Biased Recommendations

Wonsuk Chung and Rick Harbaugh*

August 2012

Abstract

We integrate recent results from the cheap talk literature into a simple model of an expert who recommends one of two actions to a decision maker who might instead take no action. We then provide the first experimental tests of predictions from the literature that recommendations are "persuasive" in that they reduce the chance the decision maker takes no action, that decision makers "discount" a recommendation for the more incentivized action, that lack of "transparency" about expert incentives increases lying by both biased and unbiased experts, and that experts "pander" by recommending the action the decision maker already favors. Subject behavior is consistent with each prediction.

JEL Classification: D82, C92, M3.

Key Words: cheap talk, discrete choice, deception, persuasion, discounting, transparency, pandering, lying aversion

---

*Chung: Economics Department, Indiana University, woschung@indiana.edu; Harbaugh: Department of Business Economics and Public Policy, Kelley School of Business, Indiana University, riharbau@indiana.edu. For helpful comments we thank Alexei Alexandrov, Mike Baye, Tim Cason, Archishman Chakraborty, Bill Harbaugh, Tarun Jain, Wooyoung Lim, Dmitry Lubensky, Michael Raith, Eric Rasmusen, Joel Sobel, Yossi Spiegel, Jimmy Walker, and seminar and conference participants at Indiana University's Kelley School of Business, the University of Rochester's Simon School of Business, the Indian School of Business, the Western Economics Association meetings, and the Fifth Conference on the Economics of Advertising and Marketing at Qinghua University's School of Management and Economics.

1 Introduction

Experts often provide recommendations to decision makers – salespeople recommend products to customers, lobbyists recommend spending proposals to legislators, and consultants recommend projects to clients. Often the expert benefits from some of the choices but does not benefit at all if decision maker fails to act, e.g., if a customer does not buy anything. When expert and decision maker incentives diverge in this way, can a recommendation still persuade the decision maker to take an action? What if the expert is known to have a stronger monetary incentive to push one of the actions, e.g., receives a higher commission on one of the products, or is suspected to have such a bias? And what if the expert knows that the decision maker might already be leaning toward one choice?

Understanding these issues is important to decision makers and also to companies, institutions, and regulators that help structure the incentive and information environment in which experts provide advice. In recent years the incentives of mortgage brokers to recommend high cost loans, the incentives of credit rating agencies to overrate risky bonds, the incentives of stock analysts to recommend certain stocks, and the incentives of doctors to recommend expensive treatments have all come under scrutiny. To resolve such problems is it necessary to eliminate biased incentives? Or is it sufficient to make incentives more transparent by requiring disclosure of any conflicts of interest?[1]

To gain insight into these types of problems, recent papers have extended the cheap talk model of Crawford and Sobel (1982) to make it applicable to discrete choice environments where an expert helps a decision maker choose between multiple actions that may also benefit the expert (Chakraborty and Harbaugh, 2007, 2010; Inderst and Ottaviani, 2012; Che, Dessein and Kartik, forthcoming). Recommending an action is a form of "comparative cheap talk" that induces the decision maker to have a more favorable impression of the recommended action but also a less favorable impression of the other actions. For instance, a salesperson's recommendation to buy one of several products on display makes a customer more favorably inclined toward the product but also less likely to buy the other products. Because of this endogenous opportunity cost from recommending one action or another, a recommendation may be credible even without reputational or other constraints on lying.

In this paper we develop a simple model that is based on this literature and then use the

[1] Regulations can impose more equal incentives, e.g., requirements for "firewalls" that limit the incentive of stock analysts to push their firm's clients, and incentives may also be adjusted voluntarily to increase credibility, e.g., Best Buy promotes its "Non-commissioned sales professionals" whose "first priority is to help you make the right purchasing decision". Similarly, conflict of interest disclosure may be imposed as the SEC does for investment advisors, or voluntarily adopted as many medical journals have done for authors.

model to provide the first experimental tests of the literature's main findings. We assume that an expert has private information on the values of two actions to a decision maker and recommends one of them. The expert benefits to some extent if either action is taken, but the expert does not benefit if the decision maker takes neither action, e.g., if a customer does not buy any of a salesperson's products. Hence, like in the Crawford-Sobel model but in a very different discrete choice setting, the expert's and decision maker's preferences are neither completely aligned nor completely opposed. For simplicity we assume that one action is good and one action is bad so that the only uncertainty is regarding which action is the good one.[2] Despite its simplicity, the model captures several phenomena that are the focus of research on recommendations by biased experts.

First, recommendations are not only credible in that they reveal information in equilibrium, but for sufficient payoff symmetry they are also "persuasive" in that they benefit the expert by reducing the probability that the decision maker walks away without taking either action (Chakraborty and Harbaugh, 2010). Even though the expert is completely biased toward the expert taking an action rather than no action, persuasive communication is possible simply because there are two possible actions that the expert benefits from. A recommendation for one action raises the expected value of that action and at the same time lowers the expected value of the other action, but the expert still benefits since the higher expected value of one of the actions is now more likely to exceed the decision maker's outside option from taking neither action. For instance, a customer is more likely to make a purchase if a recommendation persuades him that at least one of two comparably priced products under consideration is of high quality.

Second, when the expert's incentives for each action differ, in equilibrium a recommendation for the more incentivized action must be "discounted" and be less influential (Chakraborty and Harbaugh, 2007, 2010; Chung 2011). Since the expert sometimes falsely claims that the more incentivized action is better, a recommendation for that action raises the updated expected value of that action less than a recommendation for the other action. Therefore the decision maker is more likely to ignore the recommendation and stick with the outside option when the more incentivized action is recommended. In equilibrium the expert faces a tradeoff where one recommendation generates a higher payoff if it is accepted but it is less likely to be accepted, while the other recommendation generates a lower payoff if it is accepted but it is more likely to be accepted.

[2] More realistically, the values of the two actions may vary over a true two-dimensional space so that they might be nearly equal or very different (Chakraborty and Harbaugh, 2007, 2010; Che, Dessein and Kartik, forthcoming). In this case only comparative cheap talk about which action is better is credible, so the simplifying assumption that one action is good and the other bad is not so restrictive in our context.

Third, when the prior distribution of values is asymmetric so that the decision maker is more impressed by a recommendation for one of the actions, the expert benefits by "pandering" to the decision maker and recommending that action even when the other action is better (Che, Dessein, and Kartik, forthcoming). Hence biased recommendations can result even when the expert's incentives for either action are the same. For instance, if it is known that one of the actions is for some reason preferred by the decision maker then the expert has a better chance of getting a favorable action from the decision maker by recommending that action. The decision maker anticipates such pandering and, just as in the asymmetric incentives case, discounts a recommendation for that action. Our modeling approach differs from that of Che, Dessein, and Kartik, but we find that the same pandering phenomenon extends to our model.

Finally, we examine the impact of "transparency" on communication in this setting.[3] When the decision maker is uncertain whether the expert's incentives are biased toward one action or not, a biased expert has a stronger incentive to lie since the decision maker puts some weight on the possibility that the recommendation came from an unbiased expert with an equal incentive to push either action. More interestingly, since the action favored by the biased expert is discounted, an unbiased expert also has an incentive to lie by recommending the opposite action from that favored by the biased expert. For instance, if a salesperson is suspected to benefit more from pushing one product than another product but in fact has equal incentives, then the salesperson benefits from pushing the other product. Or if an unbiased newspaper is perceived to have a possible liberal bias, then recommending to readers a more conservative rather than more liberal policy increases the odds of a policy recommendation being accepted.

These results have analogs in recommendation games where the decision maker faces a binary choice and the expert has a reputation to protect, e.g., a seller makes a recommendation whether or not to buy a single product and dishonest recommendations affect future sales. Sobel (1985) shows how a recommendation for an action can be persuasive when an expert has a reputational incentive not to mislead the decision maker, and how such a recommendation is discounted by the decision maker based on the strength of the incentive to push the action relative to the reputational costs. Gentzkow and Shapiro (2006) consider the incentive to pander to the decision maker's prior beliefs when there is uncertainty about the expert's competence and going against the priors makes the decision maker doubt the reliability of the expert for future advice. Morris (2001) considers how an unbiased expert has a "political

[3] In an online Appendix, Chakraborty and Harbaugh (2010) show the robustness of comparative cheap talk to some lack of transparency regarding expert incentives, but do not investigate the properties of communication under such conditions.

correctness" concern to avoid making the same recommendation as a biased expert so as to maintain a reputation for not being biased. In the approach we follow with two actions and the option to not take either action, results similar to these classic results hold in a simple one period model.

Recommending the wrong action has a natural interpretation as a lie in this game. Based on long-established results from the experimental literature on communication games we expect subjects to be reluctant to lie,[4] and based on recent research we also expect heterogeneity in the strength of this aversion across subjects (Gibson, Tanner, and Warner, forthcoming). Therefore, to capture this behavior and to make the model suitable for experimental testing, we depart from a "pure" cheap talk approach and assume that experts have a lying cost drawn from a distribution with support that ranges from zero up to being so high as to preclude lying.[5] Such costs could reflect a true preference against lying or be a reduced form for reputational or other concerns. Inclusion of lying costs has the quantitative effect of reducing the amount of lying, and it also has the qualitative effect of ensuring the existence of a unique equilibrium by eliminating the babbling equilibrium and other equilibria such as messages meaning the opposite of their literal meaning.

In our experimental tests we cannot control for varying subject preferences against lying, so even though there is a unique equilibrium the exact lying rates and acceptance rates cannot be predicted beforehand. However, we find that the comparative static predictions of the model for the phenomena we focus on are the same for any distribution of lying costs, so we can test these predictions even without knowing the exact distribution of subject preferences against lying. These predictions are also the same as in the most intuitive equilibrium of the limiting case of pure cheap talk, so the predictions are consistent with the insights generated by the theoretical cheap talk literature.[6]

Our assumption that lying costs are heterogeneous implies that when there is a monetary incentive to lie experts with lower lying costs strictly prefer to lie while those with higher lying

[4]Such aversion is consistent with the empirical pattern of overcommunication in cheap talk games (Dickhaut et al., 1995; Cai and Wang, 2006; Sánchez-Pagés and Vorsatz, 2007) and also with more direct evidence on pupil-dilation in such games (Wang, Spezio, and Camerer, 2010).

[5]Since messages have a direct affect on payoffs as in a (costly) signaling game, the model is a form of "costly talk" (Kartik, Ottaviani, and Squintani, 2007) or, as lying costs become arbitrarily small, a form of "almost cheap talk" (Kartik, 2009). Nevertheless, when the environment is not too asymmetric communication is still driven by the endogenous opportunity cost of messages as in the pure cheap talk game. If costs for all types preclude lying the game is a persuasion game (Milgrom, 1981).

[6]In the Crawford-Sobel model the unique equilibrium of the game with lying costs converges to the most informative equilibrium of the pure cheap talk game as such costs go to zero (Kartik, 2009). For sufficient symmetry, the same result obtains in our quite different environment.

costs strictly prefer to tell the truth. The fraction of expert types who lie then depends on the incentive to lie as determined in equilibrium by the model parameters and the distribution of lying costs. With no lying costs or homogeneous lying costs the expert is indifferent in equilibrium and uses a mixed strategy of lying or not, so the added realism of heterogeneous lying costs simplifies the game by purifying the equilibrium strategies.

2 Literature Review

The experimental literature on strategic information transmission via cheap talk has focused on testing different implications of the original Crawford and Sobel (1982) model (e.g., Dickhaut, McCabe and Mukherji, 1995; Cai and Wang, 2006, Wang, Spezio and Camerer, 2010).[7] Recently this literature has been extended to testing results by Battaglini (2002) for cheap talk by competing experts in a multi-dimensional version of the Crawford-Sobel model (Lai, Lim, and Wang, 2011; Vespa and Wilson, 2012). Our analysis is also motivated by results on multi-dimensional cheap talk, but we consider cheap talk by a single expert in a discrete choice situation. Moreover, we assume a simplified two-state environment that captures some but not all of the phenomena that arise in full multi-dimensional models.[8]

The simplicity of our biased recommendations game makes it particularly suitable for experimental testing but we are not aware of any other papers that test the game or closely related games. Blume, DeJong, Kim, and Sprinkle (1998) test a cheap talk game where in one state action A is better for both sides and in the other state action B is better for both sides. Since the players have the same preferred rankings for the actions, perfect communication is possible if the players can successfully coordinate on the equilibrium meaning of messages. In contrast, Sánchez-Pagés and Vorsatz (2007) test a cheap talk game where in one state action A is better for the expert and action B is better for the decision maker, while in the other state the payoffs are reversed. Since the preferred rankings are always opposed, the only equilibrium without lying costs is uninformative mixing or "babbling" by the expert. Following in the tradition of the Crawford-Sobel model but in a discrete choice setting, our model captures the intermediate case where preferences are partially aligned so that some communication is possible.[9]

[7] There is also a large literature on pre-play communication about strategic intentions in games with complete information, including mixed strategy games where the players have a strong incentive to lie (Crawford, 2003).

[8] With multiple dimensions comparative statements across dimensions are credible even when absolute statements within each dimension are not (Chakraborty and Harbaugh, 2010), but with only two states corresponding to one or the other action being better any statement is inherently comparative so we cannot test the relative persuasiveness of comparative versus absolute statements.

[9] Blume, DeJong, Kim, and Sprinkle (2001) test several games with partial common interests but their focus

A large behavioral economics literature following Gneezy (2005) tests the "deception game" in which an informed expert recommends one of two options to a decision maker where the payoffs are such that the expert and decision maker have the opposite preferred rankings of the options. Since the expert always has an incentive to lie, the question is how preferences for not lying might counteract this incentive. Our experiment is designed to test comparative static predictions from the cheap talk literature in a simple model so we do not analyze the range of factors that can affect lying aversion as identified in the deception game literature.[10] Methodologically we differ from this literature in that we give the subjects common knowledge information about the distribution of payoffs so we can analyze how communication changes as the expert and decision maker strategically respond to changes in the information and payoff structure.[11]

Cheap talk recommendations might be explicit but might also be implicit, e.g., one choice might be presented as a default choice, one product might be listed first on a website, or one product might be given more prominence on a storeshelf. The behavioral economics literature on "nudges" shows empirically how making one choice a default choice leads to greater adoption, e.g., employees are most likely to follow savings plans that are the default choice (Madrian and Shea, 2001). As suggested in the literature, one reason for the success of such nudges may be that decision makers infer that the default choice is often implicitly recommended by the expert. Our results add to this literature by showing how nudges can be persuasive even when the expert is self-interested and the decision maker is rightly suspicious of the expert's motives.

While most of the theory literature on cheap talk recommendations considers binary choices such as whether or not to buy a single product,[12] the early literature on credence goods examines recommendations to buy one of two versions of a product which is similar to our

is different and the games do not capture the same tradeoffs as in our game.

[10] Most notably Gneezy and others find that subject reluctance to lie increases with the (expected) impact of the lie on the decision maker. For simplicity we assume that all lies are equally costly, though the model can be adjusted to capture different ways of representing the varying impact of lies.

[11] The decision maker in the deception game is not told anything about the distribution of payoffs for either player and might (correctly) believe that the rankings are always opposed as in Sánchez-Pagés and Vorsatz (2007), that they are always aligned as in Blume et al. (1998), or that they can be either as in the intermediate case we explore. Sutter (2009) finds heterogeneity in beliefs that affects interpretation of the results.

[12] In addition to the reputation papers by Sobel (1985), Morris (2001), and Gentzkow and Shapiro (2006) discussed earlier, recent papers include Inderst and Ottaviani (2009), Jindapon and Oyarzun (2010), and Hodler, Loertscher, and Rohner (2010). Note that the Crawford-Sobel model can capture recommendations for ordered actions, e.g., recommendations to "buy", "hold", or "sell" a stock (Morgan and Stocken, 2003). It can also be a reduced form model of a recommendation to an informed decision maker regarding a binary decision (Chakraborty and Yilmaz, 2011).

approach (Darby and Karni, 1973). Most relatedly, Pitchik and Schotter (1987) consider an expert such as a doctor who recommends one of two treatments to a patient who does not know the severity of their problem – a cheap treatment solves a minor problem and an expensive treatment solves either a minor or major problem.[13] The model combines cheap talk and verifiable messages since the expert is allowed to falsely claim a more expensive treatment is necessary but, for liability reasons, is not allowed to falsely claim a cheaper treatment is sufficient. This assumption and other features of the model imply that the equilibrium probability of a lie depends only on the consumer's expected payoffs rather than on the expert's incentives. Our model builds on this analysis by showing how the probability and direction of a lie can be affected by both an incentive bias toward the more expensive treatment and also a pandering bias toward the cheaper treatment.

Recently several papers analyze the case of an intermediary such as a website or retailer who is paid by competing sellers to make their product more prominent (e.g., Armstrong, Vickers, and Zhou, 2009; Chen and He, 2011; Athey and Ellison, forthcoming). In equilibrium the seller with the better product will often win the bidding so that prominence becomes a signal of quality. Our approach differs in that the intermediary is an expert who is privately informed about the benefits of different choices, e.g., a retailer who knows about the quality of different products. We focus on how this information is communicated when the expert's recommendation is not directly up for sale but may be influenced by different incentives.

We take incentives as exogenous and test whether experts and decision makers behave consistently with comparative static predictions as we vary the incentives and information. In a related framework, Inderst and Ottaviani (2012) endogenize incentives by considering a consumer's choice between two products when a salesperson with information about which product is better for the consumer is offered commissions by the competing manufacturers of the products.[14] Our experimental results support the premise that in setting commissions it is necessary to consider how consumers discount recommendations based on their knowledge of the commissions.[15]

Our theoretical and experimental results on the benefits of transparency contrast with previous results in the literature. Cain, Loewenstein, and Moore (2005) and others find ex-

[13] Dulleck, Kerschbamer, and Sutter (2011) experimentally test several aspects of this credence good problem in a different framework without cheap talk.

[14] They find that commissions are endogenously equal in equilibrium when production costs are equal. In a model closer to the current one, Chung (2011) finds that if there are enough credulous buyers then commissions differ as part of mixed strategies by manufacturers. In a search model Armstrong and Zhou (2011) find mixed strategies in competition for prominence through different commissions.

[15] They also assume endogenous prices so that prices are adjusted in anticipation of such discounting.

perimentally that subjects do not fully discount recommendations by biased experts when the biases are disclosed, and that disclosure can even degrade communication by allowing experts to feel more license to exaggerate. Li and Madarász (2008) find theoretically that disclosure of the expert's bias in the Crawford-Sobel model often reduces communication. Inderst and Ottaviani (2012) in their model find that the salesperson's commissions for each product are endogenously lower when they are disclosed to the consumer, but also find that disclosure can inefficiently induce commissions to favor the higher cost producer. In our model we find that disclosure reduces lying by biased experts and also by unbiased experts who only lie when there is some suspicion that they might be biased. Contrary to previous experimental results, we find that decision makers tend to be insufficiently skeptical of expert recommendations when incentives are uncertain rather than when they are transparent.

In the following section we develop our formal model and provide hypotheses based on its equilibrium properties. We then outline how we test the hypotheses experimentally, and finally we report on the experiment results.

3 The Biased Recommendations Game

An expert knows the respective values v_A and v_B of two actions, A and B, available to a decision maker. The decision maker chooses one of the two actions or an outside option C, the value v_C of which is known by the decision maker but not by the expert. With equal chance one of the actions has a positive value and the other action has 0 value, $\Pr[(v_A, v_B) = (a, 0)] = \Pr[(v_A, v_B) = (0, b)] = 1/2$ where $0 < a, b \leq 1$. The value v_C of the outside option C is distributed according to the distribution F which for simplicity we assume is uniform on $[0, 1]$. The decision maker receives the value of whichever action is chosen.[16]

The expert receives $\pi_A > 0$ or $\pi_B > 0$ if the decision maker chooses action A or B respectively, but receives $\pi_C = 0$ if the outside option C is chosen. If $\pi_A = \pi_B$ we say the expert has "unbiased" incentives and otherwise say the expert has "biased" incentives, though in either case the expert's and decision maker's preferences are not fully aligned since $\pi_A, \pi_B > \pi_C$. We assume that the expert's incentives are common knowledge until the Opaque Incentives subsection below where we relax this assumption to allow for uncertainty over whether the expert is biased or not.

After observing v_A and v_B the expert sends a message $m \in \{m_A, m_B\}$. Sending m_A when

[16] The model differs from the standard logit discrete choice model, which is also tractable but less easily implemented experimentally, in using a uniform distribution for noise in the outside option C and in not including noise terms for A and B.

$v_B > v_A$ or sending m_B when $v_A > v_B$ incurs a lying cost d with distribution G. When $d = 0$ with certainty the game is a pure cheap talk game which we will analyze first but our main interest is the game where d is uncertain and the private information of the expert. The decision maker observes the message m, learns the value v_C, and chooses the action A or B or the outside option C to maximize her expected payoffs given her beliefs. The payoff matrix is shown in Table 1.

		Choose A	Choose B	Choose C
$v_A > v_B$	Send m_A	π_A, a	$\pi_B, 0$	$0, v_C$
	Send m_B	$\pi_A - d, a$	$\pi_B - d, 0$	$0, v_C$
$v_B > v_A$	Send m_A	$\pi_A - d, 0$	$\pi_B - d, b$	$0, v_C$
	Send m_B	$\pi_A, 0$	π_B, b	$0, v_C$

$$\pi_A, \pi_B > 0,\ 0 < a, b \leq 1,\ v_C \sim F,\ \ d \sim G$$

Table 1: Expert and Decision Maker Payoffs

Our equilibrium concept is Perfect Bayesian Equilibrium so along the equilibrium path the decision maker's beliefs follow Bayes Rule based on the expert's communication strategy which maps the expert's information (the realization of (v_A, v_B) and also the realization of d when it is uncertain) to the message space. Given the expert's communication strategy let the probabilities of a "lie" or "false claim" be $\alpha = \Pr[m_A | v_B > v_A]$ and $\beta = \Pr[m_B | v_A > v_B]$. For presentational simplicity we focus our discussion on equilibria where these false claim probabilities α and β are such that the updated expected value of the recommended action exceeds that of the unrecommended action,

$$\begin{aligned} E[v_A|m_A] &> E[v_B|m_A] \\ E[v_B|m_B] &> E[v_A|m_B]. \end{aligned} \tag{1}$$

When this condition is satisfied the decision maker never takes the unrecommended action and takes the recommended action if and only if its expected value exceeds v_C. Hence from the perspective of the expert who does not know the realization of v_C, the probabilities that recommendations for A and B are accepted are respectively

$$\begin{aligned} P_A &= \Pr[v_C \leq E[v_A|m_A]] = F(E[v_A|m_A]) = E[v_A|m_A] = \Pr[v_A > v_B|m_A]a \\ P_B &= \Pr[v_C \leq E[v_B|m_B]] = F(E[v_B|m_B]) = E[v_B|m_B] = \Pr[v_B > v_A|m_B]b \end{aligned} \tag{2}$$

where we have used the uniform distribution assumption for F. When payoffs are sufficiently asymmetric there can be equilibria where (1) does not hold because in equilibrium the updated expected values of the recommended and unrecommended actions are the same for one of the messages. In this case, which we will analyze in the Appendix whenever it arises, the decision maker will choose the outside option if and only if v_C is sufficiently high and otherwise mix between the two actions.

As a reference point we first consider pure cheap talk in which there is no lying aversion, and then we analyze our main model that incorporates lying aversion by introducing a distribution of lying costs.

3.1 Pure Cheap Talk

When there is no aversion to lying, $d = 0$ with certainty, if one message offers a higher expected payoff the expert will always send it. But if the expert always sends the same message then there is no transmission of information. Therefore in any informative equilibrium it must be that both messages are sent with positive probability and that the expected payoff from each message is exactly the same so that the expert is indifferent between messages, i.e., is indifferent between lying or not. The expected payoffs depend on the acceptance probabilities for each message which are a function of the probabilities of lying, so in equilibrium the expert must mix between lying or not such that the expected payoffs from the two messages are equal.

To see this, consider a candidate equilibrium where condition (1) holds so that P_A and P_B from (2) are the acceptance probabilities for each message, in which case the equilibrium condition that the expert is indifferent between messages (and hence is indifferent between lying or not) is

$$\pi_A P_A = \pi_B P_B. \tag{3}$$

This candidate equilibrium exists if there are α and β in $[0, 1]$ such that (1) and (3) hold. For this section we will focus on such equilibria in which the expert lies in only one direction if at all, so there is the additional restriction that either $\alpha = 0$ or $\beta = 0$ or both. When we consider lying aversion we will show that any equilibrium with lying costs must satisfy this restriction.

Suppose that the expert might claim that A is better when it is not, $\alpha \geq 0$, but never claims B is better when it is not, $\beta = 0$. Then the probability that A is really better when it is claimed to be better is $\Pr[v_A > v_B | m_A] = 1/(1 + \alpha)$ and the probability that B is really better when it is claimed to be better is $\Pr[v_B > v_A | m_B] = 1$. So from (2) the acceptance probabilities are just

$$P_A = \frac{a}{1 + \alpha},\ P_B = b \tag{4}$$

and the equilibrium condition (3) is then

$$\alpha = \frac{\pi_A a}{\pi_B b} - 1. \tag{5}$$

If instead the expert never falsely claims that A is better but might falsely claim that B is better, $\alpha = 0$ and $\beta \geq 0$, then by similar logic the equilibrium condition (3) is $\beta = \pi_B b/\pi_A a - 1$. In the first case $\alpha \in [0,1]$ if and only if $\pi_A a/\pi_B b \in [1,2]$ and in the second case $\beta \in [0,1]$ if and only if $\pi_B b/\pi_A a \in [1,2]$. Therefore any lying is in the direction of the overall incentive and value asymmetry, and this asymmetry cannot be too large,

$$\frac{1}{2} \leq \frac{\pi_A a}{\pi_B b} \leq 2. \tag{6}$$

Our first proposition follows from this analysis and also from checking of condition (1) for a recommended action to be better than an unrecommended action in the candidate equilibrium. As shown in the proof in the Appendix, the condition holds if the additional symmetry condition

$$\frac{a}{a+b} \leq \frac{\pi_A a}{\pi_B b} \leq \frac{a+b}{b}. \tag{7}$$

also holds.[17]

Proposition 1 *For incentives and values sufficiently symmetric to satisfy (6) and (7), there exists a cheap talk equilibrium in which (i) the expert lies with some probability in the direction of the overall incentive and value asymmetry and never lies in the opposite direction, and (ii) the decision maker accepts the expert's recommendation if the value of the outside option is sufficiently low and otherwise chooses the outside option.*

Other equilibria can also exist. In a pure cheap talk game there will always be a babbling equilibrium in which the expert randomizes to reveal no information. And when the above equilibrium exists there will also be an inverted equilibrium in which messages have the opposite meanings, and other equilibria in which the expert mixes between messages to varying degrees. There can also be equilibria with mixing by the decision maker between the recommended action and the other action. However, as we show below, for sufficient symmetry the above type of equilibrium is the only equilibrium that survives when we allow for lying aversion.[18]

[17] For the pure cheap talk case notice that only the ratios π_A/π_B and a/b are relevant for the equilibrium condition and the symmetry conditions. The distinct values become important once we introduce lying aversion.

[18] Without lying aversion influential cheap talk can break down in this model if incentives and values are too asymmetric. With a true multi-dimensional state space an influential pure cheap talk equilibrium always exists if there is noise over all the choices as in the logit model (Chakraborty and Harbaugh, 2010), but the expert is actually hurt by communication if the asymmetries are too large (Chung, 2011). Levy and Razin (2007) show how cheap talk can sometimes break down in a multidimensional version of the Crawford-Sobel model.

3.2 Lying Aversion

For our main analysis we depart from the pure cheap talk approach above and assume that $d > 0$ with positive probability. The expert knows his own lying cost d but the decision maker only knows that d has a common knowledge distribution G with no mass points and with support on $[0, \overline{d}]$ for $\overline{d} \geq \max\{\pi_A, \pi_B\}$. The distribution G can be such that the game is arbitrarily close to the limiting pure cheap talk case in which $d = 0$ with certainty.

Since the type of the expert is now defined by his private information about the state (v_A, v_B) and his own lying cost d the probabilities α and β represent the overall probabilities of lies for the distribution of expert types. The continuous variation in lying costs means that at most one expert type can be indifferent in either state, so all other expert types follow a simple pure strategy of lying or not. Hence α and β are just the fraction of expert types in the respective states who have a pure strategy of lying.

Consider a candidate equilibrium in which α and β are such that condition (1) holds. Then the decision maker either follows the recommendation or chooses the outside option depending on the realization of v_C so the probabilities P_A and P_B are as given in (2). Let d_A be the marginal expert type who for $v_B > v_A$ receives the same payoff (inclusive of lying costs) from falsely claiming A is better as from truthfully saying B is better, $\pi_A P_A - d_A = \pi_B P_B$. Similarly let d_B be the marginal expert type who for $v_A > v_B$ is indifferent between lying or not so that $\pi_A P_A = \pi_B P_B - d_B$. Clearly both such indifferent types cannot simultaneously exist for $d_A, d_B > 0$ so it must be that in one case all expert types are strictly better off telling the truth or strictly better off from lying. The latter is impossible since some types find lying is too costly, so it must be that $\alpha = 0$ and/or $\beta = 0$.

Suppose that $\beta = 0$ so that experts always tell the truth when $v_A > v_B$ and the question is what fraction α of expert types will falsely claim A is better when $v_B > v_A$. Given that d_A is defined as the marginal type, all expert types with lower lying costs lie, so $\alpha = G(d_A)$. Since G is strictly increasing in α we can consider the implicit function $d_A = G^{-1}(\alpha)$ where $d_A(\alpha)$ is continuous and strictly increasing in α and $d_A(0) = 0$. Substituting $\beta = 0$ into (2), the equilibrium condition $\pi_A P_A - d_A = \pi_B P_B$ is

$$\pi_A \frac{a}{1+\alpha} - d_A(\alpha) = \pi_B b. \tag{8}$$

Note that the condition holds only if $\pi_A a \geq \pi_B b$ so any lying is in the direction of the overall incentive and value asymmetry. If $\pi_A a = \pi_B b$ the condition holds for $\alpha = 0$ so never lying is an equilibrium and as α increases the LHS strictly decreases so this is the only equilibrium. If $\pi_A a > \pi_B b$ the LHS is larger for $\alpha = 0$ and as α increases the LHS strictly decreases and becomes negative so there exists a unique $\alpha > 0$ and corresponding indifferent type d_A such that the equilibrium condition holds. Similarly, if we suppose that $\alpha = 0$ so $\beta = G(d_B)$ and $d_B(\beta) = G^{-1}(\beta)$, the equilibrium condition $\pi_A P_A = \pi_B P_B - d_B$ is

$$\pi_A a = \pi_B \frac{b}{1+\beta} - d_B(\beta) \tag{9}$$

where $d_B(\beta)$ is strictly increasing in β and $d_B(0) = 0$. By the same logic as above, there exists a $\beta \geq 0$ (and corresponding d_B) such that the equilibrium condition holds if and only if $\pi_A a \leq \pi_B b$ and this β is unique.

The following proposition establishes the uniqueness of the equilibrium which involves extending the analysis to include mixing by the decision maker between actions A and B. A mixed strategy equilibrium has the same comparative static predictions for the behavior we examine as a pure strategy equilibrium and it only exists when asymmetries are such that there is no pure strategy equilibrium because condition (1) cannot hold. As shown in the proof in the Appendix, the condition for a pure strategy equilibrium holds as long as payoffs are sufficiently symmetric or any incentive and value asymmetries are in the same direction, $(\pi_A - \pi_B)(a - b) \geq 0$. The condition is always satisfied in our experiment.

Proposition 2 *In the unique equilibrium with lying aversion (i) the expert lies in the direction of the overall incentive and value asymmetry if and only if the expert's lying costs are sufficiently low and never lies in the opposite direction, and (ii) for sufficiently symmetric incentives and values, or for any incentive and value asymmetries in the same direction, the decision maker accepts the expert's recommendation if the value of the outside option is sufficiently low and otherwise chooses the outside option.*

Lying aversion has the qualitative effect of selecting the simple type of cheap talk equilibrium analyzed in Proposition 1, and it also has the quantitative effect of reducing the amount of lying. As lying costs become smaller for all expert types, i.e., G shifts upwards, the marginal type d_A or d_B that satisfies the relevant equilibrium condition (8) or (9) becomes smaller. Therefore the equilibrium conditions become closer to condition (3) for a pure cheap talk equilibrium and the equilibrium amount of lying also becomes quantitatively closer to that of the corresponding equilibrium of the pure cheap talk game when it exists. By the same logic, the amount of lying is arbitrarily close to that of the pure cheap talk game if the mass of lying costs is sufficiently concentrated near zero.

4 Predictions

We now consider the main testable implications of the above model. We assume there is lying aversion so the equilibrium is unique as shown in Proposition 2, but the main comparative static predictions are the same for the pure cheap talk equilibrium in Proposition 1 and we will refer to the pure cheap talk case to help illustrate some of the results.

4.1 Symmetric Incentives and Values

Communication by the expert is "persuasive" if it induces the decision-maker to act in a way that increases the expert's payoff relative to no communication.[19] To see how sufficient symmetry ensures that communication is persuasive in our environment suppose $\pi_A = \pi_B$ and $a = b = 1$. Without a recommendation the expected value of either action is 1/2, so the decision maker will take one of these actions if the outside option is less. Given the uniform distribution of the value of C, the probability that the decision maker's outside option is less than 1/2 is also 1/2, so the decision maker chooses one of the expert's actions half the time and chooses the outside option half the time.

Although the expert always wants the decision maker to avoid the outside option, the expert does not benefit from misrepresenting which of the two actions is better. Hence as seen in Proposition 1 there is a cheap talk equilibrium in which a recommendation for an action has the equilibrium meaning that the recommended action is better. And as shown in Proposition 2 this equilibrium is unique under our assumption that there is at least some lying aversion. A credible recommendation raises the expected value of the recommended action to 1, so the decision maker always chooses it instead of the outside option, thereby doubling the expert's expected payoff.

With asymmetries the effects of communication are more complicated as we will investigate in detail below, and the payoff to the expert has to also include any lying costs if the expert lies in equilibrium. But the above result for complete symmetry suggests that communication should still help within a range of parameter values around the symmetric case. This is confirmed in the following hypothesis for the unique equilibrium with lying aversion. The proof is in the Appendix.[20]

Persuasiveness Hypothesis: *For sufficiently symmetric incentives and values, every expert type is strictly better off with communication than without communication.*

In the experiment described in Section 5, in the Symmetric Baseline treatment we set

[19] We follow the definition of persuasiveness in Chakraborty and Harbaugh (2010) in which all expert types benefit. A weaker definition would allow experts to benefit in expectation before learning their type.

[20] With a true multi-dimensional state space, pure cheap talk is persuasive if expert preferences over the decision maker's estimates of the state are quasiconvex (Chakraborty and Harbaugh, 2010). In a logit model Chung (2011) finds that quasiconvexity switches to quasiconcavity as the incentives and payoffs become sufficiently skewed, so the expert becomes worse off from any communication.

expert incentives equal, $\pi_A = \pi_B$, and decision maker values equal, $a = b$. The expert has no incentive to lie in this symmetric case so no lying costs are incurred and we can focus on the expected monetary payoff which is proportional to the probability of either action being taken. Therefore in the experiment we just test whether this probability is higher with communication than that predicted without communication.

4.2 Asymmetric Incentives

A common concern with expert communication is that the expert is biased toward some choice. In our model the expert always prefers action A or B relative to the outside option C, but as shown above this does not preclude persuasive communication. Of particular concern is that the expert might benefit more from either A or B, e.g., a financial analyst might benefit more if an investor chooses one financial product rather than another, or a salesperson might benefit more if a consumer chooses one product rather than another. Chakraborty and Harbaugh (2007, 2010) show that as expert payoffs from inducing a higher estimate on a dimension become larger, the decision maker becomes less influenced on that dimension by cheap talk, and in a logit discrete choice model Chung (2011) shows how this translates into a lower probability of a recommendation being accepted.

Such discounting by the decision maker faces the expert with a tradeoff. If the incentive for successfully recommending A is higher than for B, should the expert recommend A or instead should the expert recommend B since the decision maker is less suspicious of such a recommendation and more likely to accept it? This tradeoff favors A if lying is unlikely but can favor B if lying is very likely and the decision maker is correspondingly suspicious of an A recommendation, so in equilibrium we should expect a fraction of experts to lie. We are interested in how the equilibrium fraction of experts who lie changes as the incentives change.

Suppose that incentives and values favor A overall, $\pi_A a \geq \pi_B b$, so the expert only lies in favor of action A, $\beta = 0$ and $\alpha \geq 0$. Assuming condition (1) holds, then from (4) the recommendation acceptance probabilities are $P_A = a/(1 + \alpha)$ and $P_B = b$. So the question is how changes in incentives affect the probability α of a false claim that A is better which then affects P_A. For $\pi_A a \geq \pi_B b$ an increase in π_A exacerbates the asymmetry in favor of A, so we expect more lies and more discounting. For the pure cheap talk case this is seen directly from (5) where α is increasing in π_A, and under our assumption of lying aversion the effect is the same as seen from (8). Since the LHS of (8) is decreasing in α, as π_A increases α must rise, or equivalently the marginal expert type must have higher lying costs. So there is more lying and by (4) the recommendation is more discounted.

Just as stronger incentives for an action lead it to be discounted, they also make a recommendation for the other action (weakly) more convincing. For $\pi_A a \geq \pi_B b$ an increase in π_A has no impact on P_B since the expert never lies in favor of B anyway, so the effect is weak. But if π_B increases then the asymmetry in favor of A is reduced, so as seen directly from (5) and as implied by (8), this leads to a decrease in α so P_A strictly rises while P_B remains the same.

Combining these results, an increase in the incentive for an action is always weakly in the direction of lowering the probability that a recommendation for that action is accepted and raising the probability that a recommendation for the other action is accepted. And depending on whether an increase in the incentives makes overall incentives and values more symmetric or less symmetric, i.e., makes $\pi_A a$ and $\pi_B b$ closer or not, the effects on acceptance rates are strict or not.[21] As shown in the Appendix the same effects hold when the decision maker mixes between actions, so we have the following comparative static result for how behavior changes in the unique equilibrium of our game.

Discounting Hypothesis: *A higher expert incentive for an action lowers the probability that a recommendation for the action is accepted and raises the probability that a recommendation for the other action is accepted.*

In our experiment we compare acceptance rates for A in the Symmetric Baseline treatment where incentives are equal with those in the Asymmetric Incentives treatment where incentives are higher for A and lower for B. Hence we expect a strict decrease in the acceptance rate for A in the latter treatment. We also test the implication that the acceptance rates are lower for A than for B within the Asymmetric Incentives treatment. And since the acceptance rates predictions are based on predicted differences in lying rates as discussed above, we also test lying rates both across and within treatments.[22]

4.3 Asymmetric Values

Che, Dessein and Kartik (forthcoming) analyze how the effects of differing prior beliefs by the decision maker about the values of different actions can affect communication by the expert. They find that the expert often has an incentive to "pander" to the decision maker's prior beliefs by recommending an action that the decision maker is already leaning toward, even if that action is not in fact the best action. Pandering not only undermines the value

[21] If action B is favored, $\pi_A a < \pi_B b$, increases in π_A strictly raise P_B and weakly lower P_A, so any non-marginal changes in incentives that flip the favored action between A and B have strict effects.

[22] In this treatment, and in subsequent treatments, we have chosen the parameter values so that in the limit as lying costs go to zero the expert always lies. Since we expect non-zero lying costs based on previous experimental evidence, some experts are predicted to not lie.

of communication to the decision maker, it also hurts the expert since the decision maker anticipates pandering and is suspicious of it. Hence the inability of the expert to commit to revealing truthful information can leave everyone worse off.

To capture differences in the decision maker's prior beliefs about the values of the actions, we consider the case where the distribution is asymmetric with $a \neq b$ so that one of the actions has a higher expected value to the decision maker. For instance, a consumer has a preference for the design of one product over another as long as it is a good product that works. We assume that the expert also knows this preference. Che, Dessein, and Kartik consider pandering for the case where the expert and decision maker payoffs are perfectly aligned for the two actions, but their insight about pandering also holds for our case of a self-interested expert who may have an incentive bias.[23] Just as in their model, if a truthful recommendation would raise the expected value more for one action than another, the expert has a pandering bias to push that action more by falsely claiming it to be better.

As was the case with an incentive bias, the decision maker anticipates the pandering bias so the expert faces a tradeoff. Pushing A is more beneficial to the expert if the recommendation is believed, but the decision maker is more suspicious of such a recommendation compared to a recommendation for B. To see this how the equilibrium fraction of lies changes with parameter values, suppose condition (1) holds so there is a pure strategy equilibrium and again consider without loss of generality the case $\pi_A a \geq \pi_B b$ so that $\alpha \geq 0$ and $\beta = 0$. Then, as seen from (5) and (8), an increase in a strictly increases α and has no effect on β. Similarly an increase in b strictly decreases α and has no effect on β.

This establishes the following comparative static result for the case where the unique equilibrium is in pure strategies, and as shown in the Appendix the same result holds when the decision maker mixes between actions. Note that, as with the discounting hypothesis, whether the effects are strict or weak depends on whether the change in values makes overall incentives and values $\pi_A a$ and $\pi_B b$ more symmetric or less symmetric.

Pandering Hypothesis: *A higher decision maker value for an action raises the probability of a false claim that the action is better and lowers the probability of a false claim that the other action is better.*

In the experiment we test this prediction by comparing false claim rates in the Asymmetric Values treatment where b is lower with those in the Symmetric Baseline treatment. There

[23] Their assumption of perfectly aligned interests over the two actions highlights the power of pandering to disrupt communication but does not allow a pandering bias to coexist with an incentive bias as in our model. They also differ in allowing for a more general distribution of valuations for the actions and in assuming a fixed value of the outside option.

should be strictly more false claims for A in the Asymmetric Values treatment since b is lower. Note that while changes in the values a and b have the same effect on lying rates as changes in incentives π_A and π_B, the effects on acceptance rates differ. As b falls the probability $P_B = b$ falls directly and the probability $P_A = a/(1+\alpha)$ falls as false claims rise. We also test these predictions across the two treatments.[24]

Pandering and incentive biases can interact in our model to either strengthen or dampen the incentive to lie. An important case of such interaction is in selling environments when a salesperson's commission is increasing in price. If consumers are unsure of quality and one product is more expensive this gives the seller an incentive bias toward the more expensive good. But if the expected difference in quality is sufficiently small then the cheaper good offers a higher expected value net of price so the expert also has a countervailing pandering bias toward the cheaper good. This gives the salesperson more credibility and, depending on the exact parameters, the net effect may go against the incentive bias.

Pandering has largely unexplored implications for whether or not a decision maker benefits from revealing her payoff information to the expert. Suppose there is a chance that $a - b = c$ and an equal chance that $b - a = c$ for some constant $c \neq 0$ where the decision maker knows which case holds. For instance, a consumer might have a preference for the design of either product A or product B. If the expert does not know which case holds then for equal incentives the expert has no incentive to lie. But as seen in the above analysis the expert will sometimes lie for either case if the payoff difference is known. Hence the decision maker is better off from not disclosing her private information.[25]

4.4 Opaque Incentives

If an expert's incentives are biased, the results on discounting show that communication is degraded relative to the unbiased case, but communication is still possible. This assumes that the expert's incentives are "transparent" so the decision maker can adjust for them. To understand better the role of transparency we now extend the basic model to allow the decision maker to be uncertain about the expert's incentives so that the expert might be biased or not. Such uncertainty means that there are now two types of experts facing equilibrium tradeoffs between pushing A or B. In particular we assume that the expert is equally likely to be biased toward A with $\pi_A > \pi_B$ or to be unbiased with $\pi_A = \pi_B = \pi_u$ for some $\pi_u > 0$. The expert

[24] An increase in a when $\pi_A a \geq \pi_B b$ raises P_A and has no effect on P_B. We do not test this case experimentally.

[25] Whether this effect of pandering on disclosure incentives extends more generally beyond this simplifed model is an open question as Che, Dessein and Kartik note. In a discrete choice model where a seller communicates about the relative strengths of different attributes of a product, to avoid pandering a buyer benefits from withholding information about which attributes she values more (Chakraborty and Harbaugh, 2011).

knows the realized incentives but the decision maker only knows that the expert might be biased toward A or not. To focus on the incentives issue and simplify the analysis we set $a = b$.

We are interested in how the equilibrium tradeoffs for each type are different from the previous cases when the decision maker knew the expert was biased or not. For the biased expert who favors A, the decision maker is not sure they are biased so recommending A is less suspicious than before and we might expect more lying. For the unbiased expert who has equal incentives for A or B, the decision maker's uncertainty means that a recommendation for A is now suspicious, so we might expect some reluctance to recommend A. Of course, any such effects are anticipated by the decision maker so we have to examine the tradeoffs in equilibrium.

By the same arguments as for an expert with known incentives, neither a biased nor unbiased type of expert can be indifferent between telling the truth or not for both the $v_A > v_B$ and $v_B > v_A$ cases, so there is lying in at most one direction for each type. Extending previous notation, let the false claim probabilities be α_b and β_b for a biased type and α_u and β_u for an unbiased type. First suppose neither type ever lies in favor of B so $\beta_b = \beta_u = 0$. If either type lies in favor of A then the expected value of A given a recommendation is lower, so unbiased types will instead want to lie in favor of B, and there is a contradiction. If neither type ever lies in favor of A then the expected value of A or B given a recommendation is the same, so biased types will want to lie in favor of A, which is again a contradiction. Now suppose neither type ever lies in favor of A, $\alpha_b = \alpha_u = 0$. In this case there is again a contradiction if neither type or either type lies in favor of B. So it must be that either biased types sometimes lie in favor of A, $\alpha_b > 0$ and $\beta_b = 0$, and unbiased types in favor of B, $\alpha_u = 0$ and $\beta_u > 0$, or the opposite. Since biased types have more incentive to lie in favor of A, the former case must hold.

Now consider whether lying is more likely when incentives are opaque or when they are transparent. Unbiased experts do not lie when incentives are transparent and do lie when incentives are opaque, so the question is how lying by biased experts in favor of A is affected by being mixed in with unbiased experts. First, since biased experts sometimes lie in favor of A while unbiased experts never do, the expected value of A given a recommendation is higher if there are unbiased types mixed in. Second, since biased experts never lie in favor of B while unbiased experts sometimes lie in favor of B, the expected value of B given a recommendation is lower if there are unbiased types mixed in. Both of these effects increase the incentive to lie by biased types in favor of A, so they lie more when they are mixed in with unbiased types than when incentives are transparent.

As shown in the Appendix where these arguments are developed more formally, with opaque incentives there need not be a unique equilibrium with this extension to the model because the incentive to lie by each type is increasing in the probability that the other type lies. However any equilibrium has the above properties so we have the following hypothesis.

Transparency Hypothesis: *For symmetric decision maker values, if the expert might be biased toward an action or might be unbiased, then biased and unbiased experts lie in opposite directions and are more likely to lie than if incentives are transparent.*

In the experiment we test the prediction that biased and unbiased experts lie in opposite directions within the Opaque Incentives treatment where the decision maker does not know if the expert is biased or not.[26] We also test whether unbiased experts in this treatment are more likely to lie than in the Symmetric Baseline treatment where the decision makers knows the experts are unbiased, and whether biased experts are more likely to lie in this treatment than in the Asymmetric Incentives treatment where the decision makers knows the experts are biased. Regarding acceptance rates, since both types lie in the Opaque Incentives treatment the acceptance rates P_A and P_B are predicted to fall relative to the Symmetric Baseline case. Relative to the Asymmetric Incentives case, as discussed above the presence of unbiased types who sometime honestly indicate A is better raises P_A, and since these types sometimes falsely claim B is better their presence lowers P_B.

These results show how lack of transparency affects not just biased types but also unbiased types. A similar phenomenon arises in Morris's (2001) model of "political correctness" in which unbiased experts are reluctant to make the same recommendation as biased experts because it makes the decision maker more suspicious that they are biased and therefore reduces their future value as an expert. In our model this phenomenon happens in a single period because a recommendation for the same action as the biased type is less likely to be persuasive.

5 Experiment

5.1 Experimental Design

We conduct an experiment with four treatments, each one allowing for a test of one of the hypotheses generated by the model. The first Symmetric Baseline (SB) treatment sets $\pi_A = \pi_B$ and $a = b$ and tests the Persuasiveness Hypothesis that, since the expert has no reason to lie, the expert's recommendation is likely to be followed. The second Asymmetric Incentives (AI)

[26] Our case where there is an equal probability that the expert is biased or not and $a = b$ is most unfavorable to communication and each type always lies in the limit as lying costs go to zero, but we expect less lying due to non-zero lying costs.

treatment changes incentives to $\pi_A > \pi_B$ and tests the Discounting Hypothesis that, since the expert has an incentive to falsely claim A is better, the decision maker is less likely to follow a recommendation for A. The third Opaque Incentives (OI) treatment randomizes incentives between $\pi_A = \pi_B$ and $\pi_A > \pi_B$ and tests the Transparency Hypothesis that both biased an unbiased experts lie but in opposite directions. The final Asymmetric Values (AV) treatment sets $\pi_A = \pi_B$ and $a > b$ and tests the Pandering Hypothesis that the expert is more likely to recommend A.

We follow a within-subject design in which the same subjects are exposed to all four treatments in the above sequence. This sequence gives subjects the opportunity to learn by starting with the simplest symmetric treatment. To keep treatment effects from being confounded with learning effects and leading to false positives for the tested predictions, the sequence is designed so that experts in each treatment are predicted to behave in the opposite direction as in the previous treatment.[27] In the first Symmetric Baseline treatment experts have no incentive to lie. In the second Asymmetric Incentives treatment experts have an incentive to lie when $v_B > v_A$. In the third Opaque Incentives treatment unbiased experts have an incentive to lie when $v_A > v_B$. And in the fourth Asymmetric Values treatment experts have an incentive to lie when $v_B > v_A$. This sequencing ensures that observed results reflect a response to the new treatment and, to the extent that learning is incomplete during a treatment, it biases the results against our comparative static predictions.

We conduct four sessions with 20 subjects each drawn from undergraduate business classes at the Kelley School of Business at Indiana University. The experiments are conducted on computer terminals at the university's Interdisciplinary Experimental Lab using z-Tree software (Fischbacher, 2007). In each session the subjects are randomly assigned to be one of 10 "consultants" (experts) and 10 "clients" (decision makers).[28] Each treatment lasts 10 rounds and in each treatment every client is matched with every consultant once and only once. Subjects are never told with whom they are randomly matched in any round. At the end of the experiment one round from each treatment of ten rounds is randomly chosen as the round that the subjects are paid for.[29] Subjects learn which rounds were chosen and are paid privately in cash. No record of their actions mapped to their identity is maintained.

[27] For this reason we do not alternate between different sequences in different sessions. A between-subject design would eliminate any confounding of treatment and learning effects, but it would introduce the potential for treatment effects being confounded with other effects due to random subject differences across sessions.

[28] We chose this relatively general terminology to minimize the effect of prior beliefs about communication in more specific relationships such as salesperson-customer, lobbyist-politician, or financial advisor-investor.

[29] If only one round is randomly chosen then in theory players should act as if they are risk neutral and there should be no effect of cumulative (expected) earnings on behavior. Paying for a small number of randomly chosen rounds is a common method for limiting but not eliminating such earning effects.

	Choose A	Choose B	Choose C
Symmetric Baseline Treatment			
A Better	$8, 10$	$8, 0$	$0, [0, 10]$
B Better	$8, 0$	$8, 10$	$0, [0, 10]$
Asymmetric Incentives Treatment			
A Better	$10, 10$	$5, 0$	$0, [0, 10]$
B Better	$10, 0$	$5, 10$	$0, [0, 10]$
Opaque Incentives Treatment			
A Better	10 or $8, 10$	5 or $8, 0$	$0, [0, 10]$
B Better	10 or $8, 0$	5 or $8, 10$	$0, [0, 10]$
Asymmetric Values Treatment			
A Better	$8, 10$	$8, 0$	$0, [0, 10]$
B Better	$8, 0$	$8, 5$	$0, [0, 10]$

Table 2: Consultant (Expert) and Client (Decision Maker) Monetary Payoffs

The basic setup is the same in each treatment. The consultant has two projects, A and B, one of which is "good" for the client and one of which is "bad". The probability that a project is the good project is 1/2 and the computer randomly chooses which one is good independently each round. The client has an alternative project C with a value that is drawn by the computer from a uniform distribution independently each round. The consultant knows the realized values of projects A and B and the client knows the realized value of the alternative project C. After learning the realized values of projects A and B the consultant sends one of two messages via the computer to the client, either "I recommend Project A" or "I recommend Project B". The client learns the realized value of the alternative project C and sees the client's message. The client then chooses project A, B, or C and the decision and the outcome are displayed.

Incentives and values are set 10 times larger in US dollar terms than in the model as follows. Symmetric Baseline treatment: $\pi_A = \pi_B = \$8$ and $a = b = \$10$; Asymmetric Incentives treatment $\pi_A = \$10$, $\pi_B = \$5$, and $a = b = \$10$; Opaque Incentives treatment: $\pi_A = \pi_B = \$8$ for half of experts, $\pi_A = \$10$ and $\pi_B = \$5$ for half of experts, and $a = b = \$10$; Asymmetric Values treatment: $\pi_A = \pi_B = \$8$, $a = \$10$, and $b = \$5$. In every treatment the value v_C of the alternative project C is uniformly distributed between \$0 and \$10. These monetary payoffs are shown in Table 2. Note that the monetary payoffs are independent of the consultant's recommendations given each choice and the table does not include any costs to lying which are the private information of the subjects. Given the simplicity of the game the payoffs in each treatment are explained to subjects as described below rather than presented in matrix form as in Table 2. The average earnings in the 90-minute experiment are, including a $5 show-up payment, $27.02 for consultants and $32.78 for clients.

5.2 Instructions and Procedures

Subjects are led into the computer room and told to sit in front of any of 20 terminals. Instructions are provided on the terminals for the consultants and clients as the experiment progresses, and summaries are read aloud to ensure that information is common knowledge. A screenshot of the main introductory screen is in Appendix B and a summary is read aloud after subjects are seated.

After all subjects have clicked through the introductory instructions on their own terminals they see a screen which assigns them their roles for the rest of the experiment and provides instructions for the first Symmetric Baseline treatment (or "first series" as described to subjects). Screenshots of these instructions for both the consultant and client are included in Appendix B. The following summary is read aloud.

> **First Series.** The consultant has two projects - Project A and Project B. One is a good project worth $10 to the client and the other is a bad project worth $0 to the client. Each round the computer randomly assigns one project to be good and tells the consultant. The client does not know which project is good and which is bad.
>
> The client has his/her own project - Project C. Each round Project C is randomly assigned by the computer to be worth any value between $0.00 and $10.00. Any such value is equally likely. The computer tells the client how much Project C is worth, but the consultant does not know.
>
> The consultant will give a recommendation to the client via the computer. The recommendation will be "I recommend Project A" or "I recommend Project B".
>
> After getting the recommendation, the client will make a decision. The client earns the value of the project that is chosen. If the client chooses Project A or Project B the consultant will earn $8 in that round. However, if the client chooses his/her own Project C instead, the consultant will earn $0. One round from this series will be randomly chosen at the end as the round you are actually paid for.

After subjects click through the instructions on their screen, in each round consultants see the values of A and B and make their recommendations, and clients see the recommendations and the values of V_C and make their choices. Screenshots for these recommendations and choices are also included in Appendix B. The realized values and choices are then revealed to each consultant-client pair at the end of each round. Subjects only see the realized values and choices for their own pairing in that round. After the 10 rounds are completed subjects see a summary of their actions and payoffs in each round. They then see a screen with instructions for the next Asymmetric Incentives treatment ("second series") and the following summary is read aloud.

> **Second Series:** Everything is the same as the first series, except if the client chooses Project A the consultant will earn \$10 and if the client chooses Project B the consultant will earn \$5.

After the Asymmetric Incentives treatment is completed in the same manner as the Symmetric Baseline treatment, the Opaque Incentives treatment ("third series") begins. This is the most complicated treatment but its description is facilitated by its being a mix of the first two treatments. Each consultant is assigned the incentives of the first or the second treatment with equal chance. The assignment is block random and is for the duration of the treatment. The consultants know their incentives, but the clients do not know whether the consultant in any given round is biased or not. The following summary of the treatment is read aloud while subjects see a screen with the detailed description.

> **Third Series:** Everything is the same as before, except the Computer will randomly assign the consultant's payoff scheme. Half of the consultants will earn \$8 if either Project A or B is chosen, and half of consultants will earn \$10 if project A is chosen but only \$5 if project B is chosen. The consultant knows his/her payoff scheme but the client does not know which payoff scheme the Computer assigned the consultant.

After the Opaque Incentives treatment is completed, the Asymmetric Values treatment ("fourth series") begins. The following summary of the treatment is read aloud while subjects see a screen with the detailed description.

> **Fourth Series:** Everything is the same as the first series, including the consultant's payoff scheme of earning \$8 if either Project A or B is chosen, except the value of Project B to the client if it is good is only \$5 instead of \$10. If project A is good its value to the client is still \$10. A bad project is still worth \$0 to the client.

After the Asymmetric Values treatment is completed, subjects see a screen summarizing their actions and payoffs across all 40 rounds. On this screen they are also told what rounds have been randomly chosen as the rounds they will be paid for. They enter into the computer a number 1-20 that has been randomly placed next to the terminal and they are called out in sequence by this number to be paid by the experimenter. They do not see the payments for any other subjects, they sign for their payment next to their number without a printed name, and the record of their signature is preserved separately by the university's bookkeeping department. Other than the signature, no personally identifying records are maintained.

5.3 Results and Analysis

There are 10 rounds for each treatment with 10 pairs of recommendations and decisions in each round, so over the four sessions we observe a total of 400 recommendations and 400 decisions for each treatment. As in most experiments, behavior varies during the course of a session as subjects learn about the game. Since we do not have any predictions about the exact learning process, we focus attention on behavior in the last 5 rounds of each treatment. Therefore we report on observed behavior for a total of 200 recommendations and 200 decisions for each treatment. These data points are not independent since players make multiple decisions and learn from interactions with other players. For our statistical analysis we will follow the most conservative approach of treating the frequency of a behavior in a treatment for each session as the unit of analysis so that there are only four independent data points for a behavior, one for each session. These frequencies are reported in Table 3.

For the Symmetric Baseline treatment the experts have no incentive to lie so from the Persuasiveness Hypothesis we expect that decision makers will, on average, follow expert predictions whereas they would only do so half the time without communication. Focusing on the last 5 rounds of the treatment, we find that experts lie about 11% of the time and decision makers follow the expert's recommendation about 79% of the time. Figure 1(a) shows these frequencies where the "Expert Lies" bars show the frequency that the expert recommends A given that B is better (α^{SB}) and the frequency that the expert recommends B given that A is better (β^{SB}). Figure 2(a) shows the frequency of acceptance for each recommendation as a function of the value of the outside option v_C. As expected, the probability of acceptance for A and B falls as the outside option increases.

	Var.	Ses. 1	Ses. 2	Ses. 3	Ses. 4	Overall
Symmetric Baseline Treatment						
Recommend A When B Better	α^{SB}	.14	.13	.11	.07	.11
Recommend B When A Better	β^{SB}	.14	.15	.04	.04	.10
A Recommendation Accepted	P_A^{SB}	.63	.77	.92	.96	.81
B Recommendation Accepted	P_B^{SB}	.39	.71	.88	1.00	.76
A or B Recommendation Accepted	P_{AB}^{SB}	.52	.74	.90	.98	.74
Asymmetric Incentives Treatment						
Recommend A When B Better	α^{AI}	.35	.31	.52	.53	.41
Recommend B When A Better	β^{AI}	.21	.24	.04	.06	.13
A Recommendation Accepted	P_A^{AI}	.36	.32	.53	.53	.45
B Recommendation Accepted	P_B^{AI}	.68	.80	.71	1.00	.77
Opaque Incentives Treatment						
Biased: Recommend A When B Better	α_b^{OI}	.45	.45	.58	.78	.56
Biased: Recommend B When A Better	β_b^{OI}	.36	.07	.00	.31	.19
Unbiased: Recommend A When B Better	α_u^{OI}	.33	.15	.00	.00	.13
Unbiased: Recommend B When A Better	β_u^{OI}	.38	.58	.58	.50	.51
A Recommendation Accepted	P_A^{OI}	.58	.28	.68	.54	.52
B Recommendation Accepted	P_B^{OI}	.54	.64	.80	.92	.72
Asymmetric Values Treatment						
Recommend A When B Better	α^{AV}	.46	.68	.63	.70	.62
Recommend B When A Better	β^{AV}	.00	.05	.00	.09	.03
A Recommendation Accepted	P_A^{AV}	.46	.40	.68	.75	.51
B Recommendation Accepted	P_B^{AV}	.62	.30	.44	.40	.45

Table 3: Summary Data: Lying and Acceptance Frequencies by Session

Regarding the statistical significance of these results, since we treat the data generating process as being at the session level for each treatment, we are interested in whether the rate of acceptance is significantly above 1/2 in the four treatments. From Table 3, we see that this frequency P_{AB}^{SB}, which is an average of P_A^{SB} and P_B^{SB}, is above 1/2 in each session.[30] For any distribution the probability that this would occur if the hypothesis were false is $(1/2)^4 = 1/16$,

[30]Since the frequencies of an A or B recommendation vary, the overall acceptance rate P_{AB}^{SB} is a weighted average of P_A^{SB} and P_B^{SB} and is 26/50, 37/50, 45/50, and 49/50. Lower acceptance rates in the first two sessions might reflect that subjects happened to be drawn disproportionately from lower level classes. Our within-subject design prevents such random differences in the subject pool from biasing the results, and our use of session-level data prevents correlated error terms for subjects within a session from biasing the standard errors.

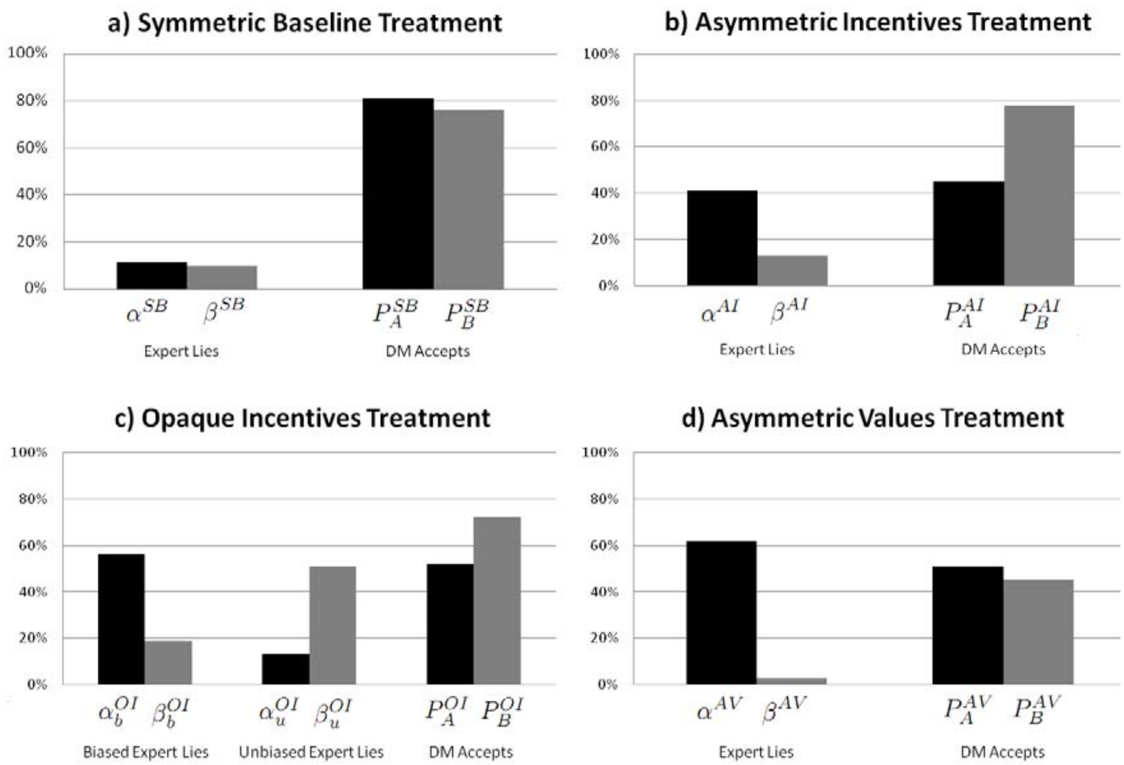

Figure 1: Overall Lying and Acceptance Frequencies

which is the p-value for the one-sided Wilcoxon signed rank test as seen in Table 4. If we assume that the frequencies are normally distributed then, as seen in Table 4, the one-sided t-test (which, in contrast with the subsequent tests, is a one sample test so it is not paired) indicates the difference is significant at the 5% confidence level.

In the Asymmetric Incentives treatment where action A is more incentivized, the Discounting Hypothesis is that decision makers are less likely to accept a recommendation for A since experts are more likely to falsely claim A is better when in fact B is better. Consistent with the predictions for false claim rates in this hypothesis, from Table 3 and as seen in Figure 1(b) we find that experts falsely claim A to be better 41% of the time and falsely claim B to be better 13% of the time. From Table 3 we see that false claims are higher for A than B in each session for the treatment, which for any distribution would occur if the hypothesis were false with probability $(1/2)^4 = 1/16$. This is the p-value for the Wilcoxon signed rank test as seen in Table 4. We also see that false claims for A are always higher than in the Symmetric Baseline treatment, which again implies a p-value of 1/16. Since this pattern occurs in almost

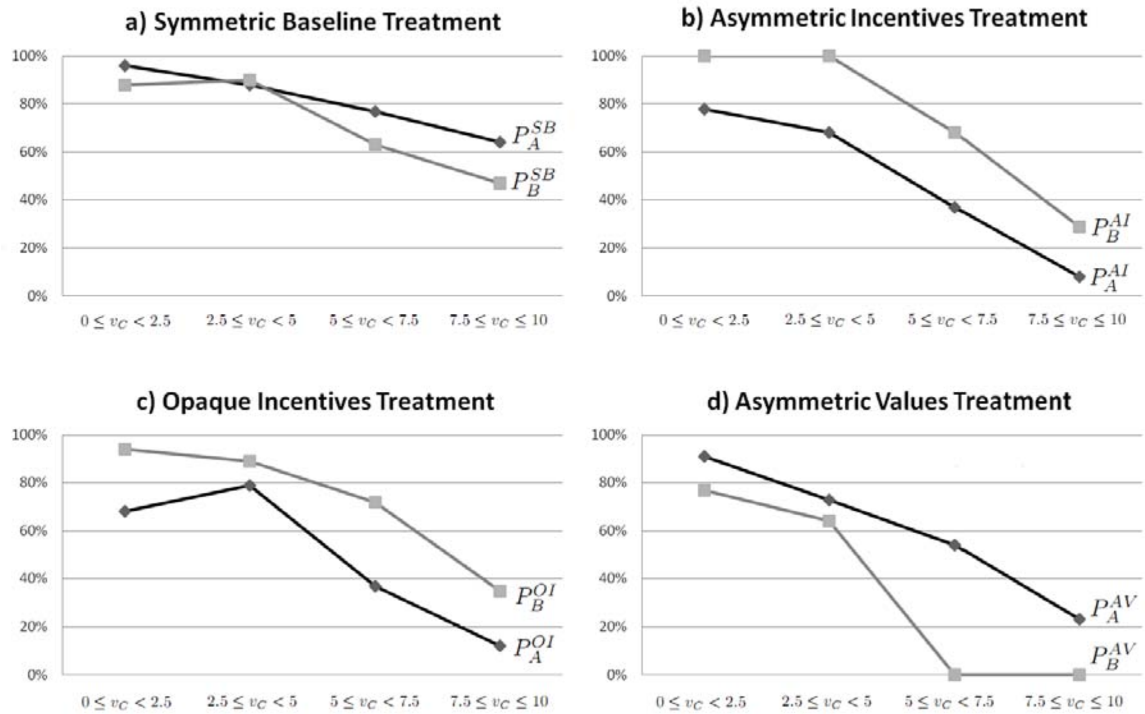

Figure 2: Overall Acceptance Frequencies by Value of Outside Option

every case, we will only mention the Wilcoxon signed rank test results when the pattern does not hold. Using a paired one-sided t-test both of these false claim differences are significant at the 5% or lower level.

Regarding acceptance rates, for the Asymmetric Incentives treatment decision makers accept the recommendation 45% of the time when A is recommended and 77% of the time when B is recommended. As seen in Figure 2(b) decision makers are very unlikely to accept a recommendation for A when the outside option is favorable. From Table 3 acceptance rates in the treatment are lower for A than for B in each session, and also acceptance rates are lower for A in the treatment than they are in the Symmetric Baseline treatment. As seen in Table 4, for the paired t-tests both of these differences in acceptance rates are significant at the 1% level.

	Hypothesis	Wilcoxon t-test	Paired t-test
Persuasiveness Hypothesis			
A and B Acceptance Rate SB Treatment	$P_{AB}^{SB} > 1/2$	.063	.032
Discounting Hypothesis			
A vs. B Lying Rates AI Treatment	$\alpha^{AI} > \beta^{AI}$	.063	.037
A Lying Rates AI vs. SB Treatments	$\alpha^{AI} < \alpha^{SB}$	.063	.010
A vs. B Acceptance Rates AI Treatment	$P_A^{AI} < P_B^{AI}$	.063	.007
A Acceptance Rates AI vs. SB Treatments	$P_A^{AI} < P_A^{SB}$	.063	.001
Transparency Hypothesis			
A vs. B Lying Rates Biased OI Treatment	$\alpha_b^{OI} > \beta_b^{OI}$	.063	.017
A vs. B Lying Rates Unbiased OI Treatment	$\alpha_u^{OI} < \beta_u^{OI}$	.063	.022
A Lying Rates Biased OI vs. SB Treatments	$\alpha_b^{OI} > \alpha^{SB}$	.063	.008
B Lying Rates Unbiased OI vs. SB Treatments	$\beta_u^{OI} > \beta^{SB}$	.063	.004
A Lying Rates Biased OI vs. AI Treatments	$\alpha_b^{OI} > \alpha^{AI}$	.063	.018
B Lying Rates Unbiased OI vs. AI Treatments	$\beta_u^{OI} > \beta^{AI}$	.063	.008
A Acceptance Rates OI vs. SB Treatments	$P_A^{OI} < P_A^{SB}$	.063	.027
A Acceptance Rates OI vs. AI Treatments	$P_A^{OI} > P_A^{AI}$	.188	.123
B Acceptance Rates OI vs. SB Treatments	$P_B^{OI} < P_B^{SB}$	.438	.372
B Acceptance Rates OI vs. AI Treatments	$P_B^{OI} < P_B^{AI}$	.188	.137
Pandering Hypothesis			
A vs. B Lying Rates AV Treatment	$\alpha^{AV} > \beta^{AV}$	.063	.000
A Lying Rates AV vs. SB Treatments	$\alpha^{AV} > \alpha^{SB}$	.063	.002
A Acceptance Rates AV vs. SB Treatments	$P_A^{AV} < P_A^{SB}$	.063	.005
B Acceptance Rates AV vs. SB Treatments	$P_B^{AV} < P_B^{SB}$	.125	.096

Table 4: p-values for One-Sided Hypothesis Tests, $n = 4$

For the Opaque Incentives treatment, from the Transparency Hypothesis we expect that a biased expert type is more likely to falsely claim A is better and, of particular interest, an unbiased expert type is more likely to falsely claim B is better. Consistent with the prediction, from the summary data in Table 3 and as seen in Figure 1(c) we find that biased experts falsely claim A is better 56% of the time and falsely claim B is better 19% of the time, while unbiased experts falsely claim A is better 13% of the time and falsely claim B is better 51% of the time. From the paired t-test results in Table 4 both of these false claim differences are significant at the 5% level. If we compare behavior with that in the Symmetric Baseline treatment, we find that in each session in the Opaque Incentives treatment biased experts are more likely to falsely claim A is better and unbiased experts are more likely to falsely claim B is better, and both differences are significant at the 1% or better level. Comparing behavior with that in the Asymmetric Incentives treatment, we find that in the Opaque Incentives treatment biased experts are more likely to falsely claim A is better, and unbiased types are more likely to falsely claim B is better, and that the differences are significant at the 5% and 1% levels respectively for the paired t-tests.

For acceptance rates in the Opaque Incentives treatment the results are more mixed. Theory predicts that decision makers should be more suspicious of an A recommendation than in the Symmetric Baseline treatment, and less suspicious of an A recommendation than in the Asymmetric Incentives treatment, but the differences are significant in only the former case. And theory predicts that decision makers should be more suspicious of a B recommendation than in either the Symmetric Baseline or Asymmetric Incentives treatment, but neither difference is statistically significant. Hence it seems that decision makers are correctly suspicious of an A recommendation, unbiased experts correctly anticipate this and falsely claim that B is better, but not all decision makers fully anticipate such lying by unbiased experts.[31]

Finally, for the Asymmetric Values treatment, from the Pandering Hypothesis we expect that experts are more likely to falsely claim that A is better than to falsely claim that B is better.[32] Consistent with the prediction, from the data in Table 3 and as seen in Figure 1(d) we find a difference of 62% vs. 3% in the summary data, which from the paired t-test results in Table 4 is significant at the 1% level. Comparing the frequency of false claims for A in this treatment with that in the Symmetric Baseline treatment, the difference of 62% vs. 11% is also significant at the 1% level for the paired t-tests. Regarding acceptance rates in the Asymmetric Values treatment, consistent with theory the acceptance rates for both A and B fall relative to the Symmetric Baseline treatment, but as seen in Table 4 only the fall for A is statistically significant.

The above results are for paired tests which are most appropriate for our within-subject design since they do not assume independence across treatments. For unpaired tests we find the same basic patterns as the paired tests. The unpaired t-tests all have lower p-values except for the comparison of A false claim rates in the OI and AI treatments. With this same exception, for the unpaired Wilcoxon (Mann-Whitney) rank order test the p-values for false

[31] Such behavior may reflect level-k thinking in which subjects do not fully consider the entire chain of strategic interactions (e.g., Stahl and Wilson, 1995; Nagel, 1995; Crawford, 2003).

[32] Recall that in our model the expert is self-interested rather than sharing the decision maker's preferences over A and B as in Che, Dessein and Kartik (forthcoming). We do not test their prediction that pandering can arise even with aligned preferences.

claim rates are all lower at $p = .014$ rather than $p = .063$. The differences in acceptance rates also all remain significant with equal or lower p-values, except that there is still no significant tendency in the Opaque Incentives treatment for decision makers to appropriately discount recommendations that are in the opposite direction of the bias held by biased experts.

5.4 Expert Lying Aversion and Decision Maker Best Responses

The model assumes the existence of some lying costs, but the comparative static predictions of the model do not depend on the exact distribution and the experiment is not designed to estimate the distribution. Nevertheless it is interesting to check what lying cost by the marginal expert – the one who is just indifferent between lying and not – would be consistent with best responses by the experts to the behavior of decision makers. For the Asymmetric Incentives treatment, substituting the observed frequencies $P_A^{AI} = .45$ and $P_B^{AI} = .78$ from Table 3 for the probabilities in (8), we get $10(.45) - d_A = 5(.78)$, or $d_A = .6$.[33] For the Opaque Incentives treatment we get higher numbers for both types of experts. For biased experts we get $10(.52) - d_A = 5(.72)$, or $d_A = 1.60$, and for unbiased experts (who have an incentive to recommend B instead of A) we get $8(.52) = 8(.70) - d_B$, or $d_B = 1.44$. For the Asymmetric Values treatment we get the lowest marginal lying cost, $8(.51) - d_A = 8(.45)$, or $d_A = .48$.

The finding of positive lying costs in these treatments is consistent with the experimental literature starting with Dickhaut et al. (1995) which finds that experts "overcommunicate" and reveal more information than would be expected in a pure cheap talk equilibrium. The pattern of lowest lying costs in the Asymmetric Values treatment and highest lying costs in the Opaque Incentives treatment also appears to be consistent with the deception game literature which finds that the consequences of lies affect the willingness to lie (Gneezy, 2005). Lies in the Asymmetric Values treatment are least damaging because the decision maker would only earn \$5 from a truthful recommendation for B. And lies in the Opaque Incentives treatment are most damaging because decision makers are more likely to believe a lie due to the mixture of biased and unbiased types. However, our model does not formally incorporate the idea that lying costs vary with their impact.[34] Note that lying costs in our experiment could reflect a true preference against lying, but may also be a reduced form for other considerations such as an altruistic concern for the payoff to the decision maker, which can also explain the pattern

[33] These and subsequent calculations assume that the expert's aversion to lying and benefit from lying are equally scaled by the chance that the lie is for a payoff-relevant round. Note that most experts appear to have a lying cost in a range that is sensitive to the benefits of lying – only 4 out of 40 experts never lie.

[34] Different ways to model the impact could be appropriate in different contexts, e.g., whether or not to adjust the perceived damage of a lie by the equilibrium probability that the decision maker acts on the lie.

of lies having differing costs depending on their effects (Hurkens and Kartik, 2009).[35]

While overcommunication is most common there is also some undercommunication which may just reflect noise in the behavior of experts or may reflect unmodeled aspects of preferences. In the Symmetric Baseline treatment the experts lie 11% of the time even though there is no incentive to lie, and in the other treatments there is always some lying in the opposite of the predicted direction. In our model such undercommunication would appear as negative lying costs for some experts, but just as overcommunication could reflect altruism, such undercommunication could also reflect competitive or "nasty" preferences (Abbink and Sadrieh, 2009). In the experiment an expert who is concerned with his payoff relative to that of the decision maker benefits from misleading them into the wrong choice.

Regarding decision maker behavior, acceptance rates on average are 8% below best response acceptance rates which is not surprising since the outside option C is certain while choosing A or B is risky. The treatments where we observe the greatest difference from the general pattern of best responses are suggestive of failure by decision makers to properly account for expert behavior. In the Asymmetric Incentives treatment the lying rates of 41% for A and 13% for B imply the expected value of A when recommended is $10(.87)/(.87 + .41) = 6.80$ and the expected value of B when recommended is $10(.59)/(.59 + .13) = 8.20$. The empirical acceptance rates are 45% rather 68% for A, and 78% rather than 82% for B. Decision makers discount expert recommendations for A so much that some decision makers do not appear to understand that communication is still possible even from a biased expert. In the Opaque Incentives treatment the combined lying rates for biased and unbiased experts of 33% for A and 35% for B imply the expected value of A when recommended is $10(.65)/(.65 + .33) = 6.63$ and the expected value of B when recommended is $10(.67)/(.67 + .35) = 6.57$. However the actual acceptance rates are 52% rather than 66% for A, and 72% rather than 66% for B. Insufficient discounting of B recommendations supports the inference from Table 4 that some decision makers do not appear to recognize that even unbiased experts can have an incentive to lie. Hence, contrary to previous results on transparency based on different models, we find that insufficient discounting of expert claims is a problem when expert incentives are opaque rather than transparent.

[35] The ordering of lying costs in our experiment might also reflect the differing complexity of the treatments since lying is cognitively difficult (Wang, Spezio, and Camerer, 2010).

6 Conclusion

This paper brings together several insights from the literature on cheap talk recommendations into one simple model. These insights were developed in different papers under varying assumptions, but they are sufficiently general that their main implications continue to hold in our simplified framework. The model shows how recommendations through cheap talk can be persuasive even when the expert always wants the decision maker to take an action rather than no action, how asymmetries in incentives and values can distort communication but need not preclude it, and how lack of transparency about the expert's incentives can lead even an expert with unbiased incentives to offer a biased recommendation. When experts are lying averse there is a unique equilibrium with testable comparative static predictions about expert and decision maker behavior that do not depend on the exact distribution of lying costs.

In the first experimental tests of these predictions from the literature, we find that for every hypothesis regarding expert behavior we can reject the null hypothesis of no change in the hypothesized direction. The false claim rates by experts change as predicted overall and in every session. Of particular interest is that when incentives are opaque we find that unbiased experts appear to recognize that they are more persuasive if they sometimes lie in order to avoid the recommendation favored by a biased expert. For the hypotheses regarding decision maker behavior, the acceptance rates also change as predicted and the changes are also statistically significant except that when incentives are opaque decision makers do not consistently discount the recommendation that is the opposite of that favored by the biased expert. Hence some decision makers do not fully anticipate how lack of transparency warps the incentives of even unbiased experts.

These results provide theoretical and empirical support for common measures that try to reduce the problem of biased recommendations. Consistent with theory we find that biased incentives make communication less reliable so attempts to eliminate or reduce biased incentives are likely to facilitate communication. And consistent with theory we find that lack of transparency further undermines communication so rules that require disclosure of expert incentives are also likely to help decision makers. However, there are two important caveats. First, as shown by Inderst and Ottaviani (2012), disclosure requirements can lead to endogenous changes in incentives that affect the equilibrium of the subsequent communication game, and there may also be other endogenous changes in the payoff or information structure in response to disclosure requirements or other measures. Second, as seen from the theoretical results by Che, Dessein and Kartik (forthcoming) that our experiment supports, an expert will often inefficiently pander to decision maker beliefs even when the problems of biased incentives and lack of transparency about incentives are both solved.

7 Appendix A – Proofs

Proof of Proposition 1: We need to verify condition (1). Given the probabilities α and β from the expert's communication strategy, the decision maker's updated estimates of the values of actions A and B given messages m_A and m_B are

$$\begin{aligned} E[v_A|m_A] &= \Pr[v_A > v_B|m_A]a = \frac{1-\beta}{1-\beta+\alpha}a \\ E[v_B|m_A] &= \Pr[v_B > v_A|m_A]b = \frac{\alpha}{1-\beta+\alpha}b \\ E[v_A|m_B] &= \Pr[v_A > v_B|m_B]a = \frac{\beta}{1-\alpha+\beta}a \\ E[v_B|m_B] &= \Pr[v_B > v_A|m_B]b = \frac{1-\alpha}{1-\alpha+\beta}b. \end{aligned} \tag{10}$$

If $\pi_A a \geq \pi_B b$ so that $\alpha \geq 0$ and $\beta = 0$ the second part of the condition $E[v_B|m_B] > E[v_A|m_B]$ always holds and the first part $E[v_A|m_A] > E[v_B|m_A]$ holds if $\alpha \leq a/b$. This always holds if $a \geq b$ and from (5) it also holds if $\pi_A/\pi_B \leq (a+b)/a$, or $\pi_A a/\pi_B b \leq (a+b)/b$. Similarly if $\pi_B b \geq \pi_A a$ so that $\alpha = 0$ and $\beta \geq 0$ the condition holds if $b \geq a$ or if $\pi_B b/\pi_A a \leq (a+b)/a$. Therefore, combining these cases, one sufficient condition is that any incentive or value asymmetries are in the same direction,

$$(\pi_A - \pi_B)(a-b) \geq 0, \tag{11}$$

and another is condition (7) that incentives and values are sufficiently symmetric. When (7) is not met but (11) is met, the equilibrium constraint (6) is violated so (11) is never relevant for our current case of pure cheap talk. Hence for the candidate equilibrium to exist it is sufficient for the incentives and values to be sufficiently symmetric to satisfy (6) and (7). ■

Proof of Proposition 2: When the decision maker follows a pure strategy in equilibrium the text shows that the equilibrium is unique. Now consider equilibria in which the decision maker is indifferent upon hearing m_A or m_B and follows a mixed strategy. Conditioning on taking any action at all, let ρ_A be the probability that the decision maker chooses A upon hearing m_A and let ρ_B be the probability that the decision maker chooses B upon hearing m_B. Since either $\beta = 0$ and/or $\alpha = 0$ by the same arguments as before, it cannot be that both $E[v_A|m_A] = E[v_B|m_A]$ and $E[v_B|m_B] = E[v_A|m_B]$ so the decision maker can only mix for at most one of the actions. Suppose the expert lies only if B is better, $\alpha \geq 0$ and $\beta = 0$, implying $E[v_B|m_B] > E[v_A|m_B]$. The indifference condition $E[v_A|m_A] = E[v_B|m_A]$ can only hold if $a < b$ and in particular requires $\alpha = a/b$. The equilibrium condition allowing for mixing is $(\rho_A\pi_A + (1-\rho_A)\pi_B)\,P_A - d_A = \pi_B P_B$, or substituting $\beta = 0$ and $\alpha = a/b$,

$$(\rho_A\pi_A + (1-\rho_A)\pi_B)\,\frac{a}{1+a/b} - G^{-1}(a/b) = \pi_B b. \tag{12}$$

Since $a/(1+a/b) < b$ a necessary and sufficient condition that the equality holds for a unique $\rho_A \in (0,1)$ is that $\pi_A a/(1+a/b) - G^{-1}(a/b) > \pi_B b$, which requires $\pi_A a > \pi_B b$. If the condition does not hold then (8) holds for some $\alpha \leq a/b$ and there is only a pure strategy equilibrium. Applying the same arguments to the case of $\alpha = 0$, we also find that a strictly mixed strategy equilibrium exists if and only if a pure strategy equilibrium does not. ■

Proof of Persuasiveness Hypothesis: Without communication, the decision maker will either take the outside option C or take the action A or B offering the highest expected payoff where $E[v_A] = a/2$ and $E[v_B] = b/2$. Since v_C is uniform on $[0,1]$ the probability the decision maker takes action A is $\Pr[E[v_A] \geq v_C] = a/2$ if $a > b$ and 0 if $a < b$. Similarly the probability the decision maker takes action B is $b/2$ if $b > a$ and 0 if $b < a$. Therefore, assuming the decision maker randomizes if $a = b$, the expert's expected payoff is $\pi_A a/2$ if $a > b$, $\pi_B b/2$ if $b > a$, and $(\pi_A + \pi_B)a/4$ if $a = b$.

With communication, suppose that condition (1) holds and consider the case where $\pi_A a \geq \pi_B b$ so that $\alpha \geq 0$ and $\beta = 0$ where α is given by the equilibrium condition (8). Expert types with a lying cost lower than d_A lie and receive a higher payoff than $\pi_B b$, while expert types with a higher lying cost tell the truth and receive $\pi_B b$, so the lowest payoff for any type is $\pi_B b$. If $a < b$ then the expected payoff without communication is $\pi_B b/2$, so communication always helps the expert. If instead $a > b$ then the expected payoff without communication is $\pi_A a/2$, so communication always benefits the expert if $\pi_B b > \pi_A a/2$. Considering also the case $\pi_A a \leq \pi_B b$, we find that communication strictly benefits every expert type if $1/2 < \pi_A a/\pi_B b < 2$, which is the symmetry condition (6). Since sufficient symmetry to meet (7) ensures that (1) holds, the result follows. ■

Proof of Discounting Hypothesis: Extending the analysis in the text to mixed strategy equilibria, consider WLOG the case $\pi_A a \geq \pi_B b$ so the mixed strategy equilibrium condition is given by (12). An increase in π_A raises the LHS so to maintain the condition ρ_A must decrease, which implies that the decision maker is less likely to accept a recommendation for A. The probability $P_B = b$ of acceptance of a B recommendation is unchanged so the effect holds weakly. An increase in π_B raises both sides but since $a/(1+a/b) < b$ it raises the RHS more so ρ_A must rise. Again there is no change in P_B, so the stated effect holds weakly. ■

Proof of Pandering Hypothesis: Extending the analysis in the text to mixed strategy equilibria, again consider WLOG the case $\pi_A a \geq \pi_B b$ so the mixed strategy equilibrium condition is given by (12). Since $\alpha = a/b$ and $\beta = 0$, as a increases the probability α of a false claim that A is better must strictly increase, and as b increases this probability must strictly fall, while the probability of a false claim that B is better is fixed, so the stated effect holds weakly. ■

Proof of Transparency Hypothesis: The expected values of each action conditional on a recommendation for that action are now

$$\begin{aligned} E[v_A|m_A] &= \Pr[v_A > v_B|m_A]a = \frac{\frac{1}{2}(1-\beta_b)+\frac{1}{2}(1-\beta_u)}{\frac{1}{2}(1-\beta_b)+\frac{1}{2}(1-\beta_u)+\frac{1}{2}\alpha_b+\frac{1}{2}\alpha_u}a \\ E[v_B|m_B] &= \Pr[v_B > v_A|m_B]b = \frac{\frac{1}{2}(1-\alpha_b)+\frac{1}{2}(1-\alpha_u)}{\frac{1}{2}(1-\alpha_b)+\frac{1}{2}(1-\alpha_u)+\frac{1}{2}\beta_b+\frac{1}{2}\beta_u}b \end{aligned} \tag{13}$$

and, using the assumption that $a = b$, the expected values of the unrecommended actions are $E[v_B|m_A] = (1-\Pr[v_A > v_B|m_A])a$ and $E[v_A|m_B] = (1-\Pr[v_B > v_A|m_B])a$. By the arguments in the text both types cannot be indifferent for any acceptance probabilities. Suppose that either $\beta_b = \beta_u = 0$ or $\alpha_b = \alpha_u = 0$. In this case condition (1) holds so the acceptance probabilities are $P_A = E[v_A|m_A]$ and $P_B = E[v_B|m_B]$ and there is always a contradiction following the arguments in the text. Therefore we are left with $\alpha_b > 0, \beta_b = 0, \alpha_u = 0, \beta_u > 0$ and the opposite case $\alpha_b = 0, \beta_b > 0, \alpha_u > 0, \beta_u = 0$, both of which imply condition (1) holds. The latter requires that the financial benefit is higher for biased types to recommend B and for unbiased types to recommend A, $\pi_A P_A < \pi_B P_B$ and $\pi_u P_A > \pi_u P_B$, which cannot hold for $\pi_A > \pi_B$, so only the former case is possible. The equilibrium indifference conditions for this case are

$$\pi_A P_A - d_A = \pi_B P_B \qquad \text{for } v_B > v_A \tag{14}$$

$$\pi_u P_A = \pi_u P_B - d_B \qquad \text{for } v_A > v_B \tag{15}$$

for the biased and unbiased types respectively where $d_A = G^{-1}(\alpha_b)$ and $d_B = G^{-1}(\beta_u)$. Regarding existence of equilibria, for any β_u starting at $\alpha_b = 0$ the LHS of (14) is larger than the RHS. From (13) note $\frac{d}{d\alpha_b}P_A, \frac{d}{d\alpha_b}P_B < 0$ and also $\frac{d}{d\beta_u}(P_A - P_B) = (\beta - 1)(\alpha - \beta)^2 + 4(\alpha - 1) < 0$, so as α_b increases the LHS of (14) falls faster until by the assumption on the support of G it eventually is smaller than the RHS. Therefore for every $\beta_u \in [0,1]$ there is a unique $\alpha_b \in (0,1)$ that solves (14). Similarly, note also that $\frac{d}{d\beta_u}P_B < \frac{d}{d\beta_u}P_A < 0$ so an increase in β_u pushes the RHS of (15) down faster than the LHS for any given α_b, so starting at $\alpha_b = 0$ in which case $\beta_u = 0$ solves (15), for every $\alpha_b \in [0,1]$ there is a unique $\beta_u \in [0,1)$ that solves (15). Considering the implicit functions $\alpha_b(\beta_u)$ and $\beta_u(\alpha_b)$ so defined, they must cross at least once in the interior so an equilibrium with lying by both types exists. By the same properties $\frac{d}{d\alpha_b}P_A < \frac{d}{d\alpha_b}P_B < 0$ and $\frac{d}{d\beta_u}P_B < \frac{d}{d\beta_u}P_A < 0$ these functions are both increasing so there may be more than one such fixed point. Despite the possibility of non-uniqueness we can still use the equilibrium properties $\alpha_b > 0, \beta_b = 0, \alpha_u = 0, \beta_b > 0$ to compare equilibria with the case of transparent preferences. Comparing (13) with (2), we find that P_A falls and P_B rises, so lying by biased experts increases as per the discussion in the text. ■

8 Appendix B – Sample Screenshots

Main introductory screen:

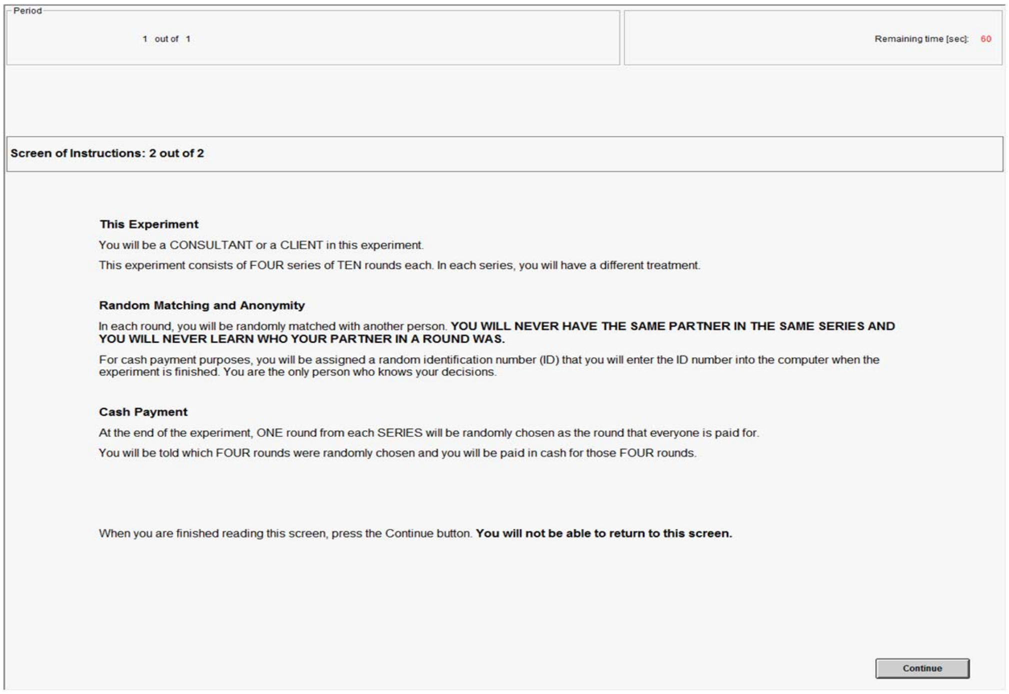

Period

1 out of 1

Remaining time [sec]: 60

Screen of Instructions: 2 out of 2

This Experiment

You will be a CONSULTANT or a CLIENT in this experiment.

This experiment consists of FOUR series of TEN rounds each. In each series, you will have a different treatment.

Random Matching and Anonymity

In each round, you will be randomly matched with another person. **YOU WILL NEVER HAVE THE SAME PARTNER IN THE SAME SERIES AND YOU WILL NEVER LEARN WHO YOUR PARTNER IN A ROUND WAS.**

For cash payment purposes, you will be assigned a random identification number (ID) that you will enter the ID number into the computer when the experiment is finished. You are the only person who knows your decisions.

Cash Payment

At the end of the experiment, ONE round from each SERIES will be randomly chosen as the round that everyone is paid for.

You will be told which FOUR rounds were randomly chosen and you will be paid in cash for those FOUR rounds.

When you are finished reading this screen, press the Continue button. **You will not be able to return to this screen.**

Continue

Consultant instructions for Symmetric Baseline treatment:

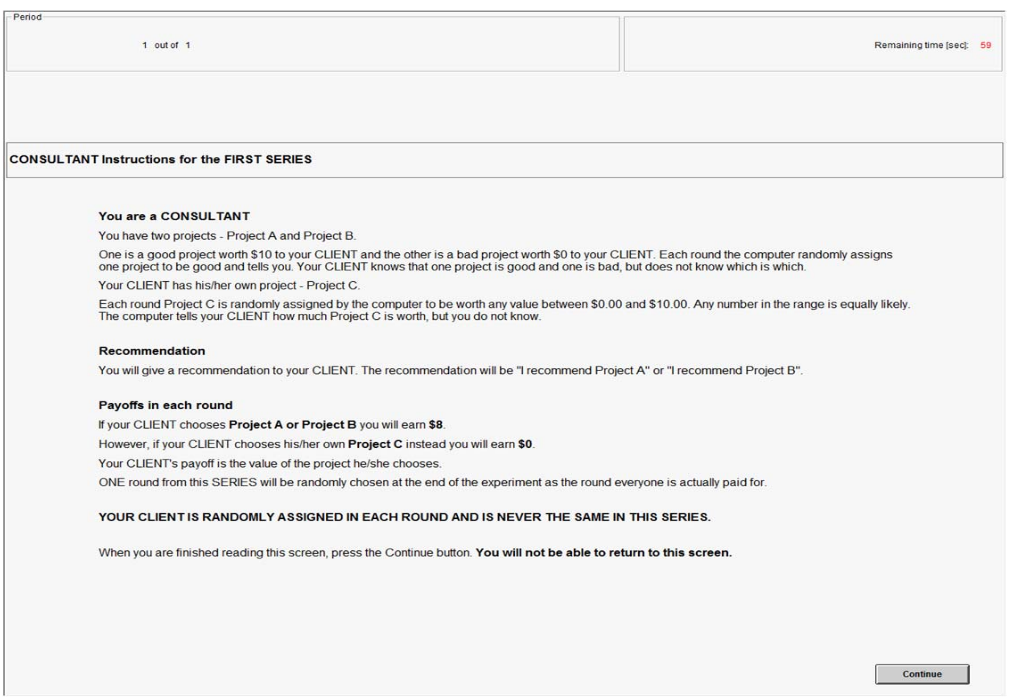

Period

1 out of 1

Remaining time [sec]: 59

CONSULTANT Instructions for the FIRST SERIES

You are a CONSULTANT

You have two projects - Project A and Project B.

One is a good project worth $10 to your CLIENT and the other is a bad project worth $0 to your CLIENT. Each round the computer randomly assigns one project to be good and tells you. Your CLIENT knows that one project is good and one is bad, but does not know which is which.

Your CLIENT has his/her own project - Project C.

Each round Project C is randomly assigned by the computer to be worth any value between $0.00 and $10.00. Any number in the range is equally likely. The computer tells your CLIENT how much Project C is worth, but you do not know.

Recommendation

You will give a recommendation to your CLIENT. The recommendation will be "I recommend Project A" or "I recommend Project B".

Payoffs in each round

If your CLIENT chooses **Project A or Project B** you will earn **$8**.

However, if your CLIENT chooses his/her own **Project C** instead you will earn **$0**.

Your CLIENT's payoff is the value of the project he/she chooses.

ONE round from this SERIES will be randomly chosen at the end of the experiment as the round everyone is actually paid for.

YOUR CLIENT IS RANDOMLY ASSIGNED IN EACH ROUND AND IS NEVER THE SAME IN THIS SERIES.

When you are finished reading this screen, press the Continue button. **You will not be able to return to this screen.**

Continue

Consultant decision screen for Symmetric Baseline treatment:

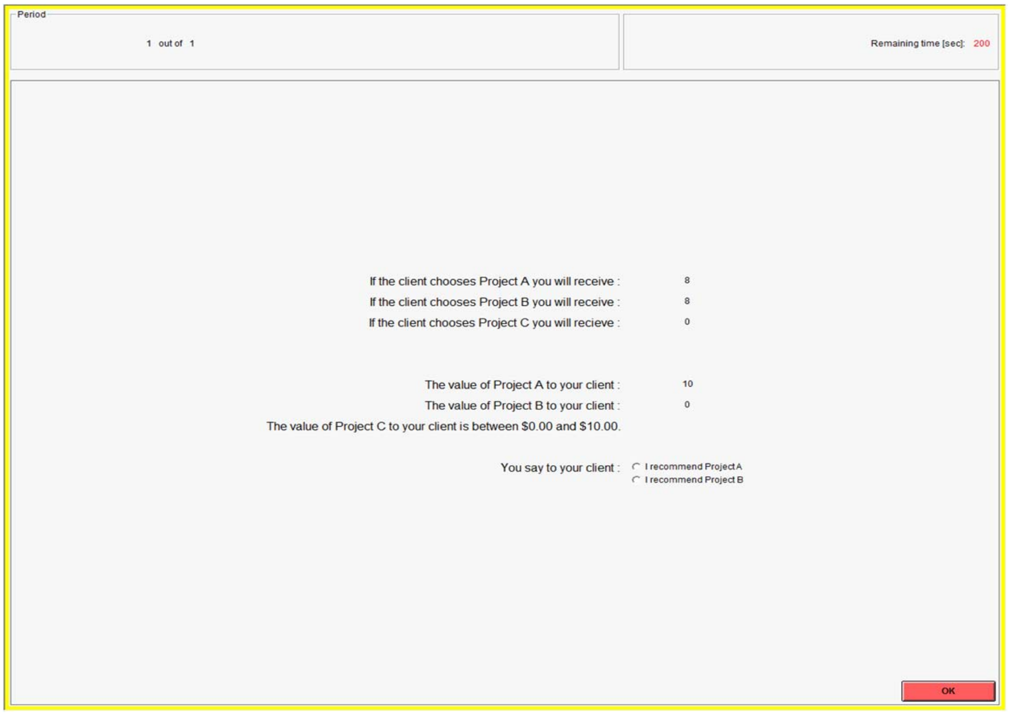

Client instructions for Symmetric Baseline treatment:

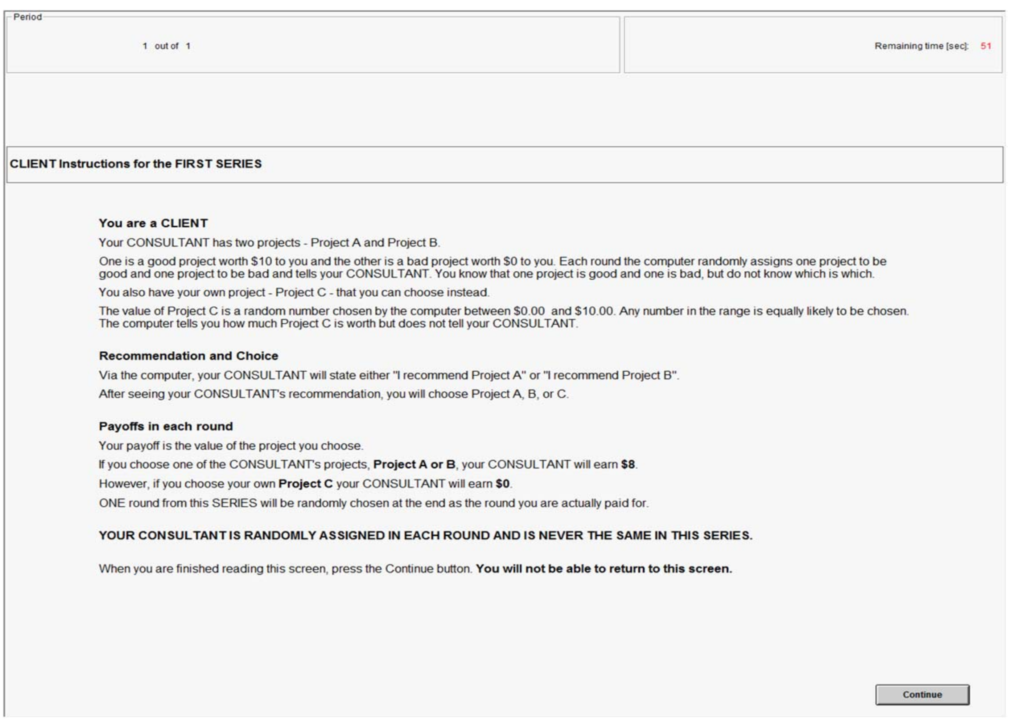

Client decision screen for Symmetric Baseline treatment:

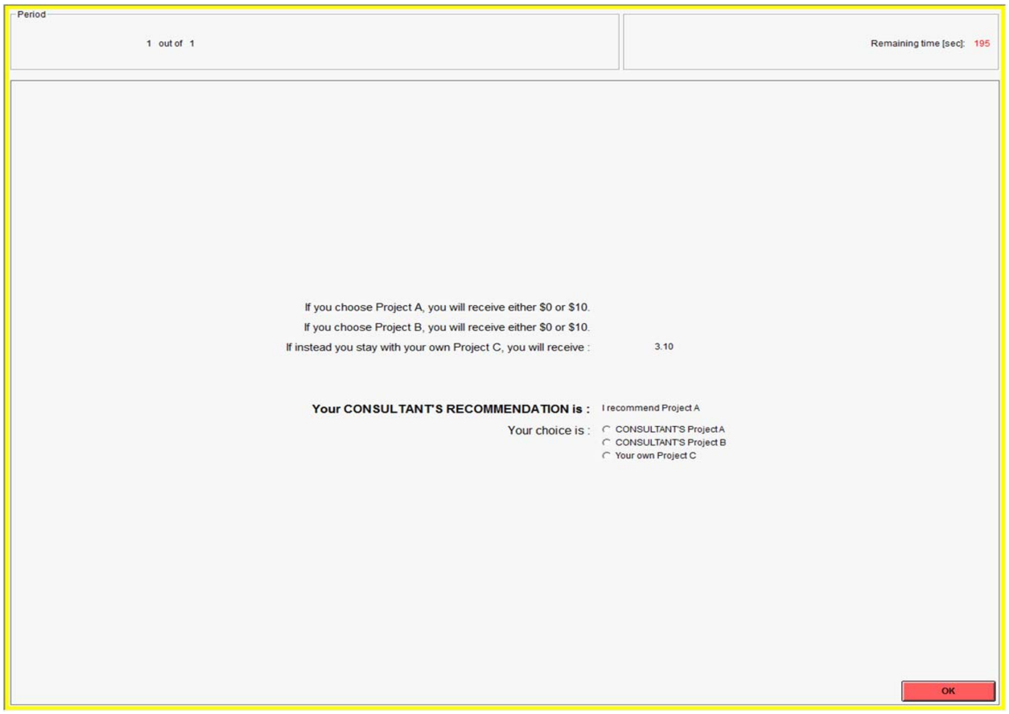

Client payoff screen for Symmetric Baseline treatment:

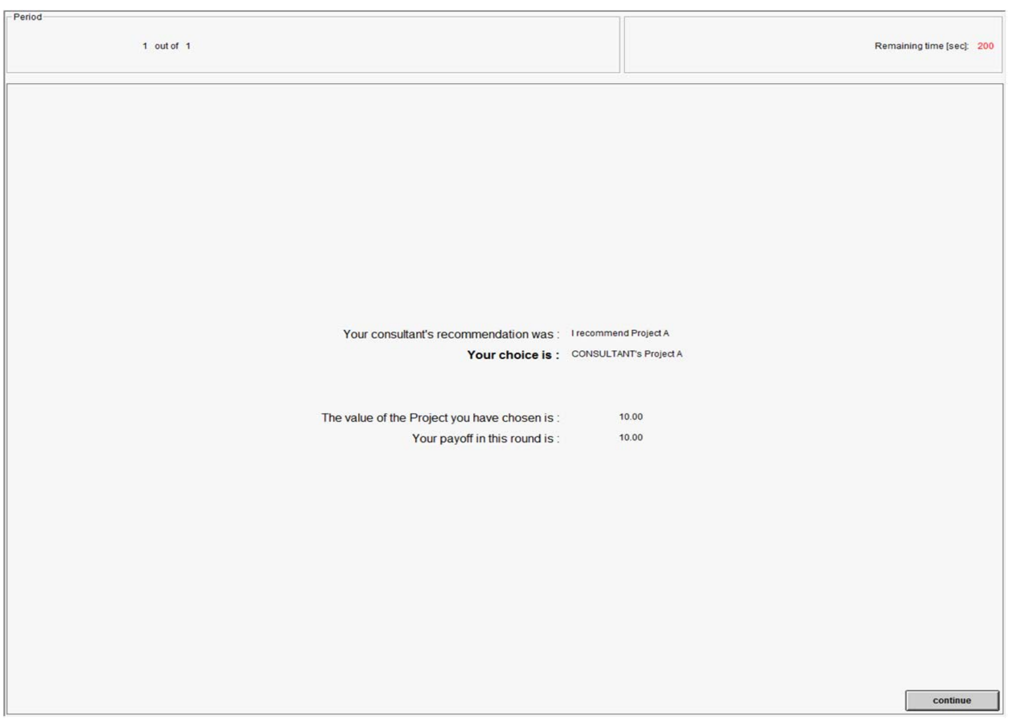

References

[1] Abbink, Klaus and Abdolkarim Sadrieh. 2009. The pleasure of being nasty. *Economics Letters* 105(3): 306–308.

[2] Armstrong, Mark and Jidong Zhou. 2011. Paying for prominence. *Economic Journal* 121(556): 368–395.

[3] Armstrong, Mark, John Vickers, and Jidong Zhou. 2009. Prominence and consumer search. *RAND Journal of Economics* 40(2), 209–33.

[4] Athey, Susan and Glenn Ellison. 2012. Position auctions with consumer search. *Quarterly Journal of Economics*, forthcoming.

[5] Battaglini, Marco. 2002. Multiple referrals and multidimensional cheap Talk. *Econometrica* 70(4): 1379–1401.

[6] Blume, Andreas, Douglas V. DeJong, Yong-Gwan Kim, and Geoffrey B. Sprinkle. 1998. Experimental evidence on the evolution of meaning of messages in sender-receiver games. *American Economic Review* 88(5): 1323–1340.

[7] Blume, Andreas, Douglas V. DeJong, Yong-Gwan Kim, and Geoffrey B. Sprinkle. 2001. Evolution of communication with partial common interest. *Games and Economic Behavior* 37: 79–120.

[8] Cai, Hongbin, and Joseph Tao-Yi Wang. 2006. Overcommunication in strategic information transmission games. *Games and Economic Behavior* 56(1): 7–36.

[9] Cain, Daylian M., George Loewenstein, and Don A. Moore. 2005. The dirt on coming clean: perverse effects of disclosing conflicts of interest. *Journal of Legal Studies* 34(1): 1–25.

[10] Chakraborty, Archishman and Rick Harbaugh. 2007. Comparative cheap talk. *Journal of Economic Theory* 132(1): 70–94.

[11] Chakraborty, Archishman and Rick Harbaugh. 2010. Persuasion by cheap talk. *American Economic Review* 100(5): 2361–2382.

[12] Chakraborty, Archishman and Rick Harbaugh. 2011. Persuasive puffery. Working paper.

[13] Chakraborty, Archishman and Bilge Yilmaz. 2011. Authority, consensus, and governance. Working paper.

[14] Che, Yeon-Koo, Wooter Dessein, and Navin Kartik. 2011. Pandering to persuade. *American Economic Review*, forthcoming.

[15] Chen, Yongmin and Chuan He. 2011. Paid placement: advertising and search on the internet. *Economic Journal* 121(556): 309–328.

[16] Chung, Wonsuk. 2011. Product recommendations by a biased seller. Working paper.

[17] Crawford, Vincent. 2003. Lying for strategic advantage: Rational and boundedly rational misrepresentations of intentions. *American Economic Review* 93(1): 133–149.

[18] Crawford, Vincent P. and Joel Sobel. 1982. Strategic information transmission. *Econometrica* 50(6): 1431–1450.

[19] Darby, Michael R. and Edi Karni. 1973. Free competition and the optimal amount of fraud. *Journal of Law and Economics* 16(1): 67–88.

[20] Dickhaut, John, Kevin McCabe, and Arijit Mukherji. 1995. An experimental study of strategic information transmission. *Economic Theory* 6(3): 389–403.

[21] Dulleck, Uwe, Rudolf Kerschbamer, and Matthias Sutter. 2011. The economics of credence goods: an experiment on the role of liability, verifiability, reputation, and competition. *American Economic Review 101* (2): 526–555.

[22] Fischbacher, Urs. 2007. z-Tree: Zurich toolbox for ready-made economic experiments. *Experimental Economics* 10(2): 171–178.

[23] Gentzkow, Matthew, and Jesse M. Shapiro. 2006. Media bias and reputation. *Journal of Political Economy*, 114(2): 280–316.

[24] Gibson, Rajna, Carmen Tanner, and Alexander F. Wagner. 2012. Preferences for truthfulness: heterogeneity among and within individuals. *American Economic Review*, forthcoming.

[25] Gneezy, Uri. 2005. Deception: the role of consequences. *American Economic Review* 95(1): 384–394.

[26] Hodler, Roland, Simon Loertscher, and Dominic Rohner. 2010. Biased experts, costly lies, and binary decisions. Working paper.

[27] Hurkens, Sjaak and Navin Kartik. 2009. Would I lie to you? On social preferences and lying aversion. *Experimental Economics* 12(2): 180–192.

[28] Inderst, Roman and Marco Ottaviani. 2009. Misselling through agents. *American Economic Review* 99(3): 883–908.

[29] Inderst, Roman and Marco Ottaviani. 2012. Competition through commissions and kickbacks. *American Economic Review* 102(2): 780–809.

[30] Jindapon, Paan and Carlos Oyarzun. 2011. Persuasive talk: delegation, competition, and announcements. Working paper.

[31] Kartik, Navin. 2009. Strategic communication with lying costs. *Review of Economic Studies* 76(4): 1359–1395.

[32] Kartik, Navin, Marco Ottaviani, and Francesco Squintani. 2007. Credulity, lies, and costly talk. *Journal of Economic Theory* 134(1): 93–116.

[33] Lai, Ernest K., Wooyoung Lim, and Joseph Tao-Yi Wang. 2011. Experimental implementations and robustness of fully revealing equilibria in multi-dimensional cheap talk. Working paper.

[34] Levy, Gilat and Ronny Razin. 2007. On the limits of communication in multidimensional cheap talk: a comment. *Econometrica* 75(3): 885–893.

[35] Li, Ming and Kristóf Madarász. 2008. When mandatory disclosure hurts: expert advice and conflicting interests. *Journal of Economic Theory* 139(1): 47–74.

[36] Madrian, Brigitte C. and Dennis F. Shea. 2001. The power of suggestion: inertia in 401(k) participation and savings behavior. *Quarterly Journal of Economics* 116 (4): 1149–1187.

[37] Milgrom, Paul R. 1981. Good news and bad news: representation theorems and applications. *Bell Journal of Economics* 12(2): 380–391.

[38] Morgan, John and Phillip C. Stocken. 2003. An analysis of stock recommendations. *RAND Journal of Economics* 34(1): 183–203.

[39] Morris, Stephen. 2001. Political correctness. *Journal of Political Economy* 109(2): 231–265.

[40] Nagel, Rosemarie. 1995. Unraveling in guessing games: an experimental study. *American Economic Review* 85(5): 1313–1326.

[41] Pitchik, Carolyn and Andrew Schotter. 1987. Honesty in a model of strategic information transmission. *American Economic Review* 77(5): 1032–1036.99

[42] Sánchez-Pagés, Santiago and Marc Vorsatz. 2007. An experimental study of truth-telling in a sender–receiver game. *Games and Economic Behavior* 61(1): 86–112.

[43] Sobel, Joel. 1985. A theory of credibility. *Review of Economic Studies* 52(4): 557–573.

[44] Stahl, Dale O. and Paul W. Wilson. 1995. On players' models of other players: theory and experimental evidence. *Games and Economic Behavior* 10(1): 218–254.

[45] Sutter, Matthias. 2009. Deception through telling the truth?! Experimental evidence from individuals and teams. *Economic Journal* 119(534): 47–60.

[46] Thaler, Richard H. and Cass R. Sunstein. 2008. *Nudge: Improving Decisions about Health, Wealth, and Happiness.* Yale University Press.

[47] Vespa, Emanuel and Alistair Wilson. 2012. Communication with multiple senders and multiple dimensions: an experiment. Working paper.

[48] Wang, Joseph Tao-yi, Michael Spezio, and Colin F. Camerer. 2010. Pinocchio's pupil: using eyetracking and pupil dilation to understand truth telling and deception in sender-receiver games. *American Economic Review* 100(3): 984–1007.