\beta \leq \frac{2}{N}$. When shifting the value of β from $\frac{2}{N} - \epsilon$ to $\frac{2}{N} + \epsilon$, for ϵ small enough, the finest incentive compatible partitional communication thus changes from N to $N - 1$. Note also that given $\beta \in (0, 1)$ and $N \geq 2$, ϕ_2 is incentive compatible, given $\{t_r(\beta, N, 2)\}_{r=1}^{N-1}$, if c is small enough.

Step 2: We know the following from Lemma 6. First, for any $N \geq 2$, $V_p(\beta, N, 2)$ is continuous and increasing in β. Second, for any $N \geq 2$ and $\beta \in (0, 1)$, $V_p(\beta, N, 2) \geq V_p(\beta, N - 1, 2)$. Given the discreteness of the intervals number N, $V_p(\beta, N, 2)$ increases in discontinuous jumps as N increases. Consider now the inequality

$$V_p\left(\frac{2}{N} - \epsilon, N, 2\right) - \left[1 - \left(\frac{2}{N} - \epsilon\right) + t_{N-1}\left(\frac{2}{N} - \epsilon, 2, N\right)\right] c$$
$$> V_p\left(\frac{2}{N} + \epsilon, N - 1, 2\right) - \left[1 - \left(\frac{2}{N} + \epsilon\right) + t_{N-2}\left(\frac{2}{N} + \epsilon, 2, N - 1\right)\right] c.$$

The left-hand side corresponds to the net expected payoff of R given $\beta = \frac{2}{N} - \epsilon$, the N-intervals profile $\{t_r(\beta, N, 2)\}_{r=1}^{N-1}$ and the consultation strategy ϕ_2. The right-hand side corresponds to the net expected payoff of R given $\beta = \frac{2}{N} + \epsilon$, the $(N-1)$-intervals profile $\{t_r(\beta, N-1, 2)\}_{r=1}^{N-2}$, and ϕ_2. Clearly, for ϵ and c small enough, the above is satisfied. In the limit, for c and ϵ arbitrarily small, the inequality is satisfied because

$$V_p\left(\frac{2}{N}, N, 2\right) > V_p\left(\frac{2}{N}, N-1, 2\right),$$

which is true given that $V_p(\beta, N, 2)$ is increasing in N. ■

H. *Proof of Proposition* 4

We first prove point (i). If only one sender is available, R can obtain $V_{sr}(\beta, 1)$, which is a continuous and monotonously increasing function of β tending to one for β tending to one. Recall also that for any β and finite N,

$$\lim_{n\to\infty} V_p(\beta, N, n) = 1 + \frac{1-3N}{6N^2} < 1.$$

We now prove point (ii). Recall that $\overline{N}(\beta, 2) = 3$ if $\beta \in \left(\frac{1}{2}, \frac{2}{3}\right]$. On the other hand, $\overline{N}(\beta, n) = 2$ given any (β, n) s.t. $\beta \geq \frac{1}{3}$ and $n \geq 3$. Finally, it is easily shown that for $\beta \in \left(\frac{1}{2}, \frac{2}{3}\right]$, $V_p(\beta, 3, 2) > \lim_{n\to\infty} V_p(\beta, 2, n)$. ■

I. *Proof of Proposition* 5

Step 1: We first consider the semi-revealing scenario. We build on a generalization of Lemma 2 to the case of more than two experts. Assume that all M senders use $\theta^*(\beta, 1)$. As in point (i) of Lemma 2, for c sufficiently low a unique consultation is not incentive compatible because consulting again after the high message $m = \theta^*(\beta, 1)$ is informative. The equivalent of point (ii) reads as follows: Given $c = \tilde{v}^2_{sr}(\beta, M) + \epsilon$, for β large enough and ϵ sufficiently small a unique consultation by R is incentive compatible. The proof of this is as follows. First, given $c = \tilde{v}^2_{sr}(\beta, M) + \epsilon$, there is by definition no sequentially rational consultation strategy $\phi_n \in \{\phi_2, .., \phi_M\}$, so that R does not deviate to a second consultation after a first consultation yielding message $m = \theta^*(\beta, 1)$. Second, note that $\lim_{\beta\to 1} v^1_{sr}(\beta) = \frac{1}{12}$. Third, for any finite M, $\lim_{\beta\to 1} \tilde{v}^2_{sr}(\beta, M) = 0$. This holds true because the expected payoff of R conditional on stopping consultation after receiving $m = \theta^*(\beta, 1)$ in the first consultation tends to one for β tending to one. We may now conclude. Given β large enough, it thus holds true that $v^1_{sr}(\beta) > \tilde{v}^2_{sr}(\beta, M)$ and it follows that one can find an ϵ small enough such that $v^1_{sr}(\beta) \geq \tilde{v}^2_{sr}(\beta, M) + \epsilon \geq \tilde{v}^2_{sr}(\beta, M)$.

Step 2: We now consider the case of partitional communication strategies. For $n \geq 2$ and $\beta \geq \frac{2}{3}$, $\overline{N}(\beta, n) = 2$. Thus, if there exists an equilibrium featuring a partitional communication strategy and multiple consultations, this involves two intervals. For $c = 0$, assuming a two-intervals strategy profile, the sequential consultation strategy ϕ_M is furthermore trivially incentive compatible.

Step 3: For $c = \tilde{v}^2_{sr}(\beta, M) + \epsilon$, R thus obtains expected utility $V_{sr}(\beta, 1) - \tilde{v}^2_{sr}(\beta, M) - \epsilon$ while for $c = 0$, she obtains expected utility $V_p(\beta, 2, M)$. Note that

$$(13) \qquad V_{sr}(\beta, 1) - \tilde{v}^2_{sr}(\beta, M) - \epsilon > V_p(\beta, 2, M) \quad \Leftrightarrow$$

$$V_{sr}(\beta, 1) - V_p(\beta, 2, M) > \tilde{v}^2_{sr}(\beta, M) + \epsilon.$$

Recall that $\lim_{\beta \to 1} V_{sr}(\beta, 1) = 1$. Furthermore, for any finite M, $\lim_{\beta \to 1} V_p(\beta, 2, M) < 1$ and also

$$\lim_{M \to \infty} \lim_{\beta \to 1} V_p(\beta, 2, M) < 1.$$

Finally, for any finite M, $\lim_{\beta \to 1} \tilde{v}^2_{sr}(\beta, M) = 0$. It follows that for β large enough, the inequality (13) holds for ϵ small enough. ■

J. *Proof of Proposition* 6

Step 1: Assuming that R follows the consultation strategy ϕ_n, the N-intervals communication strategy $\{t_r(\beta, N, n)\}_{r=1}^{N-1}$ is incentive compatible for senders if and only if $\beta \leq \bar{\beta}(N, n)$. The function $\bar{\beta}(N, n)$ is furthermore decreasing in N. Given ϕ_n, when shifting the value of β from $\bar{\beta}(N, n) - \epsilon$ to $\bar{\beta}(N, n) + \epsilon$, for ϵ small enough, the finest incentive compatible partitional communication changes from N intervals to $N - 1$ intervals. Note also that given $\beta \in (0, 1)$ and $N \geq 2$, ϕ_n is incentive compatible given $\{t_r(\beta, N, n)\}_{r=1}^{N-1}$ if c is small enough.

Step 2: Recall the following two facts. For a given $N \geq 2$ and $n \geq 1$, $V_p(\beta, N, n)$ is continuous and increasing in β. For any $N \geq 2$ and $\beta \in (0, 1)$, $V_p(\beta, N, n)$ increases in N in discontinuous jumps. Consider the inequality

$$\begin{aligned} & V_p\big(\bar{\beta}(N, n) - \epsilon, N, n\big) - \big[1 - \big(\bar{\beta}(N, n) - \epsilon\big) + t_{N-2}\big(\bar{\beta}(N, n) - \epsilon, N, n\big)\big]\, c \\ > \ & V_p\big(\bar{\beta}(N, n) + \epsilon, N - 1, n\big) \\ & - \big[1 - \big(\bar{\beta}(N, n) + \epsilon\big) + t_{N-2}\big(\bar{\beta}(N, n) + \epsilon, N - 1, n\big)\big]\, c, \end{aligned}$$

where each side corresponds to the expected payoff of R given the assumed value of β, the finest feasible partitional strategy and implied expected consultation cost. Clearly, for ϵ and c small enough, the above is satisfied. Indeed, for c and ϵ arbitrarily small, the inequality is satisfied because

$$V_p\big(\bar{\beta}(N, n), N, n\big) > V_p\big(\bar{\beta}(N, n), N - 1, n\big),$$

which is always true as $V_p(\beta, N, n)$ increases in N. ■

REFERENCES

Avery, Christopher, and Margaret Meyer. 2012. "Reputational Incentives for Biased Evaluators." http://www.nuffield.ox.ac.uk/teaching/Economics/Incentive/amfrank.pdf.

Battaglini, Marco. 2002. "Multiple Referrals and Multidimensional Cheap Talk." *Econometrica* 70 (4): 1379–1401.

Battaglini, Marco. 2004. "Policy advice with imperfectly informed experts." *B. E. Journal of Theoretical Economics* 4 (1).

Blackwell, D. 1951. Comparison of experiments." In *Proceedings of the Second Berkeley Symposium on Mathematical Statistics and Probability*, edited by J. Neeman, 93–102. Los Angeles. University of California Press.

Crawford, Vincent P., and Joel Sobel. 1982. "Strategic Information Transmission." *Econometrica* 50 (6): 1431–51.

Ely, Jeffrey C., and Juuso Välimäki. 2003. "Bad Reputation." *Quarterly Journal of Economics* 118 (3): 785–814.

Farrell, Joseph. 1993. "Meaning and Credibility in Cheap-Talk Games." *Games and Economic Behavior* 5 (4): 514–31.

Gerardi, Dino, Richard McLean, and Andrew Postlewaite. 2009. "Aggregation of expert opinions." *Games and Economic Behavior* 65 (2): 339–71.

Kawamura, Kohei. 2013. "Eliciting information from a large population." *Journal of Public Economics* 103 (1): 44–54.

Krishna, Vijay, and John Morgan. 2001. "A Model of Expertise." *Quarterly Journal of Economics* 116 (2): 747–75.

Li, Ming. 2010. "Advice from Multiple Experts: A Comparison of Simultaneous, Sequential, and Hierarchical Communication." *B. E. Journal of Theoretical Economics* 10 (1).

Matthews, Steven A., Masahiro Okuno-Fujiwara, and Andrew Postlewaite. 1991. "Refining cheap talk equilibria." *Journal of Economic Theory* 55 (2): 247–73.

Morgan, John, and Phillip C. Stocken. 2003. "An Analysis of Stock Recommendations." *RAND Journal of Economics* 34 (1): 183–203.

Morgan, John, and Phillip C. Stocken. 2008. "Information Aggregation in Polls." *American Economic Review* 98 (3): 864–96.

Morris, Stephen. 2001. "Political Correctness." *Journal of Political Economy* 109 (2): 231–65.

Ottaviani, Gianmarco, and Peter Norman Sørensen. 2001. "Information aggregation in debate: Who should speak first?" *Journal of Public Economics* 81 (3): 393–421.

Sarvary, Miklos. 2002. "Temporal Differentiation and the Market for Second Opinions." *Journal of Marketing Research* 39 (1): 129–36.

Sarvary, Miklos, and Philip M. Parker. 1997. "Marketing Information: A Competitive Analysis." *Marketing Science* 16 (1): 24–38.

Sobel, Joel. 1985. "A Theory of Credibility." *Review of Economic Studies* 52 (4): 557–73.

Wolinsky, Asher. 2002. "Eliciting information from multiple experts." *Games and Economic Behavior* 41 (1): 141–60.