

Manuscript version: Working paper (or pre-print)

The version presented here is a Working Paper (or 'pre-print') that may be later published elsewhere.

Persistent WRAP URL:

<http://wrap.warwick.ac.uk/134298>

How to cite:

Please refer to the repository item page, detailed above, for the most recent bibliographic citation information. If a published version is known of, the repository item page linked to above, will contain details on accessing it.

Copyright and reuse:

The Warwick Research Archive Portal (WRAP) makes this work by researchers of the University of Warwick available open access under the following conditions.

Copyright © and all moral rights to the version of the paper presented here belong to the individual author(s) and/or other copyright owners. To the extent reasonable and practicable the material made available in WRAP has been checked for eligibility before being made available.

Copies of full items can be used for personal research or study, educational, or not-for-profit purposes without prior permission or charge. Provided that the authors, title and full bibliographic details are credited, a hyperlink and/or URL is given for the original metadata page and the content is not changed in any way.

Publisher's statement:

Please refer to the repository item page, publisher's statement section, for further information.

For more information, please contact the WRAP Team at: wrap@warwick.ac.uk.

Manipulative Disclosure

Claudio Mezzetti

March 2020

No: 1250

Warwick Economics Research Papers

ISSN 2059-4283 (online)

ISSN 0083-7350 (print)

Manipulative Disclosure *

Claudio Mezzetti

University of Queensland

School of Economics

&

University of Warwick

Department of Economics

claudiomezzetti@outlook.com

March 5, 2020

Abstract

This paper studies disclosure of verifiable information by a privately informed expert. It shows that if the direction of the expert's bias is uncertain, then a positive measure of expert types manipulate the decision maker fully, inducing her to choose their ideal outcome. Most other types manipulate partially. The decision maker obtains her first best outcome only if the expert is unbiased or the state of the world is a boundary point of the state space and the expert prefers a more extreme outcome. Experts benefit from being poker faced and the decision maker's lack of familiarity with the problem.

Journal of Economic Literature Classification Numbers: D44, D82.

Keywords: Verifiable disclosure of information, experts and decision makers, manipulation, pooling, senders and receivers, skepticism, uncertain bias, unravelling.

* I would like to thank the participants at seminars and workshops where previous versions of this paper were presented for their helpful comments and suggestions. This work has been funded in part by Australian Research Council grant DP190102904.

1 Introduction

In perfectly competitive markets prices convey all the information agents need to make their consumption and production decisions, but many important economic decisions are not made through the market and many markets are not perfectly competitive. Internal transactions within firms and organizations rarely rely on the price system; governments regulate and intervene in many sectors of the economy; agents (e.g., managers, stock analysts, lobbyists, politicians, professionals experts) with market power and private information are pervasive. Decision makers must thus often rely on information transmitted by other, more informed, parties.

Informed parties sometimes have no stake in the different alternatives the decision maker may select. They may even have the best interest of the decision maker at heart, as when one moves to a job in a new city and a new colleague gives advice about the city's attractions. Other times, experts have stakes in the decision. In this latter case, how should one interpret the information coming from an expert? If the expert is self-serving, will she be able to manipulate the decision maker? Can she get what she wants?

There is a large economic literature devoted to the study of this general problem. It is useful to distinguish between three strands: disclosure of verifiable information, cheap talk and Bayesian persuasion. The latter has the least in common with this paper. Springing from the contributions of Kamenika and Gentzkov (2011) and Rayo and Segal (2010), the Bayesian persuasion, or information design, literature focuses on the case when a designer with commitment power and no information about the state of nature controls the information received by the decision maker. The designer's goal in this literature is typically to induce the decision maker to undertake a specific action independently of the state of the world.

On the contrary, and as in most of the cheap talk literature, this paper is concerned with the case when the expert has private information about the state of the world and her ideal decision depends on it. The decision maker is modeled as a single agent,

but in some applications it could be viewed as representing “the market.” In the cheap talk literature, the informed agent simply talks. The messages sent have no verifiable meaning. Anything that expert A says could be said by expert B , irrespective of the underlying information the two have. The general insight from the cheap talk literature is that only coarse information can be transmitted and the outcome chosen by the decision maker is not sensitive to small changes in the state of the world (e.g., see the survey by Sobel, 2013). When their preferences are imperfectly aligned, neither the expert nor the decision maker will be able to reach their ideal outcome.

The third strand of the literature studies the disclosure of verifiable information. Experts that possess verifiable information can decide to disclose all or part of it, but they cannot manufacture and disclose facts that are not true. The fundamental insight from the literature on disclosure of verifiable information is that information has little value to the expert; the expert cannot turn it to her advantage. Any news that is not reported is viewed as bad news for the expert. By putting together the evidence disclosed and the information not reported, the decision maker is able to determine which alternative is best.

Many real life examples have both elements of cheap talk and verifiable information. For example news and media companies report facts and opinions. Partial verifiability is common. If rating agency A has discovered hidden liabilities in the balance sheet of company C , it can fully or partially report them. If there are no such liabilities in the balance sheet of company D , then rating agency B cannot mimic the disclosure of A . Media companies can decide what facts to report, but they may also report opinions. Scientists working for a biotechnology firm may or may not publish the results of their research, but they cannot publish research they have not undertaken.

The purpose of this paper is to take a new look at the disclosure of information when the expert may decide how much information to disclose, with full disclosure, partial disclosure and cheap talk all being possible. Its novelty is to show that there

are natural circumstances, not explored in the existing literature, under which an expert can make good use, in a self-serving way, of her private information. Indeed, often the expert may be able to manipulate the decision maker to the point of inducing her to choose the expert's favorite alternative.

The standard assumption in the literature on disclosure of verifiable information and cheap talk is that the preferences of the expert are well known. A buyer, for example, knows that the salesman is trying to sell as much as possible of a given product. The buyer can then safely assume that all positive verifiable information about the quality of the product will be disclosed, because it will induce the buyer to buy more. The privacy of the salesman information unravels; all information will be disclosed and the buyer will be able to select her first best alternative (Grossman, 1981, Grossman and Hart, 1980, and Milgrom, 1981).

The assumption that the preferences of the expert are fully known nicely fits many situations, but not all. Consider a rating agency discovering bad news about a large company; it will have an incentive to investigate thoroughly and report the bad news in order to preserve its reputation, but it will also have an incentive to gloss over the bad news in order to appease the large company and get its business again in the future. A biotechnology firm will have an incentive to let its scientists publish all preliminary research in order to gain financing, but it also has an incentive not to publish to make it more difficult for competitors to imitate. Even in the case of a salesman, it is not always clear what her incentives are. When shopping for a washing machine, for example, a buyer may not know whether the salesman stands to gain more by selling brand *A* or brand *B*. In general, when information is not unidimensional it is difficult to be sure about the precise direction of the expert's preferences.

So, what can the expert accomplish if the direction of her bias is not known by the decision maker? This paper shed lights on this question. It takes a minimalist modeling approach, by changing as little as possible from an otherwise standard

model. The model used is the work-horse model of the literature on cheap talk, in which the expert knows the state of the world, which coincides with the first best alternative for the decision maker. Rather than pure cheap talk messages, the expert must send messages that contain the truth, but the truth can be swamped by cheap talk to the point that the message does not contain any information. The one fundamental departure from the existing literature is that the decision maker does not know the direction of the expert's bias. The expert may want to distort choices up or down.

Many new insights arise. Full disclosure is never an equilibrium. Quite to the contrary, almost all expert types are able to manipulate the decision maker and induce the choice of an alternative that is close to their favorite outcome. Indeed, a positive measure of expert types will be able to obtain their first best outcome. We call this full manipulation. Expert types who observe a non-extreme state are all able to manipulate the decision maker. Those with non-extreme biases will achieve full manipulation, while those with extreme biases will partially manipulate. The only experts that obtain the same outcome as under full disclosure are the ones whose preferences are perfectly aligned with the decision maker, and the ones that have observed an extreme state and would prefer an even more extreme outcome. There is no room for these types to distort the decision in a direction that is favorable to them.

The intuition for why manipulation occurs is simple and instructive. Because the direction of the expert's bias is uncertain, the decision maker cannot fully decode partial disclosures that do not fully reveal the state. A partial disclosure is consistent with an expert that has observed a low state but would prefer a higher choice, and with an expert that has observed a high state but would prefer a lower choice. The compromise choice of the decision maker ends up being exactly what both types of expert want.

We show that, if given a choice, experts would want to maintain uncertainty

about their biases and incentives. Decision makers, on the other hand should not worry about getting advice from a strongly biased expert, as long as the direction of the bias is known. An increase in the size of the expert's mean bias benefits the decision maker and hurts the expert. This is because disclosures from more highly biased experts are easier to read.

The paper proceeds as follows. Section 2 introduces the model. Section 3 shows that any equilibrium involves full manipulation by a positive measure of expert types. Section 4 studies the main features of the equilibrium. Section 5 contains comparative statics results. Section 6 discusses related literature and Section 7 concludes. Proofs missing from the main text are in the Appendix

2 The Model

There are two agents, Dema is a decision maker and Knowyn is an expert. They care about the match of the choice of action made by Dema with the true state of the world. The state of the world $\omega \in [\underline{\omega}, \bar{\omega}]$ is known by Knowyn but not by Dema, who views it as a random variable with prior distribution $F(\omega)$. We assume that F admits a positive density $f(\omega)$ at all points $\omega \in [\underline{\omega}, \bar{\omega}]$.

The game the two agents play has two stages. In the first stage, Knowyn decides the information to disclose to Dema. To allow full disclosure, partial disclosure and pure cheap talk as feasible options, we postulate that the expert's disclosure policy consists of choosing an interval $[a, b]$ which must contain the true state ω . If the expert chooses an interval that contains the support $[\underline{\omega}, \bar{\omega}]$ of the state of the world, then her is a pure cheap talk message. If she chooses a non-degenerate interval $[a, b]$ that is a strict subset of $[\underline{\omega}, \bar{\omega}]$, then she makes a partial disclosure, as she rules out that the state belongs to $[\underline{\omega}, \bar{\omega}] \setminus [a, b]$. If she chooses interval $[\omega, \omega]$, then she fully discloses. In the second stage, after having observed Knowyn's disclosure, Dema chooses an action, or alternative. The set of possible alternatives is the real line.

After observing the expert's disclosure, Dema knows for sure that the true state of the world belongs to the disclosed interval. This assumption is what distinguishes our model from a pure cheap talk model, without ruling out that all messages sent are cheap talk. Naturally, Dema's posterior probability about the state of the world needs also to take into account the strategy chosen by Knowyn; using Knowyn's equilibrium strategy allows Dema to further refine her beliefs.

We assume that the decision maker's goal is to choose the alternative that minimizes the Euclidean distance from the state of the world. That is, we take the state of the world to represent the decision maker's ideal choice. Thus, if x is the choice and ω is the state, then Dema's payoff is

$$V(x, \omega) = -(x - \omega)^2.$$

We could adopt a more general formulation; what we will use is that Dema's optimal choice is the alternative which equals the expected value of the state, conditional on the information directly disclosed by the expert and the expert's equilibrium strategy.

The expert's preferences are not perfectly aligned with Dema's preferences. The simplest way to model this is to assume that the expert has a bias β and that her payoff depends on the distance between the chosen alternative x and $\omega + \beta$:

$$U(x, \omega, \beta) = -(x - \omega - \beta)^2.$$

The ideal choice of the expert diverges from the ideal choice of Dema, state ω , by the bias β .

Our payoff functions coincide with those in the literature on cheap talk started by Crawford and Sobel (1982), but unlike most of the literature on cheap talk and disclosure of verifiable information, we assume that the bias β is private information of the expert and viewed by the decision maker as a random variable with support $[\beta_L, \beta_H]$. As we shall see, the interesting case is when $\beta_L < 0 < \beta_H$; in such a case Dema is uncertain about the direction (as opposed to just the size) of Knowyn's

bias. We allow β to be correlated with ω and denote with $G(\beta | \omega)$ its conditional distribution. We assume that G admits a positive density g for all β and ω .

An expert's *type* is a pair (ω, β) of a state and a bias; let $\mathcal{T} = \{(\omega, \beta) : \omega \in [\underline{\omega}, \bar{\omega}], \beta \in [\beta_L, \beta_H]\}$ be the set of possible types of the expert. Let \mathcal{C} be the set of closed subintervals of the interval $[\underline{\omega}, \bar{\omega}]$. A strategy for the expert is a function $s_K : \mathcal{T} \rightarrow \mathcal{C}$ with the restriction that $\omega \in s_K(\omega, \beta)$; the true state must belong to the reported interval. A strategy for Dema is a function $s_D : \mathcal{C} \rightarrow \mathbb{R}$. We are interested in the (pure-strategy) perfect Bayesian equilibria (PBE) of the game. A PBE is a triple $\langle s_K, s_D, \mu \rangle$, where s_K and s_D are the agents' strategies and μ is a map that associates to each $I \in \mathcal{C}$ a probability density over \mathcal{T} , representing Dema's posterior beliefs about Knowyn's type after Knowyn's disclosure of I . Let $s_K^{-1}(I) = \{(\omega, \beta) \in \mathcal{T} : s_K(\omega, \beta) = I\}$ be the inverse image of I . To be a PBE the triple $\langle s_K, s_D, \mu \rangle$ must satisfy the following conditions:

$$\mu(\omega, \beta | I) = \begin{cases} \frac{f(\omega) g(\beta | \omega)}{\int_{(\tilde{\omega}, \tilde{\beta}) \in s_K^{-1}(I)} f(\tilde{\omega}) g(\tilde{\beta} | \tilde{\omega}) d\tilde{\omega} d\tilde{\beta}} & \text{if } \omega \in I \text{ and } s_K^{-1}(I) \neq \emptyset \\ 0 & \text{if } \omega \notin I \end{cases} \quad (1)$$

$$s_D(I) = \mathbb{E}_\mu[\omega | I] =: \int_{(\omega, \beta) \in s_K^{-1}(I)} \omega \mu(\omega, \beta | I) d\omega d\beta \quad (2)$$

$$s_K(\omega, \beta) \in \arg \min_{\{I \in \mathcal{C}: \omega \in I\}} (s_D(I) - \omega - \beta)^2. \quad (3)$$

Condition (1) says that on the equilibrium path Dema's posterior beliefs about Knowyn's type are consistent with Knowyn's strategy and put zero mass on states outside the disclosed interval.¹ Condition (2) says that Dema's equilibrium strategy is to choose the expected value of the state conditional on her posterior beliefs and

¹Thus, the only restriction on beliefs that follow disclosure of an interval I that is not on the equilibrium path is that all mass be on I .

the observed disclosure. Condition (3) requires that Knowyn chooses the disclosure interval containing the state ω that maximizes her payoff, given Dema's strategy.

Define the composite outcome map $\alpha : \mathcal{T} \rightarrow \mathbb{R}$ as $\alpha = s_D \circ s_K$; $\alpha(\omega, \beta) = s_D(s_K(\omega, \beta))$ is the alternative chosen in equilibrium by Dema when Knowyn's type is (ω, β) . It is the norm for disclosure games to have multiple equilibria that induce the same outcome map. For this reason, in discussing equilibrium we will focus on the equilibrium map and will not present all possible outcome equivalent equilibria that generate it.²

It is instructive to end this section with the benchmark case in which the direction of the bias is common knowledge; that is, β_L and β_H have the same sign. Without loss of generality, take $0 \leq \beta_L \leq \beta_H$. It is then commonly known that Knowyn would like to push Dema's choice upward. Given this, it is natural for Dema to have pessimistic beliefs; given any disclosed interval $[a, b]$, Dema puts all probability mass on $\omega = a$ and hence chooses $s_D([a, b]) = a$. Given these beliefs, it is a best reply for Knowyn to disclose an interval with the true state as the left boundary: $s_K(\omega, \beta) = [\omega, b]$ with $b \geq \omega$.³ The equilibrium outcome is equivalent to the expert choosing full disclosure (i.e., choosing interval $[\omega, \omega]$). This is the same result as when the bias is fully known (e.g., see Milgrom, 1981 and Grossman, 1981); Dema always obtains the information needed to select her first best alternative.

Proposition 1 *If $0 \leq \beta_L \leq \beta_H$, then the unique equilibrium outcome map is $\alpha(\omega, \beta) = \omega$ for all $(\omega, \beta) \in \mathcal{T}$; the expert's verifiable disclosure fully reveals the true state of the world and the decision maker achieves her first-best outcome.*

Proof. First, note that in all equilibria it must be $\alpha(\omega, \beta) \geq \omega$ for all (ω, β) . If it were $\alpha(\omega, \beta) < \omega$ for some type (ω, β) , then type (ω, β) could profitably deviate by

²See footnote 3.

³Note that there are multiple, outcome equivalent, equilibria, that only vary in the right boundary of the interval $[\omega, b]$.

disclosing $[\underline{\omega}, \omega]$, which would certainly lead Dema to choose $s_D([\omega, \omega]) = \omega$. Second, suppose that $\alpha(\omega', \beta') = x > \omega'$ for some (ω', β') . Then, the equilibrium disclosure I' made by Knowyn's type (ω', β') must also have been made by at least another type (ω'', β'') with $\omega'' > x$, because it must be $x = \mathbb{E}_\mu[\omega | I']$. But then type (ω'', β'') would gain by deviating and disclosing $[\omega'', \omega'']$. ■

Knowing the direction of Knowyn's bias is as good for Dema as knowing the true value of the bias. In the remainder of the paper, we will assume that $\beta_L < 0 < \beta_H$ and focus on the case in which the bias direction is uncertain.

3 Full Manipulation

In this section, we derive some general properties of the equilibrium outcome map α . The first lemma is obvious: unbiased expert types induce the decision maker to choose the alternative that is first best optimal for both agents. When the interests of Knowyn and Dema are perfectly aligned, the expert could always fully disclose the state, thus making sure that the decision maker picks the first best outcome.

Lemma 1 *For all $\omega \in [\underline{\omega}, \bar{\omega}]$, $\alpha(\omega, 0) = \omega$.*

Proof. Suppose, to the contrary, that $\alpha(\omega, 0) = x \neq \omega$. By (1), disclosing interval $[\omega, \omega]$ induces Dema to choose ω , and is thus a profitable deviation for type $(\omega, 0)$, a contradiction. ■

Next we show that for any given ω the map $\alpha(\omega, \beta)$ must be continuous and weakly increasing in β , because types with the same state can mimic each other and disclose the same interval. More precisely, the outcome α must either stay constant as β increases, or increase linearly with β and select Knowyn's first best outcome.

Lemma 2 *For all $\omega \in [\underline{\omega}, \bar{\omega}]$, $\alpha(\omega, \beta)$ is: (i) continuous and weakly increasing in β ; (ii) either constant in an interval around β or equal to $\omega + \beta$.*

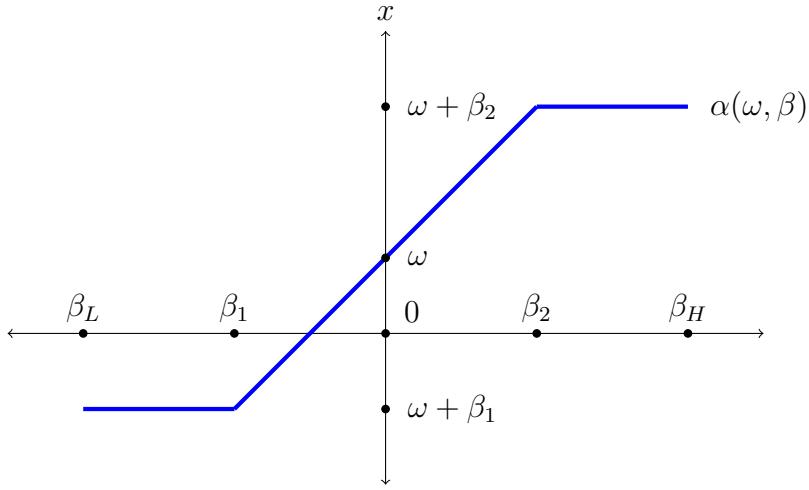


Figure 1: $\alpha(\omega, \beta)$ as a function of β for given ω .

By Lemma 2, it must be either $\alpha(\omega, \beta) = \omega + \beta$ and the expert gets its favorite's outcome, or $\alpha(\omega, \beta)$ is constant in an interval around β . The next lemma shows that, for any given state ω , the outcome map $\alpha(\omega, \beta)$ can only be constant for expert types with an “extreme” bias – that is, types with a bias in an interval including one of the boundary points β_L and β_H . Expert types with a “moderate” bias in an interval $[\beta_1, \beta_2]$ around $\beta = 0$, on the other hand, obtain their ideal outcome $\omega + \beta$.⁴ As shown in Figure 1, extreme left biased types with bias in the interval $[\beta_L, \beta_1]$ induce Dema to choose $\omega + \beta_1$, while extreme right biased types with bias in the interval $[\beta_2, \beta_H]$ induce choice $\omega + \beta_2$, so they are still able to manipulate Dema's choice moving it in the direction of their bias.

Lemma 3 *For all $\omega \in [\underline{\omega}, \bar{\omega}]$, it must be: (i) $\alpha(\omega, \beta) = \omega + \beta$ in an interval $[\beta_1, \beta_2]$ with $\beta_L \leq \beta_1 \leq 0 \leq \beta_2 \leq \beta_H$, (ii) $\alpha(\omega, \beta) = \omega + \beta_1$ if $\beta \leq \beta_1$, and (iii) $\alpha(\omega, \beta) = \omega + \beta_2$ if $\beta \geq \beta_2$.*

The next lemma shows that $\alpha(\omega, \beta)$ must be monotone in ω .

⁴Lemma 3 does not rule out the possibility that the interval is degenerate and only includes the unbiased expert type with $\beta = 0$. This possibility will be ruled out later, by Proposition 2. In general, the boundary points β_1, β_2 of the interval depend on ω .

Lemma 4 For all $\beta \in [\beta_L, \beta_H]$, the map $\alpha(\omega, \beta)$ is weakly increasing in ω .

We say that the equilibrium outcome map α involves full manipulation by expert type (ω, β) (or, alternatively, type (ω, β) fully manipulates Dema) if $\alpha(\omega, \beta) = \omega + \beta$; that is, Knowyn type (ω, β) induces the choice of her ideal outcome. We now prove the main result of this section: a positive measure of Knowyn's types achieve their first best outcome. The intuition for the proof of the result is the following. Suppose interval $[a, b]$ is disclosed. Dema will choose an outcome $x \in [a, b]$ because any state outside the interval has zero probability of being the true state. Then, all Knowyn's types that observe a state $\omega \in [a, b]$ and have x as their first best outcome (i.e., types such that $\omega + \beta = x$) are able to fully manipulate and make sure that x is chosen by Dema.

Proposition 2 In all equilibria, a positive measure of the expert's types (ω, β) fully manipulates the decision maker; that is, for a positive measure of types it is $\alpha(\omega, \beta) = \omega + \beta$.

Proof. The proof is by contradiction. Suppose, to the contrary, that at most a zero measure of Knowyn's types (ω, β) induce their first best outcome $\omega + \beta$. By Lemmas 1, 2 and 4, $\alpha(\omega, \beta)$ is continuous almost everywhere in an open rectangle around any point $(\omega_1, 0)$. Let $\varepsilon > 0$ and $\omega_1 \in (\underline{\omega} + \varepsilon, \bar{\omega} - \varepsilon)$. Take the open rectangle of types $R = \{(\omega, \beta) : \omega_1 - \varepsilon < \omega < \omega_1 + \varepsilon, -\varepsilon < \beta < \varepsilon\}$. If only a zero measure set of types (ω, β) in R are such that $\alpha(\omega, \beta) = \omega + \beta$, then by Lemmas 1 and 2 it must be $\alpha(\omega, \beta) = \omega$ for all types in R . Consider the interval $I = [\omega_1 - \varepsilon, \omega_1 + \varepsilon]$. Disclosing interval I must induce an outcome $\omega_2 \in [\omega_1 - \varepsilon, \omega_1 + \varepsilon]$, and thus it is a profitable deviation for all types $(\omega, \beta) \in R$ with $\omega + \beta = \omega_2$; a contradiction. ■

It is important to emphasize once again that full manipulation by the expert is quite different from the standard full disclosure outcome that obtains when there is no uncertainty about the direction of the expert's bias. In essence, in our model partially verifiable disclosures (i.e., disclosures of a non degenerate interval) cannot be fully

deciphered by the decision maker. Disclosure $[a, b]$ could be made by a Knowyn type that has observed a state as low as a and has a positive bias, or by a type that has observed a state as high as b and has a negative bias. No matter what alternative x the decision maker picks after such a disclosure, there will be Knowyn's types that find x their first best choice and that can all pool and disclose $[a, b]$. Full manipulation is made possible by this ability to pool and disclose the same interval by expert's types with positive and negative bias that have a given outcome x as their ideal choice.

Our final result of this section pins down the conditions that the equilibrium map must satisfy at the boundary points of the state space. Expert types that observe state $\omega = \underline{\omega}$ and have a negative bias and types that observe state $\omega = \bar{\omega}$ and have a positive bias cannot do better than fully disclosing the state. The intuition is straightforward. Types that observe state $\omega = \underline{\omega}$ and have a negative bias prefer alternative $x = \underline{\omega}$ to any alternative $x > \underline{\omega}$. Their ideal outcome is some $x < \underline{\omega}$, but Dema would never choose such an outcome. Thus, among the outcomes that Knowyn can feasibly induce Dema to choose, Knowyn's preferences are aligned with Dema's. It is thus optimal for Knowyn to disclose interval $[\underline{\omega}, \underline{\omega}]$. Similarly, in the set $[\underline{\omega}, \bar{\omega}]$ of “feasible” alternatives, the preferences of Knowyn's types that observe state $\omega = \bar{\omega}$ and have a positive bias are perfectly aligned with Dema's preferences.

Proposition 3 *In all equilibria, it is $\alpha(\underline{\omega}, \beta) = \underline{\omega}$ if $\beta \leq 0$ and $\alpha(\bar{\omega}, \beta) = \bar{\omega}$ if $\beta \geq 0$.*

Proof. Consider Knowyn's type $(\underline{\omega}, \beta)$ with $\beta \leq 0$. Disclosing interval $[\underline{\omega}, \underline{\omega}]$ is possible for this type and induces Dema to select $x = \underline{\omega}$. Dema never selects $x < \underline{\omega}$ because it is known that $\omega \in [\underline{\omega}, \bar{\omega}]$ and Knowyn prefers $x = \underline{\omega}$ to any outcome $x > \underline{\omega}$. Similarly, disclosing interval $[\bar{\omega}, \bar{\omega}]$ is possible for Knowyn's type $(\bar{\omega}, \beta)$ with $\beta \geq 0$ and induces Dema to select $x = \bar{\omega}$, which is preferred by Knowyn to any $x < \bar{\omega}$. No $x > \bar{\omega}$ is ever selected by Dema. ■

4 The Uniform Model

To elaborate on the main results and insights and describe the equilibrium structure in the most transparent way, in this section we simplify the model by assuming that the expert's bias and the state of the world are independent, uniformly distributed, random variables.

4.1 The Unbounded Case

By Proposition 3, at the boundary of the state space types whose ideal decision is outside $[\underline{\omega}, \bar{\omega}]$ cannot do better than fully disclosing the state. The boundedness of the set of states of the world also constrains the equilibrium outcome for interior states that are close to the boundary. In order to gain intuition for the amount of manipulation the expert is able to achieve, and to focus on the equilibrium configuration in the interior of the state space, in this sub-section we make an additional modification to the model so as to eliminate the boundary effects. We assume that $\underline{\omega} = -\infty$, and $\bar{\omega} = \infty$. When combined with the assumption that the state is uniformly distributed, this implies that ω has an improper, diffuse distribution over the real line. In the next sub-section, we will go back to analyze the model with a bounded state space.

We begin with the additional assumption that $\beta_L = -\beta_H$. This implies that the expected bias is zero. More importantly, it implies that the expected value of the state of all Knowyn's types that have x as their favorite choice is x , $\mathbb{E}[\omega | \omega + \beta = x] = x$. Suppose that the two “twin” Knowyn's types $(x - \beta', \beta')$ and $(x + \beta', -\beta')$ disclose the same interval (e.g., they disclose $[x - \beta', x + \beta']$). Suppose no other type discloses this interval. It is then immediate that Dema's best response to such a disclosure is to choose outcome x , because x is the average value of the state. It is also immediate that the two twin Knowyn's types have no profitable deviation, because they induce Dema to choose their ideal outcome. Finally, note that since there is no state space boundary and $\beta_L = -\beta_H$, for any Knowyn's type (ω, β) there is a twin type $(\omega + 2\beta, -\beta)$. Thus,

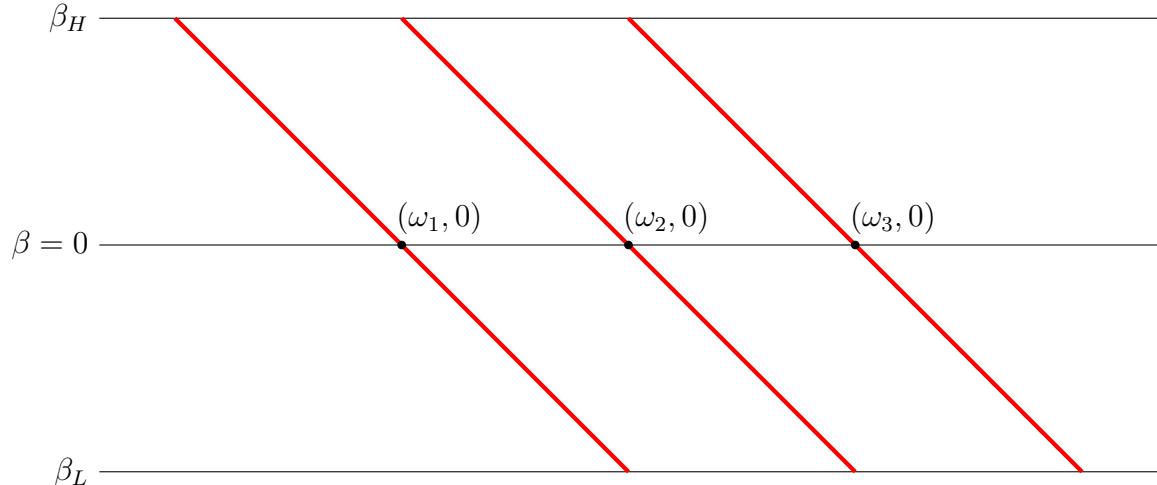


Figure 2: Unbounded Uniform Model, $-\beta_L = \beta_H$.

all Knowyn's types may just pool with their twin and fully manipulate the outcome. This argument is formalized in the following proposition.

Proposition 4 *In the unbounded uniform model, if $-\beta_L = \beta_H$, then $\alpha(\omega, \beta) = \omega + \beta$ for all expert types; the choice made by the decision maker always coincides with the expert's optimal (first-best) choice.*

Proof. We claim that the following equilibrium strategies implement the outcome map $\alpha(\omega, \beta) = \omega + \beta$. If $\beta \geq 0$, then $s_K(\omega, \beta) = [\omega, \omega + 2\beta]$ and if $\beta \leq 0$, then $s_K(\omega, \beta) = [\omega + 2\beta, \omega]$; $s_D([x, y]) = \frac{x+y}{2}$. First, note message $[x, y]$ is sent in equilibrium only by Knowyn's types $(x, \frac{y-x}{2})$ and $(y, \frac{x-y}{2})$. Hence it is $\mathbb{E}[\omega | [x, y]] = \frac{x+y}{2}$ and s_D is a best response for Dema. Second, given Dema's strategy, by following s_K all types of expert obtain their first-best outcome, hence they are playing a best response.⁵ ■

Figure 2 describes the equilibrium in this case. The state varies along the horizontal coordinate and the bias along the vertical coordinate of the two dimensional

⁵Note that there is also an equilibrium in which all Knowyn's types that have the same ideal choice x pool and disclose the same interval (e.g., they disclose interval $[x - \beta_H, x + \beta_H]$).

Cartesian space. All Knowyn's types on a negatively sloped 45 degree line have the same ideal outcome, $x = \omega_1, \omega_2, \omega_3$, and are able to induce Dema to choose it. In this version of the model, all expert types fully manipulate the decision maker.

In the unbounded version of our model, pure cheap talk messages, as we have defined them, are not available, but they could be easily reintroduced by adding an additional dimension to the expert's message. With this modification, the equilibrium in the unbounded uniform model with $\beta_L = -\beta_H$ can also be sustained by pure cheap talk messages. All that is needed is that all types (ω, β) with $\omega + \beta = x$, and only them, send the same message; no expert type would have an incentive to deviate. This is an exception. As we shall see, in all other versions of the model the equilibrium outcome cannot be sustained by pure cheap talk messages.

Suppose now that $\beta_H > -\beta_L$; that is, there are more Knowyn's types with a positive than with a negative bias. The average bias is positive and the expected value of the state of all Knowyn's types that have x as their first best choice is less than x , $\mathbb{E}[\omega | \omega + \beta = x] < x$. In this case, not all types have a twin with whom they can pool. More precisely, types with a high positive bias do not have a twin and they will only be able to partially manipulate the outcome by pooling with types that have a smaller ideal choice than them. Thus, a positive measure of expert types get their first-best outcome, while all other types get less than their first-best choice.

To aid the formal statement of the equilibrium in this case, we define the parameter β_H^* :

$$\beta_H^* = \beta_H - \sqrt{(\beta_H)^2 - (\beta_L)^2} \quad (4)$$

Proposition 5 *In the unbounded uniform model, if $-\beta_L < \beta_H$, then the equilibrium outcome map is the following:*

$$\alpha(\omega, \beta) = \begin{cases} \omega + \beta = & \text{if } \beta \leq \beta_H^* \\ \omega + \beta^* = & \text{if } \beta \geq \beta_H^* \end{cases}$$

Proof. Fix x and consider the set of Knowyn types that induce choice x according to the outcome map α ; denote this set by $T_\alpha(x)$. It is composed by the union of types $(x - \beta, \beta)$ with $\beta \in [\beta_L, \beta_H^*]$ and types $(x - \beta_H^*, \beta)$ with $\beta \in [\beta_H^*, \beta_H]$.

For all x , take the equilibrium strategy chosen by all Knowyn types in the set $T_\alpha(x)$ to be disclosing the interval $[x - \beta_H^*, x - \beta_L]$.

Compute the expected value of the state conditional on Knowyn's type belonging to $T_\alpha(x)$. Given the uniform distribution of types, it is:

$$\begin{aligned}\mathbb{E}[\omega | (\omega, \beta) \in T_\alpha(x)] &= (x - \beta_H^*) \cdot \frac{\beta_H - \beta_H^*}{\beta_H - \beta_L} + \left(\frac{x - \beta_H^* + x - \beta_L}{2} \right) \cdot \frac{\beta_H^* - \beta_L}{\beta_H - \beta_L} \\ &= x - \beta_H^* + \frac{(\beta_H^* - \beta_L)^2}{2(\beta_H - \beta_L)} \\ &= x\end{aligned}$$

Thus, given the equilibrium strategy chosen by Knowyn, after observing the disclosure of the interval $[x - \beta_H^*, x - \beta_L]$ it is a best reply for Dema to choose x . For any interval $[a, b]$ that Knowyn is not supposed to send in equilibrium, assume that Dema believes that the state is a and so she chooses a .

To see that the strategy of Knowyn is a best reply for all Knowyn types, first observe that types with bias lower than β_H^* induce the choice of their most preferred alternative, and hence could not profitably deviate. Types (ω, β) with bias $\beta > \beta_H^*$ induce choice $\omega + \beta_H^*$ by disclosing interval $[\omega, \omega + \beta_H^* - \beta_L]$; deviating to a disclosure of any interval $[a, b]$ is feasible for these types only if $a \leq \omega \leq b$, but such a disclosure can never move up the choice made by Dema. Thus the strategy of all Knowyn's types is a best reply to Dema's equilibrium strategy. This completes the proof. ■

Figure 3 describes the equilibrium outcome. The set of types that are able to induce the same choice $x = \omega_1, \omega_2, \omega_3$ (the x -decision set) includes the types on the portion of the negatively sloped 45 degrees line segment from type $(x - \beta_H^*, \beta_H^*)$ to type $(x - \beta_L, \beta_L)$, and the set of types on the vertical segment from type $(x - \beta_H^*, \beta_H^*)$ to type $(x - \beta_H^*, \beta_H)$. The critical value of the bias $\beta_H^* > 0$ is computed so that the expected value of the state on the x -decision set is equal to x . Thus, choosing x is

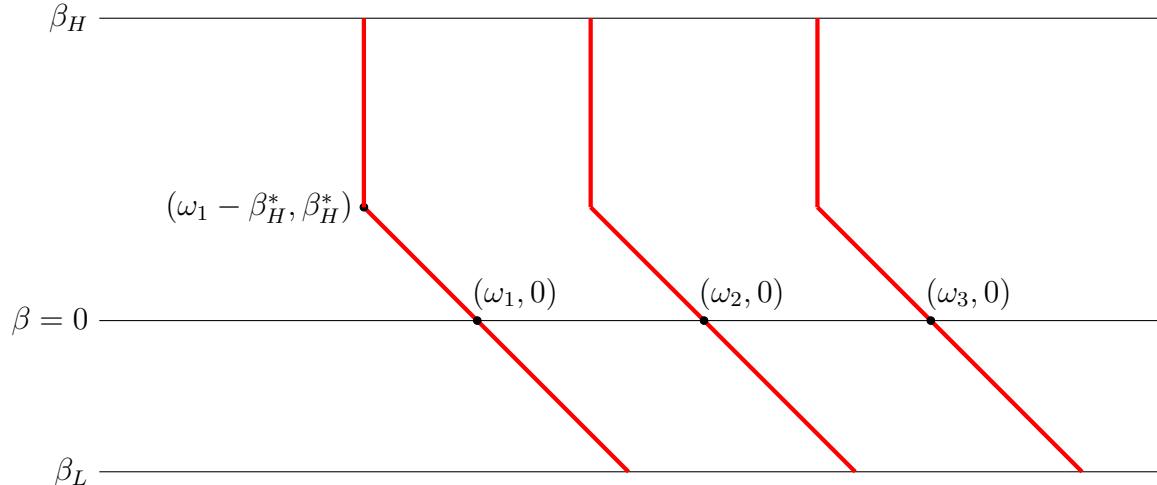


Figure 3: Unbounded Uniform Model, $-\beta_L < \beta_H$.

indeed a best response for Dema. Expert types on the diagonal portion of the x -decision set obtain their first best outcome. Hence only types on the vertical portion could possibly want to deviate, if they could induce a higher choice than x . The only way to induce a choice higher than x is to disclose an interval with a left boundary higher than $x - \beta_H^*$, which is impossible for the types of the vertical portion of the x -decision set, since $x - \beta_H^*$ is the true state that they observe.

The case $\beta_H < -\beta_L$, when there are more Knowyn's types with a negative than with a positive bias, is just the mirror image of the case just analyzed. Now $\mathbb{E}[\omega | \omega + \beta = x] > x$ and the equilibrium is described in Figure 4. Define:

$$\beta_L^* = \beta_L + \sqrt{(\beta_L)^2 - (\beta_H)^2}.$$

The x -decision set includes the types on the negatively sloped 45 degrees line segment from type $(x - \beta_H, \beta_H)$ to type $(x - \beta_L^*, \beta_L^*)$ and the types on the vertical segment from type $(x - \beta_L^*, \beta_L^*)$ to type $(x - \beta_L, \beta_L)$. Now it is types with a large negative bias (below β_L^*) that can only partially manipulate the decision maker. All types above β_L^* fully manipulate and induce Dema to pick their ideal choice.

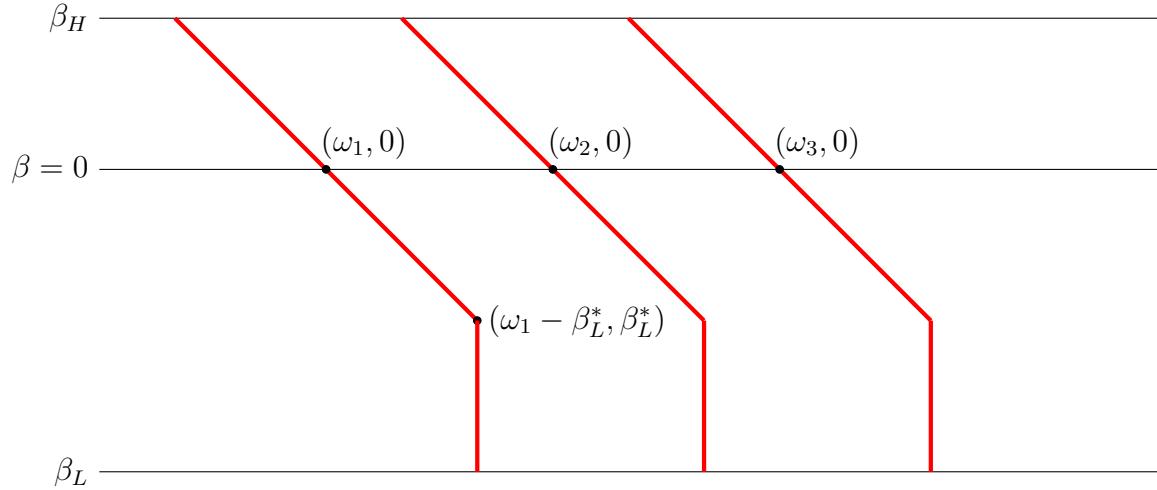


Figure 4: Unbounded Uniform Model, $-\beta_L > \beta_H$.

4.2 The Bounded Case

We now return to the bounded, uniform, model with state of the world $\omega \in [\underline{\omega}, \bar{\omega}]$. As we saw in Proposition 3, boundary types whose ideal decisions are outside the state space fully disclose. Boundary effects also play a role for types close to the boundary.

From now on, we assume that $\beta_H \geq -\beta_L$, as the case $\beta_H < \beta_L$ is just the mirror image of the case $\beta_H > -\beta_L$. Define the function $\beta_H^*(\omega)$, playing the same role as β_H^* in Proposition 5 and the function $\beta_L^*(\omega)$, playing a similar role to β_L^* :

$$\beta_H^*(\omega) = \beta_H - \sqrt{(\beta_H)^2 - \min\{(\bar{\omega} - \omega)^2, (\beta_L)^2\}} \quad (5)$$

$$\beta_L^*(\omega) = \beta_L + \sqrt{(\beta_L)^2 - \min\{(\omega - \underline{\omega})^2, (\beta_L)^2\}} \quad (6)$$

It turns out that we need to consider separately the case of zero expected bias and a state space larger than the bias space, from all other cases; hence we introduce the following mutually exclusive assumptions.

Assumption 1 *There is zero expected bias and a large state space; that is, the following two conditions hold:*

- (i) $\beta_H = -\beta_L$
- (ii) $\bar{\omega} - \underline{\omega} \geq -2\beta_L$.

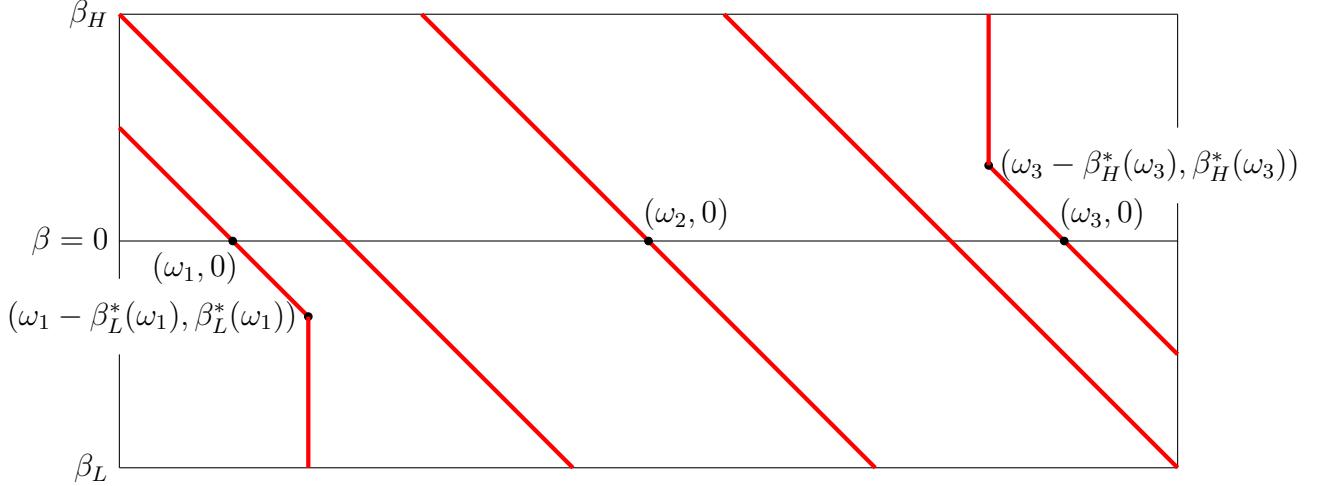


Figure 5: Bounded Uniform Model, $-\beta_L = \beta_H$, $\bar{\omega} - \underline{\omega} > 2\beta_H$.

Assumption 2 *It is either (i) $\beta_H > -\beta_L$, or (ii) $\beta_H = -\beta_L$ and $\bar{\omega} - \underline{\omega} < -2\beta_L$.*

We begin by considering the simpler case in which Assumption 1 holds, then move to the case when Assumption 2 holds instead. With zero mean bias and a large state space, the equilibrium is described in Figure 5. We can divide Knowyn's type space into three regions. All types (ω, β) in the middle region defined by $\underline{\omega} + \beta_H \leq \omega + \beta \leq \bar{\omega} - \beta_H$ are able to induce Dema to choose their ideal, or first best, alternative. This is because the expected value of the state for all types that have x as their ideal alternative is exactly x .

In the left region with $\omega + \beta < \underline{\omega} + \beta_H$, it is $\mathbb{E}[\omega | \omega + \beta = x] > x$; there are more types with a state above than below their ideal alternative x . Equilibrium is like in Figure 4; expert types with a positive bias and a moderately negative bias are able to fully manipulate Dema to choose their ideal alternative x . Types with ideal alternative x but an extremely negative bias (and hence a relative larger observed state) can only partially manipulate Dema; they induce her to choose an alternative higher than x . More precisely, the pooling region of expert types inducing the same choice x includes all types with a positive bias and x as ideal choice, types with small negative bias and x as ideal choice (i.e., types on the 45 degree segment going from

$(\underline{\omega}, x - \underline{\omega})$ to $(x - \beta_L^*(x), \beta_L^*(x))$, and types with a large negative bias and an ideal alternative smaller than x (i.e., the types on the vertical segment with state $x - \beta_L^*(x)$ and bias $\beta < \beta_L^*(x)$). Observe that $\beta_L^*(\omega)$ is strictly decreasing in ω for $\omega - \underline{\omega} < -\beta_L$. This says that the ability of experts with a negative bias to manipulate Dema gets smaller as the state ω gets closer to the left boundary $\underline{\omega}$. In the limit $\beta_L^*(\underline{\omega}) = 0$ and it is optimal for all types with negative bias and state $\omega = \underline{\omega}$ to fully disclose, as claimed in Proposition 3.

In the right region with $\omega + \beta > \bar{\omega} - \beta_H$ it is $\mathbb{E}[\omega | \omega + \beta = x] < x$; there are more types with a state below than above their ideal alternative x , and equilibrium is like in Figure 3. The pooling region of expert types inducing the same choice x includes all types with a negative bias and x as ideal choice, types with small positive bias and x as ideal choice (i.e., types on the 45 degree segment going from $(x - \beta_H^*(x), \beta_H^*(x))$ to $(\bar{\omega}, x - \bar{\omega})$), and types with a large positive bias and an ideal alternative larger than x (i.e., the types on the vertical segment with state $x - \beta_H^*(x)$ and bias $\beta > \beta_H^*(x)$). Note also that $\beta_H^*(\omega)$ is strictly decreasing in ω for $\omega > \bar{\omega} + \beta_L$. This says that the ability of experts with a positive bias to manipulate Dema gets smaller as the state ω gets closer to the right boundary. In the limit, $\beta_H^*(\bar{\omega}) = 0$ and it is optimal for all types with positive bias and boundary state $\omega = \bar{\omega}$ to fully disclose.

Proposition 6 *In the bounded uniform model, under Assumption 1 the equilibrium decision map is:*

$$\alpha(\omega, \beta) = \begin{cases} \omega + \beta & \text{if } \begin{cases} \underline{\omega} + \beta_H \leq \omega + \beta \leq \bar{\omega} - \beta_H \\ \omega + \beta \leq \underline{\omega} + \beta_H \quad \text{and} \quad \beta \geq \beta_L^*(\omega) \\ \omega + \beta \leq \bar{\omega} - \beta_H \quad \text{and} \quad \beta \leq \beta_H^*(\omega) \end{cases} \\ \omega + \beta_L^*(\omega) & \text{if } \omega + \beta \leq \underline{\omega} + \beta_H \text{ and } \beta \leq \beta_L^*(\omega) \\ \omega + \beta_H^*(\omega) & \text{if } \omega + \beta \geq \bar{\omega} - \beta_H \text{ and } \beta \geq \beta_H^*(\omega) \end{cases}$$

Proof. By Assumption 1 it is $-\beta_L = \beta_H$ and $\bar{\omega} - \underline{\omega} \geq 2\beta_H$. Consider first the set of expert types $T_\alpha(x)$ that induce Dema to choose an alternative x in the middle region

$\underline{\omega} + \beta_H \leq x \leq \bar{\omega} - \beta_H$. It is $T_\alpha(x) = \{(\omega, \beta) : \omega + \beta = x\}$. Suppose these types' strategy is to disclose interval $I_x = [x - \beta_H, x + \beta_H]$. It follows that Dema's best reply is to choose $\mathbb{E}[\omega | x \in I_x] = \mathbb{E}[\omega | \omega + \beta = x] = \frac{x - \beta_H + x + \beta_H}{2} = x$ and that these Knowyn's types have no profitable deviation, as x is their ideal choice.

The set of expert types that induce Dema to choose an alternative $x < \underline{\omega} + \beta_H$ is $T_\alpha(x) = \{(\omega, \beta) : \omega + \beta = x \text{ and } \beta \geq \beta_L^*(\omega)\} \cup \{(x - \beta_L^*(x), \beta) : \beta \leq \beta_L^*(x)\}$. Let these types' strategy be to disclose interval $I_x = [\underline{\omega}, x - \beta_L^*(x)]$. Similarly, the set of expert types that induce Dema to choose an alternative $x > \bar{\omega} - \beta_H$ is $T_\alpha(x) = \{(\omega, \beta) : \omega + \beta = x \text{ and } \beta \leq \beta_H^*(\omega)\} \cup \{(x + \beta_H^*(x), \beta) : \beta \geq \beta_H^*(x)\}$. Let these types' strategy be to disclose interval $I_x = [x - \beta_H^*(x), \bar{\omega}]$.

Let Dema's equilibrium strategy be: (i) choose x if the interval disclosed is the one that expert types in $T_\alpha(x)$ send in equilibrium. While if the interval disclosed is $[a, b]$ with no x such that $[a, b] = T_\alpha(x)$, then (ii) choose b if $a \leq \underline{\omega} + \beta_H$, $b \leq \bar{\omega} - \beta_H$; (iii) choose a if $b \geq \bar{\omega} - \beta_H$, $a \geq \underline{\omega} + \beta_H$ and (iv) choose y , with $\underline{\omega} + \beta_H \leq y \leq \bar{\omega} - \beta_H$, if $a \leq \underline{\omega} + \beta_H$, $b \geq \bar{\omega} - \beta_H$. Given Dema's equilibrium strategy, no type of Knowyn has a feasible deviation that is profitable. Deviations by types in $T_\alpha(x)$ with $x < \underline{\omega} + \beta_H$ can only lead to the choice of an alternative greater than x , while deviations by types in $T_\alpha(x)$ with $x > \bar{\omega} - \beta_H$ can only lead to the choice of an alternative smaller than x . Hence all feasible deviations lead to choices further away from Knowyn's ideal choice.

Using (6) and $-\beta_L = \beta_H$, note that Dema's best reply when observing a disclosure made by types in $T_\alpha(x)$ with $x < \underline{\omega} + \beta_H$ is to choose:

$$\begin{aligned}\mathbb{E}[\omega | x \in I_x] &= \frac{1}{x - \underline{\omega} + \beta_H} \left(\frac{\underline{\omega} + x - \beta_L^*(x)}{2} (x - \underline{\omega} - \beta_L^*(x)) + (x - \beta_L^*(x)) (\beta_H + \beta_L^*(x)) \right) \\ &= x.\end{aligned}$$

Finally, using (5), note that Dema's best reply when observing a disclosure made by types in $T_\alpha(x)$ with $x > \bar{\omega} - \beta_H$ is to choose:

$$\begin{aligned}\mathbb{E}[\omega | x \in I_x] &= \frac{1}{\beta_H + \bar{\omega} - x} \left(\frac{x - \beta_H^*(x) + \bar{\omega}}{2} (\beta_H^*(x) + \bar{\omega} - x) + (x - \beta_H^*(x)) (\beta_H - \beta_H^*(x)) \right) \\ &= x.\end{aligned}$$

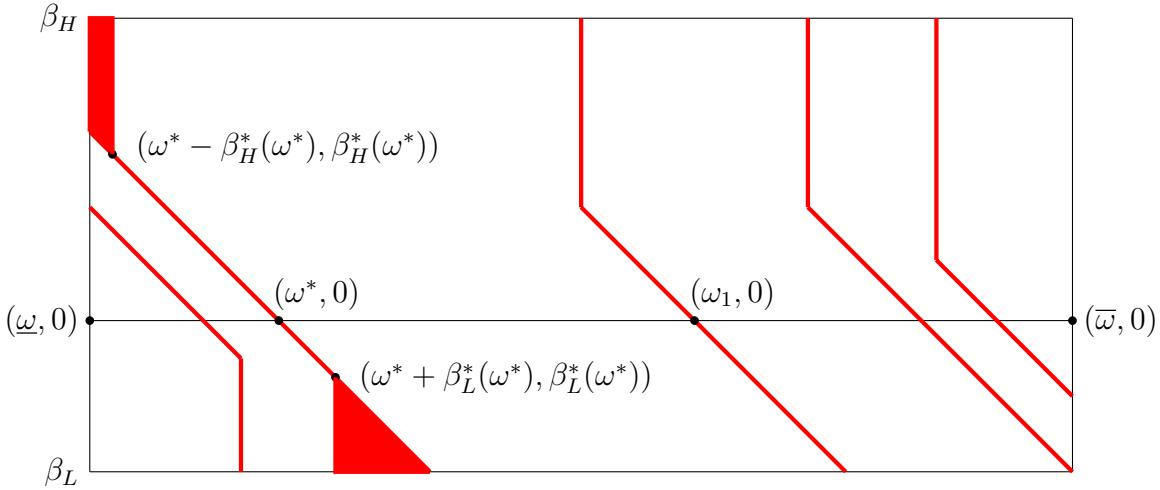


Figure 6: Bounded Uniform Model, $-\beta_L < \beta_H, \bar{\omega} - \underline{\omega} > -2\beta_L$.

This concludes the proof. ■

We now consider the case when Assumption 2 holds and either there are more expert types with a positive than a negative bias, or the state space is small. The equilibrium is depicted in Figure 6 for the case of a large state space. The main difference with the previous case is that now there is a positive measure of expert types that disclose the same interval in equilibrium and induce Dema to choose the same cut-off alternative $x = \omega^*$. More precisely, there is a positive measure of Knowyn types with a positive bias (and an observed state less than ω^*) and a positive measure of types with negative bias (and an observed state higher than ω^*) that induce Dema to choose ω^* . In Figure 6 this is represented by the two shaded areas. Types on the negatively sloped 45 degrees segment joining the two areas, which includes type $(\omega^*, 0)$, also induce Dema to choose ω^* . The two shaded areas are such that the expected state conditional on the expert types belonging to them (or to the zero measure 45 degrees segment joining them) is equal to ω^* .

To the right of the cut-off ω^* the equilibrium outcome is similar to that in Figure 3. The choice of any outcome $\omega > \omega^*$ is induced by types on the negatively sloped 45 degree segment going from type $(\omega - \beta_H^*(\omega), \beta_H^*(\omega))$ to type $(\min\{\omega - \beta_L, \bar{\omega}\}, \max\{\beta_L, \omega - \bar{\omega}\})$

and by types on the vertical segment going from $(\omega - \beta_H^*(\omega), \beta_H^*(\omega))$ to $(\omega - \beta_H^*(\omega), \beta_H)$.

Note that for $\omega \leq \bar{\omega} + \beta_L$ it is $\beta_H^*(\omega) = \beta^*$, as defined in (4). This says that for types to the right of ω^* and sufficiently far away from the right boundary $\omega = \bar{\omega}$ equilibrium takes the same form as in the case of an unbounded state space (see Figure 3).

The equilibrium map to the left of ω^* resembles the one of the equilibrium in Figure 4. Any outcomes ω to the left of the cut-off ω^* is induced by types on the negatively sloped 45 degree segment going from type $(\underline{\omega}, \omega - \underline{\omega})$ to type $(\omega - \beta_L^*(\omega), \beta_L^*(\omega))$ and by types on the vertical segment going from $(\omega - \beta_L^*(\omega), \beta_L^*(\omega))$ to $(\omega - \beta_L^*(\omega), \beta_L)$.

When Assumption 2 holds and the state space is small (i.e., $\bar{\omega} - \underline{\omega} < -2\beta_L$) equilibrium is similar to the one in Figure 6, except that now the shaded region of expert types with state greater than ω^* that induce Dema to choose ω^* extends to the right boundary with types $(\bar{\omega}, \beta)$ with $\beta \leq \beta_L^*(\omega^*)$.

As we pointed out, the cut-off value ω^* is defined so that it is equal to the expected value of the state in the shaded areas shown in Figure 6. Consider the 45 degree line that goes through type $(\omega^*, 0)$. The first shaded area is the one above the 45 degree line, starting from type $(\omega^* - \beta_H^*(\omega^*), \beta_H^*(\omega^*))$. The second is the area below the 45 degree line starting from type $(\omega^* - \beta_L^*(\omega^*), \beta_L^*(\omega^*))$. Formally, ω^* is implicitly defined by the equation $\psi(\omega) = 0$, where the function ψ is given by:

$$\psi(\omega) = \int_{\underline{\omega}}^{\max\{\underline{\omega}, \omega - \beta_H^*(\omega)\}} (x - \omega)(\beta_H - \omega + x) dx + \int_{\min\{\bar{\omega}, \omega - \beta_L^*(\omega)\}}^{\min\{\bar{\omega}, \omega - \beta_L\}} (x - \omega)(-\beta_L - x + \omega) dx.$$

The next lemma shows that ω^* exists and is uniquely defined.

Lemma 5 *In the bounded uniform model, under Assumption 2 there exists a unique state of the world ω^* such that $\psi(\omega^*) = 0$.*

We are now ready to present formally the equilibrium characterization.

Proposition 7 In the bounded uniform model, under Assumption 2, the equilibrium decision map is:

$$\alpha(\omega, \beta) = \begin{cases} \omega + \beta & \text{if } \begin{cases} \omega + \beta \geq \omega^* & \text{and } \beta \leq \beta_H^*(\omega) \\ \omega + \beta \leq \omega^* & \text{and } \beta \geq \beta_L^*(\omega) \end{cases} \\ \omega + \beta_H^*(\omega) & \text{if } \omega \geq \omega^* - \beta_H^*(\omega^*) \text{ and } \beta \geq \beta_H^*(\omega) \\ \omega - \beta_L^*(\omega) & \text{if } \omega \leq \omega^* - \beta_L^*(\omega^*) \text{ and } \beta \leq \beta_L^*(\omega) \\ \omega^* & \text{if } \begin{cases} \omega^* - \beta_H^*(\omega^*) \geq \omega \geq \omega^* - \beta \\ \omega^* - \beta \geq \omega \geq \omega^* - \beta_L^*(\omega^*) \end{cases} \end{cases}$$

Proof. For $x > \omega^*$, the set $T_\alpha(x)$ of Knowyn types inducing the choice x is the union of types $(x - \beta, \beta)$ with $\beta \in [\max\{\beta_L, x - \bar{\omega}\}, \beta_H^*(x)]$ and types $(x - \beta_H^*(x), \beta)$ with $\beta \in [\beta_H^*(x), \beta_H]$. Let the equilibrium strategy of types in the set $T_\alpha(x)$ be to disclose the interval $[x - \beta_H^*(x), x - \max\{\beta_L, x - \bar{\omega}\}]$. Setting $\gamma(x) = \min\{(\beta_L)^2, (x - \bar{\omega})^2\}$ and using (5), the expected value of the state of Knowyn's types in $T_\alpha(x)$ is:

$$\begin{aligned} & \mathbb{E}[\omega | (\omega, \beta) \in T_\alpha(x)] \\ &= \frac{(x - \beta_H^*(x))(\beta_H - \beta_H^*(x))}{\beta_H - \max\{\beta_L, x - \bar{\omega}\}} + \left(\frac{x - \beta_H^*(x) + x - \max\{\beta_L, x - \bar{\omega}\}}{2} \right) \cdot \frac{\beta_H^*(x) - \max\{\beta_L, x - \bar{\omega}\}}{\beta_H - \max\{\beta_L, x - \bar{\omega}\}} \\ &= x - \beta_H^*(x) + \frac{(\beta_H^*(x) - \max\{\beta_L, x - \bar{\omega}\})^2}{2(\beta_H - \max\{\beta_L, x - \bar{\omega}\})} \\ &= x - \beta_H^*(x) + \frac{\left(\beta_H^*(x) + \sqrt{\min\{(\beta_L)^2, (x - \bar{\omega})^2\}} \right)^2}{2 \left(\beta_H + \sqrt{\min\{(\beta_L)^2, (x - \bar{\omega})^2\}} \right)} \\ &= x - \left(\beta_H - \sqrt{(\beta_H)^2 - \gamma(x)} \right) + \frac{\left(\beta_H - \sqrt{(\beta_H)^2 - \gamma(x)} + \sqrt{\gamma(x)} \right)^2}{2 \left(\beta_H + \sqrt{\gamma(x)} \right)} \\ &= x - \beta_H + \frac{\left(\beta_H + \sqrt{\gamma(x)} \right)^2 + (\beta_H)^2 - \gamma(x)}{2 \left(\beta_H + \sqrt{\gamma(x)} \right)} \\ &= x \end{aligned}$$

Thus, after observing the disclosure of interval $[x - \beta_H^*(x), x - \max\{\beta_L, x - \underline{\omega}\}]$ with $x > \omega^*$, it is a best reply for Dema to choose alternative x .

For $x < \omega^*$, $T_\alpha(x)$ is the union of types $(x - \beta, \beta)$ with $\beta \in [\beta_L^*(x), x - \underline{\omega}]$ and types $(x - \beta_L^*(x), \beta)$ with $\beta \in [\beta_L, \beta_L^*(x)]$. Take the equilibrium strategy of types in $T_\alpha(x)$ to be disclosing $[\underline{\omega}, x - \beta_L^*(x)]$. Since $(\beta_L)^2 > (x - \underline{\omega})^2$, by (6) it is $\beta_L^*(x) = \beta_L + \sqrt{(\beta_L)^2 - (x - \underline{\omega})^2}$ and the expected value of the state of types in $T_\alpha(x)$ is:

$$\begin{aligned}
& \mathbb{E}[\omega | (\omega, \beta) \in T_\alpha(x)] \\
&= \frac{(x - \beta_L^*(x))(\beta_L^*(x) - \beta_L)}{x - \underline{\omega} - \beta_L} + \left(\frac{x - \beta_L^*(x) + \underline{\omega}}{2} \right) \cdot \frac{(x - \underline{\omega}) - \beta_L^*(x)}{x - \underline{\omega} - \beta_L} \\
&= x - \beta_L^*(x) - \frac{(x - \underline{\omega} - \beta_L^*(x))^2}{2(x - \underline{\omega} - \beta_L)} \\
&= x - \left(\beta_L + \sqrt{(\beta_L)^2 - (x - \underline{\omega})^2} \right) - \frac{\left(x - \underline{\omega} - \beta_L - \sqrt{(\beta_L)^2 - (x - \underline{\omega})^2} \right)^2}{2(x - \underline{\omega} - \beta_L)} \\
&= x - \beta_L - \frac{(x - \underline{\omega} - \beta_L)^2 + (\beta_L)^2 - (x - \underline{\omega})^2}{2(x - \underline{\omega} - \beta_L)} \\
&= x
\end{aligned}$$

Thus, it is also a best reply for Dema to choose x after observing the disclosure of $[\underline{\omega}, x - \beta_L^*(x)]$ with $x < \omega^*$.

Finally, $T_\alpha(\omega^*)$ is the union of types $(\omega^* - \beta, \beta)$ with $\beta \in [\beta_L^*(\omega^*), \beta_H^*(\omega^*)]$, types $(\omega^* - \beta, \beta')$ with $\beta' \geq \beta \in [\beta_H^*(\omega^*), \omega^* - \underline{\omega}]$, and types $(\omega^* - \beta, \beta')$ with $\beta' \leq \beta \in [\beta_L, \omega^* - \beta_L^*(\omega^*)]$. Let the equilibrium strategy of types in $T_\alpha(\omega^*)$ be to disclose $[\underline{\omega}, \omega^* - \beta_L]$. As proven in Lemma 5, ω^* is defined so that it is equal to the expected value of the state in the shaded areas shown in Figure 6, which are the set of Knowyn's types belonging to $T_\alpha(\omega^*)$. It is thus optimal for Dema to choose ω^* when the interval interval $[\underline{\omega}, \omega^* - \beta_L]$ is disclosed.

We now need to define the equilibrium strategy of Dema for all the off-the-equilibrium-path disclosures. After any such disclosure $[a, b]$, let Dema's equilibrium strategy be to choose b if $b < \omega^*$, to choose ω^* if $a \leq \omega^* \leq b$, and to choose a if

$\omega^* < a$. This strategy is optimal for some off the equilibrium path beliefs by Dema (e.g., if she believes that the state is equal to her choice with probability one).

It now only remains to prove that the equilibrium strategy of Knowyn is a best reply to Dema's strategy. First consider the types that induce a choice $x > \omega^*$. As is clear from Figure 6 these types either induce their ideal choice or would like to induce a higher choice, but if the types who would like a higher choice disclose any feasible (i.e., including the true state) interval different from their equilibrium interval, then Dema will choose a lower alternative. Hence these types have no profitable deviation. Similarly, consider the types that induce a choice $x < \omega^*$. These types either get their ideal choice or would like a lower choice, but by disclosing any feasible (i.e., including the true state) interval different from their equilibrium interval, they induce Dema to choose a higher alternative. Thus, these types also have no profitable deviation. Finally, consider the types that induce the choice of ω^* . Again, given Dema's strategy, any feasible deviation by types in this set that would like a higher choice can only induce a lower choice and any feasible deviation by types who would like a lower choice can only induce a higher choice. This completes the proof. ■

The qualitative features of the equilibrium of the bounded uniform model described in Proposition 7 remain valid in the model with a general type distribution and bounded state space, but there might be multiple regions with pooling by a positive measure of types. These positive measure regions separate two regions of the type space where the set of expert types that pool and disclose the same interval is a zero measure set. Both of these zero measure regions have a a 45 degree segment and a vertical segment; in one region the vertical segment includes types with an extreme negative bias, in the other it includes types with an extreme positive bias. With general type distributions, there is an issue of existence of a pure strategy equilibrium, but under mild regularity conditions existence is guaranteed.

5 Mean Bias and Variance

In the uniform model, the mean and variance of the state and the expert's bias fully describe the information setting. In this section, we first study the effect of changes in the mean and variance of the expert's bias. We then consider changes in the mean and variance of the state.

Intuitively, keeping the variance in Knowyn's bias fixed and maintaining the assumption that her mean bias is non negative (i.e., $\beta_H \geq -\beta_L$), an increase in mean bias makes Knowyn's preferences more predictable and should then hurt her while benefiting Dema. The effect of an increase in the bias variance keeping the mean constant is less intuitively clear. As we shall see, it hurts Dema as it expands the average distance between Dema's ideal choice and the chosen alternative, but has no effect on Knowyn. To formally derive these results in the simplest way, we consider the unbounded uniform model. We maintain the assumption that $\beta_H \geq -\beta_L \geq 0$.

Using Proposition 5, it is convenient to work with the equilibrium loss functions of Dema and Knowyn for a given type (ω, β) , defined by:

$$L_D(\omega, \beta) = \sqrt{-V(\alpha(\omega, \beta), \omega)} = \begin{cases} \sqrt{\beta^2} & \text{if } \beta \leq \beta_H^* \\ \beta_H^* & \text{if } \beta \geq \beta_H^* \end{cases}$$

$$L_K(\omega, \beta) = \sqrt{-U(\alpha(\omega, \beta), \omega, \beta)} = \begin{cases} 0 & \text{if } \beta \leq \beta_H^* \\ \beta - \beta_H^* & \text{if } \beta \geq \beta_H^* \end{cases}$$

Since the agents' losses do not depend on the state ω , we might write the expected losses for each state as follows:

$$\begin{aligned} \mathbb{E}[L_D] &= - \int_{\beta_L}^0 \beta \frac{d\beta}{\beta_H - \beta_L} + \int_0^{\beta_H^*} \beta \frac{d\beta}{\beta_H - \beta_L} + \int_{\beta_H^*}^{\beta_H} \beta_H^* \frac{d\beta}{\beta_H - \beta_L} \\ &= \frac{(\beta_L)^2}{\beta_H - \beta_L}; \end{aligned}$$

$$\begin{aligned}\mathbb{E}[L_K] &= \int_{\beta_H^*}^{\beta_H} (\beta - \beta_H^*) \frac{d\beta}{\beta_H - \beta_L} \\ &= \frac{\beta_H + \beta_L}{2},\end{aligned}$$

where the last equalities in both expressions follow from the definition of β_H^* in (4).

Since Knowyn's bias is uniformly distributed in the interval $[\beta_L, \beta_H]$, it is immediate that the mean bias is $\mathbb{E}[\beta] = \frac{\beta_H + \beta_L}{2}$ and the standard deviation in the bias is $\sigma_\beta = \frac{\beta_H - \beta_L}{2\sqrt{3}}$. Thus, the expected losses can be written as:⁶

$$\begin{aligned}\mathbb{E}[L_D] &= \frac{(\sqrt{3}\sigma_\beta - \mathbb{E}[\beta])^2}{2\sqrt{3}\sigma_\beta} \\ \mathbb{E}[L_K] &= \mathbb{E}[\beta]\end{aligned}$$

Parts (i) and (ii) of the next proposition are easily proven by differentiating the agents' expected losses.

Proposition 8 *In the unbounded uniform model with $\beta_H \geq -\beta_L \geq 0$:*

- (i) *An increase in the mean bias of the expert, keeping the standard deviation constant, decreases Dema's expected loss and increases Knowyn expected loss;*
- (ii) *An increase in the standard deviation of the expert's bias, keeping the mean constant, increases Dema's expected loss and has no effect on Knowyn expected loss.*
- (iii) *The measure of expert types that fully manipulate the outcome decreases with the expert's mean bias and increases with the standard deviation of the bias.*

To see that (iii) holds, note that the measure of expert types that fully manipulate and induce choice $\alpha(\omega, \beta) = \omega + \beta$ is:

$$\begin{aligned}\frac{\beta_H^* - \beta_L}{\beta_H - \beta_L} &= 1 - \frac{\sqrt{(\beta_H)^2 - (\beta_L)^2}}{\beta_H - \beta_L} \\ &= 1 - \sqrt{\frac{\mathbb{E}[\beta]}{\sqrt{3}\sigma_\beta}}\end{aligned}$$

⁶Note that, since $\beta_L \leq 0$ and $\beta_H \leq 2\sqrt{3}\sigma_\beta$, it is $\sqrt{3}\sigma_\beta \geq \mathbb{E}[\beta]$.

The intuition for the effect of the mean bias is clear. For a fixed support size $\beta_H - \beta_L$ in the bias, or equivalently a fixed standard deviation, if the mean bias is sufficiently large, then the direction of the bias is known (because $\beta_L = 0$ and $\beta_H > 0$) and full disclosure is the equilibrium outcome. On the other hand, if the mean bias is zero we are in the case described in Figure 1, (i.e., $-\beta_L = \beta_H$) and all expert's types fully manipulate the outcome. The proposition shows that as the mean bias increases we move monotonically from one extreme to the other.

We now consider changes in the distribution of the state of the world ω . A change in the mean of the state with a constant variance clearly has no substantive effect on the equilibrium outcome and the agents' payoff. On the other hand, since the expert's ability to manipulate the decision maker's choice is smaller near the boundaries, an increase in the variance of the state with constant mean benefits Knowyn at the expense of Dema, as it expands the middle region relative to the boundary regions. To derive this result formally in the simplest way, we consider the uniform model with zero expected bias and a large state space (i.e., we postulate that Assumption 1 holds).

Proposition 9 *In the uniform, bounded model under Assumption 1, an increase in the standard deviation of the state increases the measure of expert types that fully manipulate, decreases the expected loss of the expert, and increases the expected loss of the decision maker.*

Proof. Since the state space is large, i.e., $\bar{\omega} - \underline{\omega} > 2\beta_H = -2\beta_L$, the losses of Knowyn and Dema conditional on being either in the left or in the right boundary regions are independent from the standard deviation of the state $\sigma_\omega = \frac{\bar{\omega} - \underline{\omega}}{2\sqrt{3}}$. The losses conditional on the expert ideal choice being x (i.e., conditional on $(\omega, \beta) \in T_\alpha(x)$), with $x = \omega + \beta$ in the middle region of the type space, are also the same; Knowyn's loss is zero and Dema's loss is higher than the loss conditional on any set $T_\alpha(x)$ of Knowyn's types belonging to the two boundary regions. Since an increase in the

standard deviation of the state increases the probability that the expert’s ideal choice is in the middle region, the proposition follows. ■

Propositions 8 and 9 make it clear that experts should strive to look ex-ante unbiased, while decision makers should look for highly biased experts. This is because disclosures from experts with a large mean bias are easier to read; their preferences and incentives are more transparent than the preferences of experts with equal bias variance but a more uncertain direction of their bias. The propositions also show that the expert is in a better position to manipulate the outcome when the decision maker is more uncertain about the optimal choice (e.g., because she is less familiar with the problem).

6 Related Literature

The literature on disclosure of verifiable information is large and has many disparate applications; see Milgrom (2008) and Dranove and Jin (2010) for surveys. The seminal papers are Grossman (1981), Grossman and Hart (1980) and Milgrom (1981). They assume that the bias of the expert is common knowledge. Their main insight is that the decision maker ought to be highly skeptical. When evidence is incomplete, the decision maker assumes that any missing information is likely to be unfavourable to the expert. As a consequence, experts will want to reveal all favourable information. This unravels any attempt to hide information and leads to full disclosure.

A few papers have introduced some uncertainty in the expert’s preferences, but none have gone as far as this paper in demonstrating the potential for full manipulation by the expert when the bias direction is unknown. Seidmann and Winter (1997) allow uncertainty over the expert’s preferences, but they focus on the conditions under which there is an equilibrium with full disclosure. Giovannoni and Seidmann (2007) focus on the totally “pooling” equilibrium with no information transmission (i.e., no disclosure); such an equilibrium exists if and only if all expert types prefer the unin-

formed decision maker ideal action over the ideal action of a fully informed decision maker (a condition that cannot hold in our model with preferences à la Crawford and Sobel, 1982). In Dye (1985) and Shin (1994), full disclosure may fail because of uncertainty over whether the expert knows the true state. In such a case the decision maker's skepticism is tempered; evidence may be incomplete because the expert does not have all the information. In Wolinsky (2003), the expert's bias is unknown, but it can only take two values. The expert can fully report, or underreport favorable information. In equilibrium, the biased expert with favorable information above a threshold fully discloses, while all other expert types play a mixed strategy. In Dziuda (2011), it is uncertain whether the expert is honest or biased in favor of one of two alternatives. Her focus is showing that a biased expert may also disclose unfavorable information. In the dynamic setting of Acharya et al. (2011) and Guttman et al. (2014), full disclosure at each point in time fails because the decision maker (investors) does not know when the expert (a firm) has acquired information. Acharya et al. (2011) show that good market news slows, and bad news triggers, the release of information by firms. Guttman et al. (2014) show that later disclosures are interpreted more favorably by the market. In a different vein, Verrecchia (1983) shows that full disclosure may also fail if the expert must pay a cost to disclose information.

For a survey of the vast literature on cheap talk, see Sobel (2013). Here we should first mention three papers in which the preferences of the privately informed party (the sender) are not perfectly known by the receiver. Morgan and Stocken (2003) consider a cheap talk model in which the expert might be biased or unbiased. They show that bias uncertainty may increase information transmission compared to the case where the expert has a known but intermediate bias. Dimitrakas and Sarafidis (2005) show that revelation of the expert's bias may diminish information transmission when the bias size is uncertain, and Li and Madarasz (2008) show that revelation always decreases information transmission and the welfare of both parties when the direction of the bias is uncertain. This is very different from what happens

in our model in which verifiable disclosures are available; in our model, revealing the bias helps the decision maker and hurts the expert. More importantly from our point of view, these papers show that even if the bias of the sender is uncertain, when verifiable disclosures are not feasible and only cheap talk is available, equilibrium is partitional as in the standard cheap talk model with known bias. In the setting of our model with a bounded state space this means that one could find cheap talk messages m_1, m_2, \dots, m_n with $n \geq 1$ such that all Knowyn types with $x_{j-1} \leq \underline{\omega} + \beta \leq x_j$ send message m_j , with $1 \leq j \leq n$, $x_0 = \underline{\omega} + \beta_L$ and $x_n = \bar{\omega} + \beta_H$. Unlike in our paper in which verifiable disclosure are possible, in such an equilibrium only a zero measure of expert types is able to fully manipulate the outcome.

We should finally mention Chakraborty and Harbaugh (2010) and Kartik and Van Weelden (2019). Chakraborty and Harbaugh (2010) studied cheap talk when the sender, as in the persuasion literature, has state-independent preferences (i.e., an ideal choice). Contrary to the case of a single dimensional state-space in which information transmission would not be possible, they show that with a multidimensional state space a partially informative equilibrium may exists; the sender can communicate some information and influence the receiver's choice by trading off information along the dimensions of the state space. Kartik and Van Weelden (2019) studied a model in which the electorate is uncertain about the preferences of a politician (the sender) and showed that the politician may send informative cheap talk messages before the election, and in particular espouse views not aligned with the electorate's preferences. The reason is that voters may prefer a politician with known biased views to one whose bias is sufficiently uncertain. This resembles what happens in our model, where the decision maker always prefers an expert with known bias to one with uncertain bias direction.

7 Conclusions

We have introduced uncertainty in the direction of an expert’s bias in a model in which both verifiable information and cheap talk messages may be communicated by the expert, and shown that information unravelling fails. Full disclosure is not an equilibrium. Experts with positive bias observing a low state of the world pool with experts with a negative bias observing a high state. Manipulation is pervasive and a positive measure of expert types are able to get what they want (i.e., their ideal choice). An increase in the familiarity with the problem helps the decision maker. More interestingly, the size of an expert’s bias is less important to a decision maker than knowledge of the bias direction. Experts that are known to be strongly biased in one direction can be easily read and their disclosures decoded. Thus, experts should strive to be poker faced. They should try to conceal which way they would like to push the outcome.

A number of promising extensions of the basic idea in this paper are worth pursuing in future research. First, the value of keeping a poker face may shape the composition of expert partnerships and explain, for example, the value of diversity. If partners come from diverse backgrounds and experiences, it will naturally be more difficult for a client to discern the direction of the organization’s bias. As a suggestive example, it is broadly consistent with the theory of this paper that in 2003 the first Bush administration sent Colin Powell, and not Dick Cheney or Donald Rumsfeld, to present to the United Nations Security Council evidence about the existence of weapons of mass destruction in Saddam’s Iraq. As opposed to the general view of Cheney and Rumsfeld as hawks, Powell had a reputation for being a moderate, and hence there was more uncertainty about his willingness to selectively disclose information to justify the Iraq war. His presenting had more “manipulative power”!

Second, when the decision maker is an active agent (and not the market, or a mass of customers), it is often possible to seek a second option. Intuitively, a second opinion ought to be valuable, even with *a priori* identical experts, because there is

always a chance the new expert has a different bias. Indeed, if it is known that the decision maker will seek a second opinion, the first expert will want to change her disclosure policy. More careful analysis is needed to understand what would happen. For example, it is not clear whether the decision maker should consult two different experts simultaneously or sequentially. It is also not clear whether the decision maker should disclose the information revealed by an expert she has previously visited.

Third, manipulation and in particular market manipulation, often involves taking actions, as opposed to disclosing information. For example, transactions which create an artificial security price are regarded as market manipulation by the Securities Exchange Act of 1934. In general, taking actions involves both elements of pure information disclosure and elements of costly signaling. The difference between the two hinges on what is observable to the market and on the cost of divulging the information. The basic insight of this paper ought to go through, however. If the market is not certain whether the manipulator is trying to push price up or down, market manipulation is likely to have a good chance of succeeding, at least partially.

Finally, it is important to study ways for decision makers to improve reliability of the disclosure process. In this regard, it is instructive to look at the measures taken by the editors of several top medical journals. Because data manipulation, or “fudging the data”, is thought to be common, they have decided to stop publishing drug research sponsored by pharmaceutical companies unless the research was registered from its beginning in a public database.⁷ In essence, this is a way to put constraints on the disclosures available to the companies.

⁷Note that while it is perfectly sensible to be skeptical of sponsored research, it is not obvious that the individual researcher should be always assumed to have no integrity. Thus, it is reasonable to say that the exact direction of the bias is at least to some extent unknown.

References

- [1] Acharya, Viral V., DeMarzo, Peter and Ilan Kremer. 2011. “Endogenous Information Flows and the Clustering of Announcements.” *American Economic Review* 101(7): 2955–2979.
- [2] Chakraborty, Archishman and Rick Harbaugh 2010. “Persuasion by Cheap Talk.” *American Economic Review* 100 (5): 2361–2382.
- [3] Crawford, Vincent P., and Joel Sobel. 1982. “Strategic Information Transmission.” *Econometrica* 50(6): 1431–1451.
- [4] Dimitrakas, Vassilios and Yianis Sarafidis. 2005. “Advice from an Expert with Unknown Motives.” <http://dx.doi.org/10.2139/ssrn.797844>
- [5] Dranove, David and Ginger Zhe Jin. 2010. “Quality Disclosure and Certification: Theory and Practice.” *Journal of Economic Literature* 48(4): 935–963.
- [6] Dye, Ronald A. 1985. “Disclosure of Nonproprietary Information.” *Journal of Accounting Research* 23(1): 123–145.
- [7] Dziuda, Wioletta. 2011. “Strategic Argumentation.” *Journal of Economic Theory* 146 (4): 1362–1397.
- [8] Giovannoni, Francesco and Daniel J. Seidmann. 2007. “Secrecy, Two-sided Bias and the Value of Evidence” *Games and Economic Behavior* 59: 296 – 315.
- [9] Grossman, Sanford J. 1981. “The Informational Role of Warranties and Private Disclosure about Product Quality.” *Journal of Law and Economics*, 24(3): 461–83.
- [10] Grossman, Sanford J. and Oliver D. Hart. 1980. “Disclosure Laws and Takeover Bids.” *The Journal of Finance* 35(2): 323–334.

- [11] Guttman, Ilan, Kremer, Ilan and Andrzej Skrzypacz. 2014. “Not Only What but Also When: A Theory of Dynamic Voluntary Disclosure.” *American Economic Review* 104(8), 2400–2420.
- [12] Kamenica, Emir, and Matthew Gentzkow. 2011. “Bayesian Persuasion.” *American Economic Review*, 101(6): 2590–2615.
- [13] Kartik, Navin, and Richard Van Weelden. 2019 “Informative Cheap Talk in Elections.” *Review of Economic Studies*, 86(2): 755–784.
- [14] Li, Ming, and Kristóf Madarász. 2008. “When Mandatory Disclosure Hurts: Expert Advice and Conflicting Interests.” *Journal of Economic Theory*, 139(1): 47–74.
- [15] Milgrom, Paul. 1981. “Good News and Bad News: Representation Theorems and Applications.” *Bell Journal of Economics*, 12(2): 380–91.
- [16] Milgrom, Paul. 2008. “What the Seller Won’t Tell You: Persuasion and Disclosure in Markets.” *Journal of Economic Perspectives* 22(2): 115–131.
- [17] Morgan, John, and Phillip C. Stocken. 2003. “An Analysis of Stock Recommendations.” *RAND Journal of Economics* 34(1): 183–203.
- [18] Rayo, Luis, and Ilya Segal. 2010. “Optimal Information Disclosure.” *Journal of Political Economy*, 118 (5): 949–987.
- [19] Seidmann, Daniel and Eyal Winter. 1997. “Strategic Information with Verifiable Messages.” *Econometrica*, 65: 163–169.
- [20] Sobel, Joel. 2013. “Giving and Receiving Advice.” In *Advances in Economics and Econometrics*, D. Acemoglu, M. Arellano and E. Dekel (eds.), Cambridge, Cambridge University Press.

- [21] Shin, Hyun Song. 1994. “News Management and the Value of Firms.” *Rand Journal of Economics*, 25(1): 58–71.
- [22] Verrecchia, Robert. 1983. “Discretionary Disclosure.” *Journal of Accounting and Economics* 5: 179–194.
- [23] Wolinsky, Asher. 2003. “Information Transmission when the Sender’s Preferences are Uncertain.” *Games and Economic Behavior* 42(2): 319–326.

Appendix

Proof of Lemma 2.

(i) Expert type (ω, β') can mimic type (ω, β'') by disclosing the same interval. Hence, it must be:

$$\begin{aligned} -[\alpha(\omega, \beta') - \omega - \beta']^2 &\geq -[\alpha(\omega, \beta'') - \omega - \beta']^2, \quad \text{or} \\ 2(\omega + \beta')\alpha(\omega, \beta') - \alpha(\omega, \beta')^2 &\geq 2(\omega + \beta')\alpha(\omega, \beta'') - \alpha(\omega, \beta'')^2, \quad \text{or} \\ 2(\omega + \beta')[\alpha(\omega, \beta') - \alpha(\omega, \beta'')] &\geq \alpha(\omega, \beta')^2 - \alpha(\omega, \beta'')^2. \end{aligned} \tag{7}$$

Similarly, since type (ω, β'') can mimic type (ω, β') :

$$2(\omega + \beta'')[\alpha(\omega, \beta'') - \alpha(\omega, \beta')] \geq \alpha(\omega, \beta'')^2 - \alpha(\omega, \beta')^2 \tag{8}$$

and hence, adding up each side of the two inequalities:

$$(\beta' - \beta'')[\alpha(\omega, \beta') - \alpha(\omega, \beta'')] \geq 0.$$

This shows that $\alpha(\omega, \beta)$ is weakly increasing in β and hence continuous almost everywhere.

(ii) Take $\beta' < \beta''$ and $\beta \in [\beta', \beta'']$. By (7) and (8) we have

$$2(\omega + \beta'')[\alpha(\omega, \beta'') - \alpha(\omega, \beta')] \geq \alpha(\omega, \beta'')^2 - \alpha(\omega, \beta')^2 \geq 2(\omega + \beta')[\alpha(\omega, \beta'') - \alpha(\omega, \beta')]. \tag{9}$$

Note that the inequalities in (9) hold as equalities if $\alpha(\omega, \beta') = \alpha(\omega, \beta'')$. Suppose instead $\alpha(\omega, \beta') \neq \alpha(\omega, \beta'')$. By (i) it is $\alpha(\omega, \beta') < \alpha(\omega, \beta'')$, and we can write (9) as

$$2(\omega + \beta'') \geq \alpha(\omega, \beta') + \alpha(\omega, \beta'') \geq 2(\omega + \beta').$$

Taking limits as β' and β'' converge to β , we obtain $\alpha(\omega, \beta) = \omega + \beta$ at all points where α is continuous. This also implies that there cannot be any jump discontinuity and hence α must be continuous in β . ■

Proof of Lemma 3.

By Lemma 2, either $\alpha(\omega, \beta)$ is constant or is equal to $\omega + \beta$. Suppose $\alpha(\omega, \beta)$ is constant in an interval $[\beta', \beta'']$ with $\beta_L < \beta' < \beta'' < \beta_H$, while it is increasing for values of β in the intervals $(\beta' - \varepsilon, \beta')$ and $(\beta'', \beta'' + \varepsilon)$. By Lemma 2, it must be $\alpha(\omega, \beta') = \omega + \beta' < \omega + \beta'' = \alpha(\omega, \beta'')$. This contradicts the assumption that $\alpha(\omega, \beta)$ is constant in the interval $[\beta', \beta'']$. Thus, it can only be constant in an interval including one of the endpoints of the bias support. ■

Proof of Lemma 4.

For $\beta = 0$ the statement is true, since $\alpha(\omega, 0) = \omega$ for all ω by Lemma 1. Fix $\beta > 0$. First, if $\alpha(\omega, \beta) = \omega$, then $\alpha(\omega, \beta)$ is strictly increasing to the right of ω , since it must be $\alpha(\omega', \beta) \geq \omega'$ for all $\omega' > \omega$, as type (ω', β) could induce the choice of ω' by reporting the true state. Second, suppose $\alpha(\omega, \beta) > \omega$. Then, type (ω, β) must disclose an interval $[a, b]$ with $a \leq \omega$ and $b \geq \alpha(\omega, \beta)$. Since type (ω', β) with $\omega < \omega' < b$ could mimic (disclose the same interval as) type (ω, β) , it must be

$$\begin{aligned} [\omega' + \beta - \alpha(\omega', \beta)]^2 &\leq [\omega' + \beta - \alpha(\omega, \beta)]^2, \text{ or} \\ 2(\omega' + \beta)[\alpha(\omega, \beta) - \alpha(\omega', \beta)] &\leq [\alpha(\omega, \beta) + \alpha(\omega', \beta)][\alpha(\omega, \beta) - \alpha(\omega', \beta)], \text{ or} \\ \alpha(\omega, \beta) &\leq \alpha(\omega', \beta), \end{aligned}$$

where the third inequality holds because by $\beta > 0$ and Lemma 3, $\omega' + \beta > \omega + \beta \geq \alpha(\omega, \beta)$ and $\omega' + \beta \geq \alpha(\omega', \beta)$, which implies, adding up the inequalities, $2(\omega' + \beta) > [\alpha(\omega, \beta) + \alpha(\omega', \beta)]$. This shows that for $\beta > 0$ the map $\alpha(\omega, \beta)$ is increasing in ω . The proof for $\beta < 0$ is analogous. ■

Proof of Lemma 5.

First observe that

$$\psi(\underline{\omega}) = \int_{\underline{\omega}}^{\min\{\bar{\omega}, \omega - \beta_L\}} (x - \underline{\omega})(-\beta_L - x + \underline{\omega}) dx > 0.$$

Then observe that

$$\psi(\bar{\omega}) = \int_{\underline{\omega}}^{\bar{\omega}} (x - \bar{\omega})(\beta_H - \bar{\omega} + x) dx < 0.$$

Since ψ is continuous in ω , this already proves that there exists ω such that $\psi(\omega) = 0$. To prove uniqueness and complete the proof, we will show that $\psi(\omega)$ is decreasing in ω for all states ω such that $\psi(\omega) = 0$. We must distinguish between two main cases, each with different sub-cases. In Case A, it is $\bar{\omega} - \underline{\omega} > -2\beta_L$; in this case the size of the state space is large relative to the size of highest negative bias. In Case B, it is $\bar{\omega} - \underline{\omega} < -2\beta_L$. We begin by considering the three sub-cases of Case A.

Case A1: $\omega < \underline{\omega} - \beta_L$. In this case, it is $\beta_H^*(\omega) = \beta_H - \sqrt{(\beta_H)^2 - (\beta_L)^2} = \beta_H^*$, $\beta_L^*(\omega) = \beta_L + \sqrt{(\beta_L)^2 - (\omega - \underline{\omega})^2}$ and:

$$\psi(\omega) = \int_{\underline{\omega}}^{\max\{\underline{\omega}, \omega - \beta_H^*\}} (x - \omega)(\beta_H - \omega + x) dx + \int_{\omega - \beta_L^*(\omega)}^{\omega - \beta_L} (x - \omega)(-\beta_L - x + \omega) dx.$$

First note that if $\omega < \underline{\omega} + \beta_H^*$, then the first term on the rhs is a constant equal to zero. If $\omega > \underline{\omega} + \beta_H^*$, by setting $y = x - \omega$ and changing variable of integration the first terms becomes:

$$\int_{\underline{\omega} - \omega}^{-\beta_H^*} y(\beta_H + y) dy$$

and is decreasing in ω since $\beta_H > -\beta_L$. Now consider the second term; again change the variable of integration to be $y = x - \omega$. The term becomes:

$$\int_{-\beta_L^*(\omega)}^{-\beta_L} y(-\beta_L - y) dy$$

and it is also decreasing in ω . This completes the proof that $\psi(\omega)$ is a decreasing function in Case A1.

Case A2: $\underline{\omega} - \beta_L < \omega < \bar{\omega} + \beta_L$. In this case, it is $\beta_H^*(\omega) = \beta_H - \sqrt{(\beta_H)^2 - (\beta_L)^2} = \beta_H^*$, which implies that $\omega - \beta_H^* > \underline{\omega} - \beta_L - \beta_H^* > \underline{\omega}$. It is also, $\beta_L^*(\omega) = \beta_L$ and hence:

$$\psi(\omega) = \int_{\underline{\omega}}^{\omega - \beta_H^*} (x - \omega)(\beta_H - \omega + x) dx = \int_{\underline{\omega} - \omega}^{-\beta_H^*} y(\beta_H + y) dy$$

Differentiating with respect to ω we obtain:

$$\psi'(\omega) = (\omega - \underline{\omega})(\omega - \underline{\omega} - \beta_H) < (\omega - \underline{\omega})(\underline{\omega} - \beta_L - \underline{\omega} - \beta_H) < 0$$

where the last inequality follows from $\beta_H > -\beta_L$, by Assumption 2. Hence, in Case A2 $\psi(\omega)$ is also decreasing in ω .

Case A3: $\bar{\omega} + \beta_L < \omega$. In this case, it is $\beta_H^*(\omega) = \beta_H - \sqrt{(\beta_H)^2 - (\bar{\omega} - \omega)^2}$, $\beta_L^*(\omega) = \beta_L$ and:

$$\psi(\omega) = \int_{\underline{\omega}}^{\omega - \beta_H^*(\omega)} (x - \omega)(\beta_H - \omega + x) dx = \int_{\underline{\omega} - \omega}^{-\beta_H^*(\omega)} y(\beta_H + y) dy < 0$$

where the inequality follows from the integrand being negative (as $y < 0$ and $\beta_H > \beta_H^*(\omega)$) for all values of y . This shows that in Case A3 there does not exist an ω such that $\psi(\omega) = 0$.

We can then conclude that if we are in Case A (i.e., it is $\bar{\omega} - \underline{\omega} > -2\beta_L$), then $\psi(\omega)$ is decreasing at all ω such that $\psi(\omega) = 0$ and hence there is a unique ω^* .

We now consider Case B, when it is $\bar{\omega} - \underline{\omega} < -2\beta_L$. There three sub-cases.

Case B1: $\omega < \bar{\omega} + \beta_L$. In this case, it is $\beta_H^*(\omega) = \beta_H - \sqrt{(\beta_H)^2 - (\beta_L)^2} = \beta_H^*$, $\beta_L^*(\omega) = \beta_L + \sqrt{(\beta_L)^2 - (\omega - \underline{\omega})^2}$ and $\psi(\omega)$ is the same as in Case A1. The proof that it is a decreasing function is also the same.

Case B2: $\bar{\omega} + \beta_L < \omega < \omega - \beta_L$. In this case, it is $\beta_H^*(\omega) = \beta_H - \sqrt{(\beta_H)^2 - (\bar{\omega} - \omega)^2}$, $\beta_L^*(\omega) = \beta_L + \sqrt{(\beta_L)^2 - (\omega - \underline{\omega})^2}$. Since it is $\bar{\omega} < \omega - \beta_L$:

$$\psi(\omega) = \int_{\underline{\omega}}^{\max\{\underline{\omega}, \omega - \beta_H^*(\omega)\}} (x - \omega)(\beta_H - \omega + x) dx + \int_{\min\{\bar{\omega}, \omega - \beta_L^*(\omega)\}}^{\bar{\omega}} (x - \omega)(-\beta_L - x + \omega) dx.$$

If $\omega < \underline{\omega} + \beta_H^*(\omega)$, then the first term on the rhs is a constant equal to zero. If $\omega > \underline{\omega} + \beta_H^*(\omega)$, by setting $y = x - \omega$ and changing variable of integration the first terms becomes:

$$\int_{\underline{\omega} - \omega}^{-\beta_H^*(\omega)} y(\beta_H + y) dy$$

and differentiating with respect to ω it can be seen to be decreasing in ω . Now consider the second term; if $\omega > \bar{\omega} + \beta_L^*(\omega)$, then it is equal to zero. If $\omega < \bar{\omega} + \beta_L^*(\omega)$, by setting $y = x - \omega$ and changing variable of integration the second terms becomes:

$$\int_{-\beta_L^*(\omega)}^{\bar{\omega} - \omega} y(-\beta_L - y) dy$$

Differentiating with respect to ω we obtain:

$$\psi'(\omega) = (\bar{\omega} - \omega)(\beta_L + \bar{\omega} - \omega) + \frac{\partial \beta_L^*(\omega)}{\partial \omega} \beta_L^*(\omega)(\beta_L - \beta_L^*(\omega)) < 0.$$

Hence, to conclude that in Case B2 $\psi(\omega)$ is a strictly decreasing function, it only remains to show that it cannot be the case that both first and second term on the rhs of $\psi(\omega)$ are zero. Suppose to the contrary that both terms are zero. For the first term to be zero it must be:

$$\omega \leq \underline{\omega} + \beta_H^*(\omega) = \omega + \beta_H - \sqrt{(\beta_H)^2 - (\bar{\omega} - \omega)^2} \quad (10)$$

and since the rhs of (10) is decreasing in β_H , a necessary condition for the inequality to hold is that it holds at $\beta_H = -\beta_L$, the smallest value β_H can take. For the second term on the rhs of $\psi(\omega)$ to be zero it must be:

$$\omega \geq \bar{\omega} + \beta_L^*(\omega) = \bar{\omega} + \beta_L + \sqrt{(\beta_L)^2 - (\omega - \underline{\omega})^2} \quad (11)$$

and since the rhs is decreasing in β_L , a necessary condition for the inequality to hold is that it holds at $\beta_L = -\frac{\bar{\omega} - \underline{\omega}}{2}$, the highest value that β_L can take in this case. Thus, replacing $\beta_H = -\beta_L$ and $\beta_L = -\frac{\bar{\omega} - \underline{\omega}}{2}$ in (10) and (11), a necessary condition for both terms on the rhs of $\psi(\omega)$ to be zero is:

$$\frac{\bar{\omega} + \underline{\omega}}{2} + \sqrt{\left(\frac{\bar{\omega} - \underline{\omega}}{2}\right)^2 - (\omega - \underline{\omega})^2} \leq \omega \leq \frac{\bar{\omega} + \underline{\omega}}{2} - \sqrt{\left(\frac{\bar{\omega} - \underline{\omega}}{2}\right)^2 - (\bar{\omega} - \omega)^2}$$

which cannot hold.

Case B3: $\underline{\omega} - \beta_L < \omega$, which implies that $\omega - \beta_H^*(\omega) > \underline{\omega} - \beta_L - \beta_H^*(\omega) > \underline{\omega}$. In this case, it is $\beta_H^*(\omega) = \beta_H - \sqrt{(\beta_H)^2 - (\bar{\omega} - \omega)^2}$. It is also $\beta_L^*(\omega) = \beta_L$ and:

$$\psi(\omega) = \int_{\underline{\omega}}^{\omega - \beta_H^*(\omega)} (x - \omega)(\beta_H - \omega + x) dx = \int_{\underline{\omega} - \omega}^{-\beta_H^*(\omega)} y(\beta_H + y) dy < 0$$

where the inequality follows from the same argument used in the study of Case A3. Thus in Case B3 there does not exist an ω such that $\psi(\omega) = 0$.

We have thus shown that in all sub-cases of Cases A and B the function $\psi(\omega)$ is decreasing at all ω such that $\psi(\omega) = 0$ and hence there is a unique ω^* . This concludes the proof of the lemma. ■