Communication in a Complicated World

Steve Callander* Nicolas S. Lambert† Niko Matouschek‡

September 21, 2017

VERY PRELIMINARY AND INCOMPLETE

Abstract

Experts play a central role in many economic, political, and social interactions, and the ability of experts to sway decisions in their favor is widely documented. A key ingredient of this ability is the difficulty non-expert decision makers have in de-biasing an expert's advice. For instance, even if you know your mechanic is recommending excessive repairs, how do you know what repair your car really needs? Canonical models of expertise do not capture this property. In canonical representations, expertise consists of a single piece of information and, as a result, advice can be perfectly de-biased and experts are unable to sway decisions in their favor. We develop an alternative model of expertise, one in which expert knowledge consists of a continuum of correlated states and in which advice can be de-biased only imperfectly, even when advice is hard information. We show that this richer conception of expertise provides the expert with leverage to sway decisions in her favor. Yet to do so, the expert must communicate in a new form, conveying ostensibly redundant beyond a simple recommendation. We show that this power to sway outcomes benefits the expert for any degree of complexity beyond that assumed in canonical models, although peaking for decisions of moderate complexity.

1 Introduction

The power of experts to sway decisions in their favor has long been noted and documented. Bureaucrats manipulate their "political masters" into adopting policies they favor (Weber [1922]). Division managers manipulate headquarters into funding too many projects (Milgrom and Roberts [1988]). Realtors manipulate homeowners into selling too quickly and cheaply (Levitt and Syverson [2008]). And OBGYNs manipulate patients into having too many C-sections (Gruber and Owings [1996]).[1]

*Stanford University, Graduate School of Business; sjc@stanford.edu.

†Stanford University, Graduate School of Business; nlambert@stanford.edu.

‡Northwestern University, Kellogg; n-matouschek@kellogg.northwestern.edu.

[1]This behavior is so widespread that healthcare economists have coined the term "physician-induced demand" to refer to the additional demand physicians generate by manipulating their patients into over-consuming their services.

We contend that experts are able to sway decisions in their favor because it is difficult for laypeople to "de-bias" their advice. Even when you know that the advice you got from an expert serves her interests, it may be very difficult to deduce what course of action best serves your interests. For instance, you may know that your OBGYN has an incentive to recommend an expensive delivery procedure. But if she recommends a C-section, how do you figure out what you really need?

Standard models of communication, such as Crawford and Sobel [1982] and Milgrom and Roberts [1986], do not capture the difficulty of de-biasing advice. As a result, they predict that the power of experts to sway decisions in their favor is limited. Such models assume that if your doctor truthfully revealed the outcome of any one procedure, you would learn the outcomes of all procedures and thus be able to figure out which one is best for you. You would be able to so because the entirety of the doctor's expertise is assumed to consist of a single piece of information that can be inferred from a single decision-outcome pair. In such models, therefore, a single piece of information turns a layperson into an expert.

We argue that this assumption understates the informational advantage of experts and robs them of much of their power to sway decisions in their favor. To capture the informational advantage of experts more fully, we explore a setting in which an expert is privately informed about a continuum of correlated states, each one of which determines the outcome of a corresponding decision. For each state, the expert can reveal its value or not but she cannot misrepresent it. Her information is therefore verifiable or "hard."

Because the expert's informational advantage is so rich, her advice cannot be de-biased perfectly. And because her advice cannot be de-biased perfectly, the expert is able to sway decisions in her favor. She is most effective in doing so for decisions that are neither too simple nor too complex but somewhere in-between. To sway decisions in her favor, however, the expert has to designs her advice appropriately. She cannot simply make a recommendation by providing information about the decision she wants the advisee to make. Instead, she also has to provide information about some, but not all, of the decisions she does not want him to make.

To state the context and contributions of our paper more precisely, it is helpful to start with the standard example that is commonly used to illustrate both, the canonical model of communication with soft information in Crawford and Sobel [1982] and the canonical model of communication with hard information in Milgrom and Roberts [1986]. A sender is privately informed about state θ and a receiver has the right to make decision d. The state and the decision determine the outcome $x = d - \theta$, which, in turn, determines the receiver's utility $-x^2$. In the example with soft information, the sender's utility is given by $-(x-b)^2$, where $b > 0$ measures the sender's bias. And in the example with hard information, the receiver's utility is given by $-d$. In either case, after the sender learns her information she sends a message to the receiver who then updates his beliefs and makes the decision that he expects to be best for him.

A key feature of this setting is that the outcome function is invertible. If the sender were to reveal the outcome generated by a single decision, the receiver would be able to infer θ and thus work out the outcome generated by any decision. The sender therefore cannot induce the receiver to make her preferred decision in every state. In fact, she cannot even induce him to bias his decisions in her favor at all. The receiver either ends up making his preferred decision in every state, as in the unique equilibrium with hard information, or he makes

his preferred decision on average, as in any equilibrium with soft information (for a more detailed description of the standard example and its equilibria see, for instance, Sections 2.2 and 2.5 in Gibbons et al. [2013]).

The invertibility of the outcome function is not an artifact of the standard example. It follows from the assumptions that the receiver's preferred decision is unique, which holds in the general models in Crawford and Sobel [1982] and Milgrom and Roberts [1986] and most of the papers that build on their work, and the assumption that the receiver's preferred decision is strictly monotonic in the state, which follows from the standard sorting condition $\partial^2 u_R/\partial d\partial\theta > 0$, where u_R is the receiver's utility.[2]

The prediction that the sender cannot sway the average decision is equally common. As Kamenica and Gentzkow [2011] point out, it has to hold in any model in which the receiver is a Bayesian and his decision rule is a linear function of his posterior belief about the average state, which is the case in most applications with both soft and hard information (see, for instance, Gibbons et al. [2013]).

To focus on the role of the expert's information, we stay as close as possible to the standard example we just sketched. We find it convenient to work with the specification in which information is hard, as in Milgrom and Roberts [1986]. We also adopt the assumption that the sender's utility is given by $-d$. The expert's bias is therefore extreme, she prefers smaller decisions, no matter what the states are. As such, this assumption stacks the cards against the expert's ability to sway decisions in her favor.

The novelty of our model is the outcome function. Instead of assuming that the outcome function is $x = d - \theta$, we assume it is $x = d - \theta(d)$, or equivalently $x = d + \theta(d)$, where $\theta(\cdot)$ is the realized path of a Brownian motion with zero drift and scale $\sigma > 0$, and for which $\theta(0) = 0$. The sender knows the full mapping from decisions into outcomes. In contrast, prior to receiving advice, the receiver only knows the parameters of the Brownian motion and the single outcome value at a status quo decision.

The key feature of this specification is that the outcome function is partially invertible. If the receiver learns the outcomes generated by some set of decisions, he cannot perfectly infer the outcomes generated by the other, "unknown" decisions. He can, however, form more precise beliefs about the outcomes generated by those decisions. His beliefs will be more precise, the closer the unknown decision is to a known one. And they will be more precise, the less complex the expert's information is, that is, the smaller the variance of the Brownian motion is.

We start by exploring a class of "one-decision equilibria" in which the sender reveals the outcome of just a single decision. The outcome of the decision is better for the receiver than the status quo but not as good as his preferred outcome. The receiver then faces a choice. He can rubberstamp the sender's recommendation and ensure himself a suboptimal but certain outcome. Or he can learn from the recommendation but make a different decision, one that will give him a better outcome on average but also expose him to risk of an uncertain outcome.

[2]The literature differs in whether utilities are defined over outcomes or decisions. Both specifications are analogous. Suppose the sender's and the receiver's utilities are given by $v_S(d,\theta)$ and $v_R(d,\theta)$ and that the receiver's preferred decision, $d(\theta)$, is unique. Moreover, suppose $\partial^2 v_R/\partial d\partial\theta > 0$ so that $d(\theta)$ is strictly increasing in the state θ. Since $d(\theta)$ is unique, one can define the outcome function $x(d,\theta) = d - d(\theta)$ and redefine utilities as $u_S(x,\theta) = v_S(x(d,\theta) + d(\theta),\theta)$ and $u_R(x,\theta) = v_R(x(d,\theta) + d(\theta),\theta)$. Moreover, since $d(\theta)$ is strictly increasing, the outcome function is invertible.

Anticipating this choice, the sender recommends a decision that is just good enough for the receiver to prefer rubberstamping her recommendation to experimenting with a different decision.

The one-decision equilibrium is similar to the equilibria of the standard example we sketched above, in that the sender communicates only a single piece of information. It therefore does not exploit the rich informational advantage the sender enjoys in our setting. As a result, it too fails to sway the receiver's decision in favor of the sender: on average, the decision the receiver ends up making is no different from the one he would make in the absence of any advice.

Our central result is that the sender can do better if she designs her advice in a way that does exploit her informational advantage. To this end, we characterize a class of "interval equilibria" in which the sender once again recommends a decision that the receiver finds optimal to rubberstamp. This decision is better for the sender than the one she recommends in the one-decision equilibrium and, on average, it is better than the one the receiver would make in the absence of any advice.

Even though the sender is better off, though, the receiver is not: the recommended decision in the interval equilibrium is worse for him than the one in the one-decision equilibrium. Why then does he still prefer rubberstamping the recommendation to experimenting with a different decision?

The reason is that the sender does not only provide information about the decision she wants the receiver to make but also about neighboring decisions she does not want him to make. The purpose of this additional information is not to provide more information about the outcome of the recommended decision, which is perfectly revealed by the recommendation. Instead, the purpose is to convince the receiver that there are no decisions close to the recommended one that are any better for him and that decisions that are further away, and that receiver has not revealed, are too risky and not worth experimenting with.

In the interval equilibrium, the sender provides a lot more information than in the one-decision equilibrium. He also, however, provides a different type of information and thus gives a different type of advice. It is natural to interpret the sender's message in the one-decision equilibrium as a recommendation, as we did above. The sender's advice is therefore what Sobel [2013] calls "conative," it is an instruction on what to do. The sender's message in the interval equilibrium also contains a recommendation and therefore also has a conative element. The rest of the sender's message, however, is purely " referential," it merely provides facts about the environment (Sobel [2013]). It is the combination of conative and referential advice that gives the sender the power to sway decisions in the interval equilibrium.

Having explored why and how the sender is able to sway decisions in her favor, we conclude by examining when she is most effective in doing so. We focus on the complexity of the issue the expert is advising on, which we measure with the variance of the Brownian motion. More complexity allows the sender to induce smaller decisions in the interval equilibrium. It also, however, induces him to make smaller decisions in the absence of any advice. As a result, the sender's power to sway decisions in her favor—the difference between the average decision with and without advice—is not monotonic and instead peaks for intermediate levels of complexity. Our model therefore suggests that experts may be most powerful for issues that are neither very complex, nor very simple, but somewhat in-between.

Related Literature

(To be added.)

2 Model

A sender and a receiver play a game of strategic communication that is described by the following assumptions.

Technology: There is a set of decisions $\mathcal{D} = \{d_0, d_1, \ldots, d_n\}$, where $d_0 = 0$ and

$$d_i = d_{i-1} + \frac{1}{\sqrt{n}} \text{ for } i = 1, 2, ..., n.$$

The $n+1$ decisions therefore span the interval $[0, \sqrt{n}]$, with each being equally far from its neighbors.

Decision $d \in \mathcal{D}$ generates outcome $x \in \mathbb{R}$ according to outcome function $x(d)$. The outcome function is the realization of a random walk with drift. Specifically, $x(0) = 0$ and

$$x(d_i) = x(d_{i-1}) + \mu_n + \sigma_n \theta_i \text{ for } i = 1, ..., n,$$

with $\mu_n \equiv \mu/\sqrt{n}$ and $\sigma_n \equiv \sqrt{\sigma^2/\sqrt{n}}$, where each θ_i is independently drawn from a standard normal distribution and μ and σ are strictly positive parameters. As n goes to infinity, the set of feasible decisions becomes the non-negative half line $\mathbb{R}_+$ and the outcome function becomes the realization of a Brownian motion with drift μ and scale σ.

We denote the states of the world by $\theta = (\theta_1, ..., \theta_n)$ and the set of possible states by $\Theta \equiv \mathbb{R}^n$.

Preferences: The receiver's utility is given by $u_R(x) = -(x-b)^2$, where b is a strictly positive parameter. The sender's utility is given by $u_S(d) = -d$. The incentive conflict between the sender and the receiver is therefore extreme: while the receiver wants the decision to be adapted to the state, the sender always prefers smaller decisions to larger ones. The receiver's utility follows the canonical models of the literature.

Information: The sender is privately informed about states θ_i, and thus outcomes $x(d_i)$, for $i = 1, ..., n$. All other information is public, including the outcome of decision $d_0 = 0$ and the parameters of the process that generates the outcome function.

Communication: For each decision d_1 to d_n the sender has a piece of hard information that verifies its outcome. The sender can hide or reveal any number of these pieces of information but she cannot fake them. Her message is therefore given by $m = (m_1, m_2, ..., m_n)$, where $m_i \in \{x(d_i), \emptyset\}$.

Timing: First, nature draws the outcome function and the sender learns the realization. Second, the sender sends her message. Third, the receiver updates his beliefs and chooses a decision. Finally, the sender and the receiver realize their payoffs and the game ends. Note that neither party has commitment power: the sender cannot commit to a message rule and the receiver cannot commit to a decision rule.

Solution concept: The solution concept is Perfect Bayesian Equilibrium. We require off path beliefs to be consistent with any hard information that has been revealed by the sender. If the sender deviates and, for instance, reveals that decision d_2 generates outcome five, the receiver cannot ignore this evidence and believe that it generates outcome six instead.

Formally, a strategy for the sender is a mapping M from the set of all possible states Θ to the set of all possible messages $\mathcal{M}$. A strategy for the receiver is a mapping D from the set of all possible messages $\mathcal{M}$ to the set of all possible decisions $\mathcal{D}$. The receiver's beliefs are described by a belief mapping B that assigns belief $B(m)$—a probability distribution over states, $B(m) \in \Delta(\Theta)$—to every possible message $m \in \mathcal{M}$. Strategies M and D and belief mapping B form a Perfect Bayesian Equilibrium if *(i)* the sender's strategy M maximizes her utility given D, *(ii)* the receiver's strategy D maximizes his expected utility given B, and *(iii)* given M, on path, the receiver's beliefs satisfy Bayes' rule whenever possible and off path, they are consistent with any hard information that has been revealed.

Deviations and off-path beliefs: There are, in general, two ways in which the sender can deviate from her strategy: "on schedule" and "off schedule." To define these terms, let M^* denote a specific strategy for the sender and let $M^*(\theta)$ denote the message the strategy calls for in state $\theta \in \Theta$. Also, let $\mathcal{M}^*$ denote the set of messages that are consistent with the strategy. We say that the sender deviates "on schedule" if, in some state θ, she sends a message m that is in $\mathcal{M}^*$ but different from $M^*(\theta)$. And we say that she deviates "off schedule" if she sends a message that is different from $M^*(\theta)$ and not in $\mathcal{M}^*$. The receiver, therefore, can detect off schedule deviations but not on schedule ones.

Suppose now that the sender deviates off schedule and let d^* denote the best decision for the receiver among the status quo decision and any decision the sender may have revealed. We say that the receiver's off path beliefs are "skeptical" if, following an off schedule deviation, the following holds: *(i)* if the sender did not reveal all decisions larger than d^*, the receiver believes that the largest unrevealed decision is best for him and *(ii)* if the sender did reveal all decisions to the right of d^*, the receiver beliefs that d^* is the best decision for him.

This definition of skeptical beliefs is akin to the "skeptical posture" in Milgrom and Roberts [1986]. It captures the notion that the sender cannot induce more favorable off path beliefs by hiding information.

Discussion of the assumptions

To justify the main assumptions of the model we just described, it is helpful to relate it to what has become the standard example, and workhorse model, of games with strategic communication. The standard example was first introduced by Crawford and Sobel (1982) to illustrate their model of cheap talk. It has since been adapted to illustrate many different

types of games with strategic communication. Our model stays as close as possible to the standard model with hard information, as described, for instance, in Milgrom and Roberts [1986].

The key difference is the specification of the outcome function. In the standard example, the outcome function is given by $x(d) = d + \theta$, where θ is a single-dimensional state that is distributed on some interval. The outcome functions for different values of θ, therefore, never intersect. As we noted in the Introduction, this implies that if the receiver were to learn the outcome generated by a single decision, she would be able to infer the state and work out the outcome of any decision.

Our intention is to capture richer notions of expertise that cannot be inferred from a single piece of information. The assumption that the outcome function is the realized path of a stochastic process enables us to do just that. In our setting, learning the outcome generated by a single decision would allow the receiver to better predict the outcomes of other decisions. But he would never be able to do so perfectly. The outcome function, in other words, is only partially observable, while it is perfectly invertible in the standard example. This is the key difference between the models.

A second difference is that we assume a discrete number of decisions, which we then take to infinity, while the standard example simply assumes a continuum. We start with a discrete number of decisions to deal with a number of technical issues. The skeptical beliefs we specified above, for instance, rely on the existence of a "largest unrevealed decision," which may not exist with a continuum of decisions. Even though we state the model and our results for an arbitrary number of decisions, we focus throughout on the limit case that is closest to the standard example.

Finally, a difference that at first seems innocuous is that we assume $u_S(d) = -d$ rather than $u_S(d) = d$. This difference seems innocuous since, in the the standard example, it is irrelevant whether the sender's utility is increasing or decreasing in the decision. In ours, however, it is not. Recall that in the absence of any communication from the sender, the only decision the receiver is informed about is the smallest decision $d_0 = 0$. The assumption that the sender's utility is decreasing in the decision then implies that the sender wants to discourage the receiver from experimentation. It captures situations in which the decision that is best for the sender is safe for the receiver. In the car repair example, for instance, replacing the catalytic converter is expensive, and thus attractive for the mechanic, but also safe for you, since it will solve any problems you may have had with the existing converter. The challenge for the mechanic is to discourage you from choosing a more limited repair that will be less costly and may still solve your problem. We focus on settings in which the expert wants to discourage experimentation since we believe that they are both common and more interesting than the case in which the expert wants to encourage experimentation. We do, however, explore this alternative case as an extension.

3 Non-Strategic Advice

We start by exploring a simplified version of the game in which the sender is replaced by nature and nature reveals the outcome $x(d')$ of just one, arbitrarily chosen decision $d' \in \mathcal{D}$.

After observing $x(d')$, the receiver updates his beliefs about the outcomes of all other de-

cisions. We refer to these beliefs as "neutral," to distinguish them from the beliefs the receiver forms when the sender acts strategically. The receiver's neutral beliefs are characterized in our first lemma.

Lemma 1. *Suppose nature reveals the outcome $x(d')$ of one arbitrarily chosen decision $d' \in \mathcal{D}$. The receiver then forms the following "neutral" beliefs. For any decision $d \geq d'$, he believes that outcome $x(d)$ is normally distributed with mean and variance*

$$\mathrm{E}\left[x\left(d\right)|x(d')\right] = x(d') + \left(d - d'\right)\mu \text{ and } \mathrm{Var}\left[x\left(d\right)|x(d')\right] = \left(d - d'\right)\sigma^2.$$

For any decision $d \leq d'$, instead, he believes that outcome $x\left(d\right)$ is normally distributed with mean and variance

$$\mathrm{E}\left[x\left(d\right)|x(d')\right] = \frac{d}{d'}x\left(d'\right) \text{ and } \mathrm{Var}\left[x\left(d\right)|x(d')\right] = \frac{d\left(d' - d\right)}{d'}\sigma^2.$$

Lemma 1 shows that for any decision $d > d'$, the receiver expects outcome $x(d)$ to be larger, the larger the decision is. If $d < d'$, instead, beliefs form a Brownian bridge. The receiver then expects outcome $x(d)$ to be a convex combination of the two known outcomes $x(0)$ and $x(d')$. In either case, the receiver is more uncertain about the outcome of a decision, the further it is from the closest known actions d_0 and d'.

After updating his beliefs, the receiver makes the decision that maximizes his expected utility. It is convenient to write the receiver's expected utility as

$$\mathrm{E}\left[u_R(x(d))|x(d')\right] = -\left(\mathrm{E}\left[x(d)|x(d')\right] - b\right)^2 - \mathrm{Var}\left(x(d)|x(d')\right),$$

which highlights his two, potentially conflicting goals: getting an expected outcome that is close to b and avoiding the risk. In what follows, we first discuss the receiver's optimal choice informally and then summarize and prove the results in our first proposition.

The receiver's optimal decision depends critically on how close the revealed outcome $x(d')$ is to his preferred outcome b. If $x(d')$ is sufficiently close to b, the receiver will not experiment and simply make the revealed decision. In this case, the prospect of getting an expected outcome that is even closer to b is not worth the additional risk.

What constitutes being "sufficiently close" to b to warrant making the revealed decision d' depends on whether its outcome $x(d')$ is smaller or larger than b: if $x(d') \leq b$, the receiver will make decision d' if $x(d')$ is no further than

$$\alpha \equiv \frac{\sigma^2}{2\mu}$$

from b and if $x(d') \geq b$, he will make decision d' if $x(d')$ is no further than

$$\beta(d') \equiv \frac{1}{2}\left(\sqrt{b^2 + 2\sigma^2 d'} - b\right)$$

away from b.

The difference between α and $\beta(d')$ reflects the difference in the relevant beliefs. If $x(d') \geq b$, the receiver's best alternative to making the revealed decision is to experiment with a smaller decision and if $x(d') \leq b$, his best alternative is to experiment with a larger

decision. The receiver's beliefs in these two cases are different, as shown in Lemma 1. One difference is that the effect of a change in d on the mean and variance of its outcome $x(d)$ is independent of d' if $d \geq d'$ but depends on d' otherwise. This is why $\beta(d')$ depend on d' while α does not.

If the outcome of the revealed decision is larger than $b + \beta(d')$, it is no longer optimal for the receiver to make the revealed decision. He then finds it optimal to experiment by making the decision $d \leq d'$ whose outcome he expects to be closest to $b + \beta(d')$. The critical value $b + \beta(d')$ is therefore both, the threshold outcome above which the receiver engages in experimentation and the expected outcome he aims for if he does engage in experimentation.

To state the receiver's optimal decision more precisely, it is convenient to denote the difference between the expected outcome of decision d and $b + \beta(d')$ by

$$\epsilon_\beta(d) \equiv \mathrm{E}\left[x\left(d\right)|x(d')\right] - \left(b + \beta(d')\right).$$

If $x(d') > b + \beta(d')$, the receiver's optimal decision is then given by

$$d_\beta(d') = \arg\min_d \left|\epsilon_\beta(d)\right|.$$

If the number of decisions is finite, there may be two decisions that minimize $\epsilon_\beta(d)$, in which case we define $d_\beta(d')$ as the larger one. Moreover, it will generally be the case that $\epsilon_\beta(d_\beta(d')) > 0$, which makes the exposition slightly more tedious. Both these issues disappear in the limit in which the number of decisions goes to infinity, in which case there is a unique decision whose expected outcome is exactly equal to $b + \beta(d')$.

If the outcome of the revealed decision $x(d')$ is smaller than $b - \alpha$, but not too much so, the receiver also engages in experimentation. In line with his decision making when $x(d') > b + \beta(d')$, the receiver then makes the decision $d \geq d'$ whose outcome he expects to be closest to $b - \alpha$. To state his optimal decision more precisely, we denote the difference between $b - \alpha$ and the expected outcome of decision d by

$$\epsilon_\alpha(d) \equiv (b - \alpha) - \mathrm{E}\left[x\left(d\right)|x(d')\right].$$

If $x(d) < b - \alpha$, the receiver's optimal decision is then given by

$$d_\alpha(d') = \arg\min_d \left|\epsilon_\alpha(d)\right|.$$

If the number of decisions is finite, it may once again be the case that there are two decisions that minimize the objective, in which case we define $\epsilon_\alpha(d)$ to be the smaller one. Also, it will generally be the case that $\epsilon_\alpha(d_\alpha(d')) > 0$. Both issues once again disappear in the limit in which the number of decisions go to infinity, in which case there is a unique decisions $d \geq d'$ whose expected outcome is exactly equal to $b - \alpha$.

The key difference between $x(d')$ being above $b + \beta(d')$ and below $b - \alpha$ is that if $x(d')$ is sufficiently far below $b - \alpha$, the receiver prefers making the status quo decision to engaging in experimentation. This will be the case if the utility from the status quo is larger than the expected utility from experimenting with decision $d_\alpha(d')$, that is, if

$$(x_0 - b)^2 \leq \left(b - \mathrm{E}\left[x(d_\alpha(d'))|x(d')\right]\right)^2 + Var\left(x(d_\alpha(d'))|x(d')\right)$$

or, equivalently,

$$x(d') \leq -\frac{1}{2\alpha}\left[(\alpha - b)^2 + (x_0 - b)^2 - b^2 - \epsilon_\alpha(d_\alpha(d'))^2\right] \equiv \underline{x}(d', x_0).$$

As we noted above,$\epsilon_\alpha(d_\alpha(d')$ goes to zero as the number of decisions goes to infinity. In the limit, the expression for the threshold therefore further simplifies to $\underline{x}(d', 0) = -(b-\alpha)^2/(2\alpha)$, where we also used the fact that $x_0 = 0$.

The following proposition summarizes our discussion so far.

Proposition 2. *Suppose nature reveals the outcome $x(d')$ of one arbitrarily chosen decision $d' \in \mathcal{D}$. The receiver's optimal decision is then given by*

$$d^*(d') = \begin{cases} d_\beta(d') & if\ x(d') \geq b + \beta(d') \\ d' & if\ x(d') \in [b - \alpha, b + \beta(d')] \\ d_\alpha(d') & if\ x(d') \in [\underline{x}(d'), b - \alpha] \\ d_0 & if\ x\,(d') < \underline{x}(d'), \end{cases}$$

where α, $\beta(\cdot)$, $d_\alpha(\cdot)$, $d_\beta(\cdot)$, and $\underline{x}(\cdot)$ are defined as above.

A special case of this proposition is the one in which the revealed decision d' happens to be the status quo d_0. This case corresponds to the situation in which the receiver has to make a decision without receiving any advice. The proposition shows that the receiver then makes decision $d^*(d_0) = d_\alpha(d_0)$. To avoid the uninteresting case in which the receiver makes the sender's preferred decision even if he does not receive any advice, we assume that $d_\alpha(d_0) > 0$. This will be the case if the bias is sufficiently large relative to the degree of uncertainty and the number of decisions is sufficiently large, as stated explicitly below.

Assumption: *We assume that*

$$b - \frac{\sigma^2}{2\mu} \geq \frac{\mu}{2\sqrt{n}},$$

which ensures that the receiver makes a strictly positive decision in the absence of any advice.

The decision the receiver makes in the absence of any advice is a natural benchmark against which we will compare his decision when he receives advice from a strategic sender. Our goal in the next section is to understand when, why, and by how much such a sender is able to induce a decision that is smaller than the one the receiver makes in the absence of any advice.

4 Strategic Advice

We finished the last section by describing the receiver's decision making in the absence of any advice. We start this one by showing that if the sender is strategic, she will always provide at least some advice. To see this, suppose to the contrary that the sender's strategy is not to reveal any information. We know from the previous section that the receiver's best response is to make decision $d_\alpha(d_0) > 0$. Now suppose the outcome function happens

to be everywhere decreasing. The sender would then deviate by revealing the outcomes of all decisions. Since beliefs have to be consistent with hard information even off path, the receiver would then find it optimal to make the status quo decision, which is better for both parties than $d_\alpha(d_0)$. We therefore have the following proposition

Proposition 3. *In any equilibrium, the sender reveals the outcome of at least one decision.*

The sender's problem is therefore not whether to give advice but how to do so.

4.1 Conative Advice

To explore the sender's problem, we first consider equilibrium in which the sender reveals the outcome of a single decision which the receiver then makes. It is natural to interpret such equilibria as ones in which the sender "recommends" a decision and the receiver "rubberstamps" the recommendation. In line with this interpretation, we refer to such equilibria as "conative." The term was introduced by Jakobson [1960] who distinguishes between different functions of communication. The conative function corresponds roughly to a recommendation or instruction of what to do (Sobel [2013]).

Definition 4. A strategy profile is *conative* if, with probability one, the sender reveals the outcome of a single decision, and the receiver makes the revealed decision.

The key incentive constraint in a conative equilibrium is that the receiver has to find it optimal to rubberstamp the sender's recommendation. As in the case of non-strategic advice, the receiver should be willing to do so if the outcome of the recommended decision is sufficiently close to b. To get at what qualifies as "sufficiently close" when the sender is strategic, consider the following strategies and beliefs. The sender reveals the smallest decision that generates an outcome in $[x_c, b]$, where $x_c \in [b - \alpha, b]$. If no such decision exists, she reveals the outcomes of all decisions. The receiver's on path beliefs are consistent with Bayes' rule and his off path beliefs are skeptical. Finally, on path, the receiver makes the revealed decision that is best for him and, off path, he acts in accordance with his skeptical beliefs.

If the receiver's beliefs were neutral, it would follow from our discussion of non-strategic advice that these strategies and beliefs form an equilibrium. Since the sender is revealing information strategically, however, beliefs may no longer be neutral. To see this, suppose there exists a decision whose outcome lies in $[x_c, b]$. Suppose further that the sender sticks to her strategy and reveals the smallest decision d_c such that $x(d_c) \in [x_c, b]$. After observing $x(d_c)$, the receiver correctly infers that the outcome of any decision $d < d_c$ is less than x_c. His beliefs are therefore described by a truncated normal distribution, rather than a regular one, and are thus no longer neutral. We do not have closed form expressions for these beliefs. We know, however, that the receiver prefers d_c to any decision $d < d_c$. We know this since the expected outcome of any such decision has to be less than x_c and its variance has to be positive. The receiver therefore prefers rubberstamping $d_c$to making any decision $d < d_c$.

What about making a decision $d > d_c$? Since the sender's strategy is to reveal the *smallest* decision whose outcome lies in a particular range, revealing this decision contains no information about the outcomes of larger decisions. The receiver's beliefs about the

outcome of any decision $d > d_c$ is therefore still neutral. And since his beliefs are neutral, the receiver prefers rubberstamping d_c to any making any decision $d > d_c$, just as he would if advice were non-strategic.

The strategies and beliefs we stated above therefore form an equilibrium. For this equilibrium to be conative, however, it needs to be the case that there almost always exists a decision whose outcome lies in $[x_c, b]$. If the number of decisions is finite, there is always a strictly positive probability that no such decision exists. In this case, the equilibrium is therefore not conative, nor is there any other equilibrium that is.

Lemma 5. *For a given finite number of decisions, no equilibrium strategy profile is conative.*

In the limit in which the number of decisions goes to infinity, however, the outcome function becomes a Brownian motion with positive drift. For such an outcome function, there almost always exists a decision that generates any positive outcome, including those that lie in $[x_c, b]$. In the limit, the equilibrium therefore is conative and it implements outcome x_c almost always.

Proposition 6. *There exists a limit equilibrium that is conative and implements outcome x_c if and only if $x_c \in [b - \alpha, b]$.*

The proposition goes beyond summarizing our discussion so far. It shows that there exist conative equilibria that implement any outcome $x_c \in [b - \alpha, b]$. But it also shows that there do not exist any conative equilibria that implement outcomes that are not in $[b - \alpha, b]$. In particular, there are no conative equilibria that implement outcomes that are just above b. To see why, suppose we augment the strategies we described above so that the sender reveals the largest decision d_c such that $x(d_c)$ is in $(b, b+\beta(d_c)]$. Suppose also that such a d_c exists, as it almost always will in the limit, and that the sender sticks to her strategy and reveals $x(d_c)$. In line with our discussion above, the receiver would then find it optimal to rubberstamp the recommendation. The reason these strategies do not form an equilibrium nevertheless is that the sender will always deviate by revealing the outcomes of all decisions $d \leq d_c$. She will do so since it induces the receiver to make a decision that is strictly smaller than d_c and generates outcome b, which is better for both the sender and the receiver. There, therefore, cannot exist a conative equilibrium that implements an outcome that is even just a little bit above b.

An immediate implication of the proposition is that the conative equilibrium that is optimal for the receiver is the one that implements $x_c = b$. In this equilibrium, the sender reveals all relevant information, just as in the unravelling equilibrium of the standard example. As a result, the sender is not able to sway decisions in her favor. In fact, she is worse off than she would be if the receiver were forced to make a decision without being able to consult her.

The sender does better in equilibria in which x_c is smaller and she does best in the one in which $x_c = b - \alpha$. Even in this equilibrium, however, the average decision $\mathrm{E}[d_c]$ is just equal to the decision $d_\alpha(d_0)$ the receiver would make in the absence of any advice. The question then is whether the sender can do better by going beyond conative advice and revealing more information to the receiver.

4.2 Referential Advice

We now explore equilibria in which the sender complements her recommendation with supporting evidence that reveals the outcomes of some of the other decisions the receiver could make. The purpose of this additional information is to induce the receiver to rubberstamp the recommendation by casting doubt on the outcomes of other decisions. We refer to such an advice strategy as "referential." We borrow this term from Jakobson [1960] who introduced it to denote a function of communication that is different from the conative function of giving an instruction and that roughly corresponds to the notion of providing context for such an instruction.

Definition 7. A strategy profile is *referential* if, with probability one, the sender reveals the outcome of at least two decisions and the receiver makes the revealed decision that is best for him.

We will show that the sender is more persuasive, and thus performs better, when her advice is referential than conative. To this end, consider the following strategies and beliefs. The sender reveals the outcomes of decisions $d_1,\ d_2,\ ...,\ d'$, where d' is the smallest decision such that

$$x(d') \leq \underline{x}(d', x(d_r))$$

and where the threshold $\underline{x}(\cdot,\cdot)$ is defined in Section 3. If no such decision exists, $d' = d_n$, in which case the sender reveals the outcomes of all decisions. The receiver's on path beliefs about the outcomes of decisions $d > d'$ are neutral and his off path beliefs are skeptical. Finally, the receiver makes the best revealed decision $d_r \leq d'$.

To see that these strategies and beliefs form an equilibrium, note first that the sender's strategy does not permit any on schedule deviations. If the sender deviates, the receiver will know. And if the receiver knows the sender has deviated, he will act in accordance with his skeptical beliefs, which dissuades the sender from deviating in the first place.

If the sender does not deviate, the receiver observes the outcomes of all decisions $d_1, ..., d'$. If $d' = d_n$, the sender reveals the outcomes of all decisions, in which case it is immediate that the receiver's best response is to make decision d_r. If, instead, $d' < d_n$, the receiver forms beliefs about the outcomes of decisions $d > d'$. These beliefs are neutral since d' is the *smallest* decision whose outcome falls below a threshold that does not depend on the outcomes of any larger decisions. Learning the outcomes of decisions $d_1, ...,\ d'$ therefore does not provide the receiver with any additional information about the outcomes of decisions $d > d'$. Finally, since the receiver's beliefs are neutral, it follows from our discussion of non-strategic advice that his best alternative to making decision d_r is to experiment by making decision $d_\alpha(d')$. It also follows from that discussion that the receiver prefers d_r to experimenting with $d_\alpha(d')$ if and only if the above inequality holds.

The strategies and beliefs we stated above therefore form an equilibrium. For this equilibrium to be referential, however, it needs to be the case that the sender almost always reveals the outcomes of at least two decisions. If the number of decisions is finite, there is always a strictly positive probability that $x(d_1) \leq \underline{x}(d_1, x_0)$, in which case the sender reveals the outcome of only one decision and the receiver sticks to the status quo. If the number of decisions goes to infinity, however, the sender almost always reveals the outcomes of a continuum of decisions and the receiver makes the best revealed decision d_r.

The final step is to show that the sender performs better in this referential equilibrium than in any conative equilibrium. In the conative equilibrium that is best for the sender the receiver almost always makes decision d_c^*, the smallest decision that generates outcome $b - \alpha$. Consider now Figure [TO BE INSERTED], which plots the threshold $\underline{x}(d', x(d_r))$ as a function of $x(d_r)$. The key feature of the figure is that the threshold converges to the 45 degree line as $x(d_r)$ goes to $b - \alpha$. If $x(d_r)$ is close to $b - \alpha$, it therefore takes only a very small downturn in the outcome function for the receiver to prefer d_r to experimenting with $d_\alpha(d')$. Such a downturn almost always occurs, in which case $d_r < d_c^*$. And even if it doesn't, $d_r = d_c^*$. The receiver's decision in the referential equilibrium is therefore almost always strictly smaller than his decision in any conative equilibrium, and it is never strictly larger. Together, these observations provide an intuition for our next proposition.

Proposition 8. *There exists a limit equilibrium that is referential and in which, with probability one, the sender performs strictly better than she does in any conative equilibrium.*

A feature of the referential equilibrium we described above is that the sender reveals the outcomes for a continuum of decisions. This raises the question of why the sender has to reveal so much information and cannot simply reveal $x(d_r)$ and $x(d')$. The problem with revealing only these two pieces of information is that the sender could then deviate on schedule, for instance by revealing the outcomes of decision $d_r - \epsilon$ and an appropriately chosen $d'' > d'$. Revealing the outcomes of a continuum of decisions therefore makes the sender's strategy credible.

4.3 Comparative Statics

Having established that advice is more persuasive if it goes beyond making a recommendation, we now explore how much more persuasive it is. One might expect that an expert's ability to persuade depends on the complexity of issues she is an expert on. In exploring the relative persuasiveness of conative and referential advice, we therefore focus on how it varies with our measure of informational complexity σ.

We start by comparing the expected decision in the referential equilibrium in the previous section with that in the sender's preferred conative equilibrium. This comparison provides a lower bound for how much the sender benefits from giving referential rather than conative advice. And since the average decision in the sender's preferred equilibrium is the same as the decision the receiver would make in the absence of any advice, it also provides a lower bound for how much the sender can benefit from providing advice. Since we cannot compute the expected decision in the referential equilibrium, we rely on simulations for the comparison between the two equilibria. Specifically, we simulate the referential equilibrium 1000 times for a particular set of parameters and record the averages of the two key variables: the largest decision the sender reveals and the decision the receiver makes. We run this simulation for 20 different values of σ and report the results in Figure [TO BE INSERTED].

The figure shows that, in the referential equilibrium, the more complex the sender's information is, the less information she has to reveal to dissuade the receiver from making a risky decision. And the less information the sender has to reveal, the smaller is the decision the receiver ends up making. As one might expect, the sender is therefore more persuasive, the more complex her informational advantage is.

The same, however, is true in the sender's preferred conative equilibrium, in which case the average decision is given by $(b - \alpha)/\mu$, where $\alpha \equiv \sigma^2/(2\mu)$. This reflects the fact that, on average, the smallest decision the receiver is willing to rubberstamp is smaller, the more uncertainty there is. The difference between the two equilibria is that the average decision decreases at a constant rate in the sender's preferred conative equilibrium and at a decreasing rate in the referential equilibrium. As a result, the relative persuasiveness of referential advice is largest for intermediate degrees of complexity. For such intermediate values, the figure shows that the average decision in the referential equilibrium is less than half of what it is in the sender's preferred conative equilibrium. In contrast, the type of advice matters less if complexity is either very high or low. If it is very high, the receiver ends up making a small decision even if advice is conative. And if it is very low, he ends up making a large decision even if advice is referential.

The persuasiveness of referential advice raises the question how the receiver fares when he receives such advice. Since the receiver is a Bayesian, he has to be better off receiving advice than not. Figure y confirms this fact and shows that the receiver too benefits the most from referential advice for intermediate degrees of complexity. Since referential advice induces a smaller decision on average than the decision the receiver would make in the absence of any advice, it also induces a smaller outcome on average. For the receiver to benefit from referential advice, it therefore has to be the case that it reduces the variation in outcomes by enough to compensate for the getting a worse outcome on average. To illustrate this, Figure [TO BE INSERTED] plots the CDF of the receiver's utility in the referential equilibrium minus his utility in the absence of any advice (for the same 1000 runs we used for the previous figures). The figure shows that in more than 60% of the runs, the receiver is worse off receiving referential advice than not receiving any advice. The reason he is still better off on average, is that in a small number of cases, the receiver is much better off with referential advice than without any advice. The receiver therefore benefits from referential advice because it avoids rare disasters.

References

Vincent P Crawford and Joel Sobel. Strategic information transmission. *Econometrica*, 50 (6):1431–1451, 1982.

Robert Gibbons, Niko Matouschek, and John Roberts. Decisions in organizations. In *Handbook of Organizational Economics*, pages 373–431. Princeton University Press, 2013.

Jonathan Gruber and Maria Owings. Physician financial incentives and cesarean section delivery. *The RAND Journal of Economics*, 27(1):99–123, 1996.

Roman Jakobson. Linguistics and poetics. In Thomas Sebeok, editor, *Style in Language*, pages 350–377. MIT Press, 1960.

Emir Kamenica and Matthew Gentzkow. Bayesian persuasion. *American Economic Review*, 101(6):2590–2615, 2011.

Steven D Levitt and Chad Syverson. Market distortions when agents are better informed: The value of information in real estate transactions. *Review of Economics and Statistics*, 90(4):599–611, 2008.

Paul Milgrom and John Roberts. Relying on the information of interested parties. *RAND Journal of Economics*, 17(1):18–32, 1986.

Paul Milgrom and John Roberts. Economic theories of the firm: Past, present, and future. *Canadian Journal of Economics*, 21(3):444–458, 1988.

Joel Sobel. Giving and receiving advice. In Daron Acemoglu, Manuel Arellano, and Eddie Dekel, editors, *Advances in Economics and Econometrics: Tenth World Congress*, pages 373–431. Princeton University Press, 2013.

Max Weber. *Economy and Society*. University of California Press, 1922.