

Thesis: Essays in Information Transmission

Jesper Rudiger¹

Supervised by Antonio Cabrales
Universidad Carlos III de Madrid

¹E-mail: jrudiger@eco.uc3m.es.

Acknowledgements

First and foremost I would like to thank my advisor Antonio Cabrales for his help, for always being available and for always being constructive, whether that involved healthy skepticism in moments of exuberance or encouragement in moments of adversity. Apart from my advisor, I have in particular benefitted from the comments of Ángel Hernando-Veciana, Praveen Kujal, Diego Moreno and José Penalva, who have most diligently read earlier versions of the papers. Navin Kartik and Andrew Schotter have made comments to specific papers. Nicolas Vieille kindly hosted me a semester at HEC Paris and helped me develop more general aspects of my project. I have been fortunate to study with people such as Johannes Mueller-Trede, Robert Blotevogel, Dan Nixon and Adrien Vigier who inspired a critical view of “the dismal science”. At Universidad Carlos III de Madrid I have enjoyed many comments from other students, in particular, Daniel García and Cecilia Avramovich, as well as discussions and companionship from María Martín Rodríguez, Cecilia Vives, Zoe Kolita, and Liang Chen, as well as many others.

Resumen en Castellano

La tesis consiste en tres capítulos independientes de investigación en economía de la información. El primer capítulo estudia cómo el sesgo político se sostiene en un mercado de medios de comunicación. El segundo capítulo describe la realización de un experimento en el laboratorio en el cual investigo la utilización de recomendaciones sesgadas por determinados sujetos. El tercer capítulo busca vínculos entre dos diferentes formas de modelar heterogeneidad en juegos de transmisión de información, siendo: (i) preferencias no comunes y creencias comunes, o (ii) preferencias comunes y creencias no comunes. Se analiza cuándo y bajo qué circunstancias los dos tipos de modelos son equivalentes.

El capítulo 1 da comienzo con la observación de que votantes frecuentemente se suscriben a medios de comunicación que tienen más sesgo político que ellos mismos. Primero, muestro que eso es natural en un mercado es decir: un monopolista siempre elegirá el sesgo máximo y a pesar de que la competencia ayudará a disminuir el sesgo del mercado típicamente, no lo suprimirá por completo. En el caso límite de un mercado grande con competencia perfecta, el sesgo del mercado será exactamente igual al sesgo de los votantes. Estos resultados analizan el lado de la oferta del mercado. Añadimos a esto un análisis del lado de la demanda del mercado. Es posible, que los votantes demanden medios que son más sesgados que ellos mismos, bien porque tengan mucha confianza en sus ideas políticas o bien porque exista cierta incertidumbre sobre el sesgo de los medios.

El capítulo 2 es una prueba experimental de cómo utilizamos consejos de otras personas que tienen un sesgo comparado con nosotros. Para ello, se les pide a los sujetos del experimento que estimen la probabilidad de sacar una bola negra de una caja que

contiene 10 bolas que son ó negras ó blancas. La distribución de los colores es incierta. Cuando los sujetos toman el papel del Recibidor observan la estimación de otro sujeto (el Consejero) que tiene información privada sobre el contenido de la caja pero cuya estimación es sesgada. No obstante, es fácil ajustar la estimación para eliminar el sesgo. Encuentro que recibidores le dan más peso al consejo del consejero si el sesgo es pequeño aunque todos los consejeros son igualmente informativos a priori. Además, en una extensión les permito a los recibidores que elijan al consejero, sabiendo sólo el sesgo de éste. Los recibidores les dan más peso a los consejeros en este caso (comparado con el caso en el que los consejeros se asignan exógenamente) aunque, aquí como antes, todos los consejeros son igualmente informados a priori. Sin embargo, este efecto está sólo presente cuando recibidores eligen el consejero con menos sesgo. Interpreto el primer resultado como un tipo de homofilia - a los recibidores les gustan más los consejeros que son más similares a ellos - y el segundo resultado como un efecto de ilusión de control - los recibidores les dan más peso a consejeros cuando los pueden elegir.

El capítulo 3 trata dos diferentes formas de modular la heterogeneidad entre jugadores en juegos de transmisión de información. Tales juegos normalmente se caracterizan por dos jugadores: un Recibidor (R) y un Consejero (S). S tiene información privada sobre el estado del mundo y le manda un mensaje a R que toma una acción que afecta a los dos. Hay dos fuentes de heterogeneidad: las preferencias y las creencias de los jugadores. En la mayoría de los casos los jugadores tienen o preferencias no comunes y creencias comunes (CB) o preferencias comunes y creencias no comunes (CP). El capítulo investiga bajo qué circunstancias un modelo del tipo CB tiene un modelo del tipo CP que es equivalente. Esto es importante por dos razones: Primero, la elección entre modelos del tipo CB y del tipo CP es a veces arbitraria en el sentido de que no hay ninguna intuición que favorezca al uno o al otro. En este caso, es importante comprender cuando los dos enfoques son equivalentes. Segundo, aún cuando una situación favorece a uno de los dos, es importante saber cómo el enfoque afecta los resultados. Dos modelos son equivalentes si cumplen dos condiciones: (i) Equivalencia de Elección, que implica que en el modelo CB tanto como en el modelo CP los jugadores eligen la misma acción condicionada en la información, y (ii) Equivalencia de Estrategia, que implica que en el modelo CB tanto como en el modelo CP los jugadores eligen la misma

estrategia en equilibrio. Primero, investigo el caso de espacios de acciones y de estados del mundo que son discretos. Entonces, cualquier modelo CB siempre tiene un modelo CP equivalente si el espacio de estados es lo suficientemente grande comparado con el espacio de acciones. En el caso de que sean continuos los espacios identifico condiciones de suficiencia para que un modelo CB tenga un modelo CP equivalente. Por último, expongo que ni siquiera si los dos modelos son equivalentes según la definición empleada, no implica que otras propiedades sean iguales.

Contents

1	The Market for Biased News	11
1.1	Introduction	11
1.1.1	Motivation	11
1.1.2	Discussion of Model	14
1.2	A Model of a News Market	17
1.2.1	Payoffs	17
1.2.2	Information Structure	18
1.2.3	Distribution of Voters	20
1.2.4	Timing	20
1.3	The Information Transmission Game	21
1.3.1	Equilibrium of the Information Transmission Game	21
1.3.2	Some Comparative Statics of News Demand	23
1.4	The Market for News	26
1.4.1	Monopoly versus Duopoly	28
1.4.2	Increasing Competition	31
1.5	Demand-Side Bias: Adding Journalist Uncertainty and Voter Stubbornness to the Model	36
1.5.1	Journalist Bias Uncertainty	37
1.5.2	Voters of Different Prior Precision	40
1.6	Conclusion	43
1.7	Proofs	44

2	Using Information with a Known Bias in the Lab	53
2.1	Introduction	53
2.2	An Experiment on Probability Estimation	59
2.2.1	Setup	59
2.2.2	Scoring Rule	60
2.2.3	Sessions	61
2.2.4	Feedback and Matching	62
2.2.5	Benchmarks	62
2.2.6	Hypotheses	63
2.3	Preliminary Analysis: Aggregated Data	63
2.4	Methodology for Disaggregated Analysis	66
2.4.1	Risk Attitudes and Probability Estimates	67
2.5	Results	69
2.5.1	H1: Does Relative Bias Matter?	69
2.5.2	H2: Extreme Estimates	70
2.5.3	H3: Choosing the Sender in Stage 3	72
2.5.4	Comments	74
2.6	Discussion	78
2.6.1	Main Findings Corroborated by Literature	78
2.6.2	How Do the Results Relate to Empirics?	79
2.6.3	Policy Implications	80
2.7	Tables and Figures	81
2.7.1	Tables	81
3	Equivalence of Non-Common Priors and Non-Common Preferences in Information Transmission	82
3.1	Introduction	82
3.1.1	Motivation	82
3.1.2	Modeling Player Heterogeneity in the Information Transmission Literature	84
3.2	Setup	85
3.2.1	Model	85

3.2.2	A Measure of Equivalence	87
3.3	A Discrete World	90
3.3.1	The Simplest Case	91
3.3.2	The Size of the State Space Matters	93
3.3.3	A Sufficiency Result for Equivalence	95
3.4	Equivalence for Continuous Models	97
3.5	Preferences Over Information: An Example	101
3.6	Conclusion	105
3.7	Proofs	107

Introduction

The thesis consists of three independent chapters on information transmission in economics. The first chapter studies how political bias is sustained in a media market. The second chapter describes an laboratory experiment on using biased recommendations. The third chapter considers two different approaches to modeling player heterogeneity in information transmission games: either non-common preferences and common prior beliefs or common preferences and non-common priors. It is analyzed when and under what circumstances the two approaches can be seen as equivalent.

Chapter 1. The chapter starts with the observation that voters often subscribe to news from media that have more extreme views (bias) than the voters themselves. First, we show how this is natural in a market setting. In particular, a monopolist will always choose to be maximally biased and although competition generally helps to reduce bias it will typically not eradicate it. In the limit case of a large market with perfect competition between media, news bias will be exactly equal to voter bias. These results provide us with an supply-side explanation for why media tends to be more biased than the voters they cater to. We extend this with an analysis of the demand-side. It is possible that voters prefer more biased news media because they are either very confident in their beliefs or because there is uncertainty about media bias.

Chapter 2. The chapter is an experimental test of how we use the biased advice of others to form our own beliefs. In the experiment, subjects are asked to estimate the probability of drawing a black ball from a cage, which has an uncertain distribution of black and white balls. In the role of Receivers they observe the estimate of another

subject (the Sender) who has private information about the composition of the urn but whose estimate is biased. The bias is exogenous, observable and easily adjustable. We find that Receivers place higher weight on the Sender's estimate when the bias is small although there is no reason to do this; all Senders are equally informative a priori. In an extension we allow Receivers to choose between two Senders knowing only their bias. Receivers place more weight on the Sender's estimate in this case (compared to the case where the Sender is exogenously assigned) although, again, all Senders are equally informative a priori. However, this effect is only present whenever Receivers choose the Sender with the smallest bias. We interpret the first result as the existence of a type of homophily - more similar Senders are given higher weight - and the second result as an "illusion of control" effect - the Receiver places higher weight on the estimate of the Sender when he can choose him.

Chapter 3. The chapter digresses on two different approaches for modeling player heterogeneity in information transmission games. Such games normally feature an action taken by the Receiver and a payoff relevant state variable about which the Sender has private information. There are two sources of heterogeneity: player preferences conditional on the true state and player prior beliefs about the state. In most cases, players either have different preferences and common prior beliefs (CB) or common preferences but different prior beliefs (CP). The paper investigates under what circumstances a CB model has an equivalent CP model. This is important for two reasons. First, the choice between a CP or a CB model is sometimes arbitrary in the sense that there is no intuition to favor one or the other. In this case it is important to understand if and when the two approaches are equivalent. Second, even when the particular setting favors one of the approaches, it is important to understand how that approach affects results. To analyze equivalence we define two concepts: (i) Choice-Equivalence, which implies that in both models players want to take the same action conditional on available information, and (ii) Strategy-Equivalence, which means that in both models the players have the same equilibrium strategies. We find that when the state and action spaces are discrete then for any CB model an equivalent CP model always exists if the state space is sufficiently large relative to the action space. When the spaces are continuous we identify sufficient

conditions for the existence of equivalent models and identify a group of CB models that always have an equivalent CP model. Last, we argue that even equivalence obtains in the above sense, other properties may still be different, and in particular we show in an example that even when a CP and a CB model are equivalent they do not generate the same preferences for information.

Chapter 1

The Market for Biased News

1.1 Introduction

1.1.1 Motivation

The current US news market is abundant in news sources with a clear political bias and there is plenty of evidence of slanting. Take for example ex-CBS producer Bernie Goldberg’s [23] story of liberal slant in his former employer or the alleged personal influence on news coverage wielded by Fox News Chairman Roger Ailes.¹ Economic studies with various measures of bias have shown the news market to exhibit a general liberal bias with Fox News notably standing out on the conservative side (Groseclose and Milyo [27]; DellaVigna and Kaplan [16]; Lott and Hassett [31]²). A remarkable feature of the news market is the popularity of some of the very biased news shows. For instance, according to the Pew Research Center’s 2010 Media Consumption Survey³, 34% of Republican voters regularly watch conservative talk shows⁴ and 45% regularly watch Fox News. Yet only around 15% of Republicans characterize themselves as very conservative. Similar examples could be constructed for Democrats. We want to investigate this disconnect between the bias of voters and the news sources they subscribe to and offer

¹See <http://www.guardian.co.uk/media/2011/aug/10/roger-ailes-fox-news-murdoch>.

²Published by the conservative American Enterprise Institute.

³See <http://www.people-press.org/2010/09/12/americans-spending-more-time-following-the-news/>.

⁴The O’Reilly Factor, Rush Limbaugh, Sean Hannity, Glenn Beck.

supply- as well as demand-side explanations for why this may be so.

In particular we construct a model in which voters care about a single issue, the state of the climate, and must decide whether to vote for or against investing in renewable energy. To obtain information about climate change they can access a news market with potentially several news networks and subscribe to one of their news shows.⁵ Each news show features one journalist who observes private information (for instance a scientific report on climate change) and sends a cheap talk message. We use the pronominal convention that voters are denoted by *He* and journalists by *She*. Voters and journalists have the same preferences but non-common prior beliefs about the state of the climate. We refer to the expectation of the prior as *bias* and assume that it is observable. Networks care only about maximizing their market share. The timing is as follows. First networks observe voter and journalist bias and hire journalists. Voters then observe news show bias and choose which show (if any) to subscribe to. Journalists observe their private information and send a cheap talk message through their show. Finally voters cast their ballot. We say that voters who are a priori for/against investing in renewable energy are liberal/conservative and refer to the most biased voters as extremists, to the less biased voters as moderates and to the unbiased (in the sense of being a priori indifferent between investing or not) as neutrals.

The key intuition is that the more biased a journalist is, the stronger a signal she sends when she makes a recommendation that goes against her prior beliefs. Hence, the more conservative she is the more informative she will be if she recommends investing in renewable energy and vice versa. This implies that the voter will always prefer a journalist with the same direction of bias as himself (liberals prefer liberals, conservatives prefer conservatives). As an example, imagine a voter who is initially against investing in renewable energy but who is also a concerned citizen and wants to learn if he is right or wrong. If The New York Times (liberal) supports investing in renewable energy he puts it down to its liberal bias and he will not change his vote. On the other hand, if the Washington Post (conservative) recommends the investment the voter concludes that they must have observed a strong signal favoring investment and changes his vote. This incentive to learn when his prior is wrong implies that he will find the Washington

⁵Here we say news shows although this may also refer to newspapers or internet news sources.

Post more informative than The New York Times a priori.⁶

We show that a monopolist with two news shows chooses to be maximally biased in the sense that he captures the most extreme voters on both sides of the market. The reason for this is that if a biased voter subscribes to news from a given network so will all voters that are more moderate than him and hence a monopolist can capture the extremists without losing subscribers at the center. Hence, *moderates do not discipline a monopolist*. In the duopoly equilibrium the market bias is smaller: both networks are always at most as biased as the monopoly network since there is an incentive to move toward the center to capture market shares. Since there is no price competition the only cost of moving toward the center is the potential loss of biased subscribers. If we assume that the distribution of voter types is discrete and symmetric we can show that in fact only the relative number of extremists and neutrals matter to market bias. This provides us with a supply-side explanation for media bias and why moderate voters might be unable to induce networks to deliver less biased news although they would prefer it. Moving on to triopoly we show that the market bias is not necessarily reduced compared to duopoly and hence there is no monotonic relation between market bias and market size. In particular, if we make assumptions on the voter distribution that reflect data on US ideology we can show that at least one side of the market (and in some prominent cases both sides) will be more biased in triopoly than in duopoly. As the market grows very large each network focuses on a single voter group and hence the market bias corresponds exactly to the a priori voter bias. Thus, a social planner who cares only about market bias will not necessarily want more competition. He will have to weigh up two aspects: increased market bias implies that moderates subscribe to more biased news but it also brings some extremists (who would never subscribe to moderate news) to subscribe whereas before they did not. The net effect is not clear. If the social planner is concerned with the welfare of voters as measured by their utility, a monotonicity condition tells us that extreme voters gain and moderate voters lose when the market becomes more biased.

We also explore possible demand-side explanations for why moderates choose very biased news sources and focus on two cases: heterogeneous prior precision and un-

⁶The intuition is very similar to Cukierman and Tommasi [14] and Suen [45].

certainty about journalist bias. When prior precision is heterogeneous we show that stubborn/confident (high prior precision) voters prefer more biased media and open-minded/unconfident (low prior precision) voters prefer less biased media. The reason is that the more confident the voter is the stronger a signal he needs to change his mind and according to the previous intuition he can only get such a signal from a more biased journalist. Returning to the example, a very confident conservative with moderate bias might prefer Fox News over The Washington Post because only Fox can send him a strong enough signal to convince him to invest in renewable energy. This also carries over to the market setting where an increase in voter prior precision leads to more bias in equilibrium.

When there is uncertainty about journalist bias and only average bias is known the optimal journalist type for all voters is unambiguously more biased than in the absence of uncertainty. This is because the uncertainty makes the journalist's message less informative to the voter who therefore needs a more biased journalist to convince him. However, even though voters demand more biased news the market need not become more biased. When the market is moderately biased an increase in uncertainty leads to less market bias in equilibrium whereas when the market is very biased an increase in uncertainty leads to more market bias in equilibrium. The reason is that uncertainty has a two-fold effect on a voter's propensity to subscribe to a journalist who is more moderate than himself. On the one hand, as uncertainty increases the voter has less information about the signal observed by the journalist, which diminishes his gain from subscribing to her show. On the other hand, there is also a smaller probability that her signal was very moderate (so that he should not follow her recommendation), which augments his gain from subscribing to her show. When the voter is moderate the former effect dominates and news shows move toward the center and when he is extreme the latter effect dominates and news shows move away from the center.

1.1.2 Discussion of Model

Voters and journalists have common preferences but different priors. This assumption seems to well describe issues such as climate change: if everybody agreed that climate change is caused by fossil fuel combustion we would supposedly all want to invest in

renewable energy. What differs are our beliefs about climate change. So if a conservative wants to invest less in renewable energy than a liberal, it is because he believes climate change to be less likely. If the true state of the climate was known they would both want to do the same. Several information transmission papers incorporate non-common priors. Che and Kartik [11] consider players with different prior expectations, endogenous information acquisition and verifiable messages and show that receivers may prefer senders that have a different prior than themselves. Admati and Pfleiderer [1] model cheap talk with identical preferences but where the Sender is possibly overconfident about the informativeness of his signal and the message space is discrete. Kawamura [29] extends the model by adding bias in preferences and not restricting the message space. In the present paper we introduce a variant of our model that has a similar flavor to these latter two papers in that we allow voters and journalists to have different levels of confidence in their beliefs. Suen [45] is perhaps closest to the idea of this paper. He shows that the bias of both the optimal advisor and the set of acceptable advisors is increasing in the receivers own bias.

Voters acquire network news at zero marginal cost, which reflects that news are often either acquired through internet at zero cost or on television where different networks form part of a package of channels and therefore can be observed without incurring further costs. We assume that networks care only about maximizing their market share which corresponds to a situation in which their main source of revenue is advertising. This is important since networks will only be competing on their positioning in the market and not on prices. We therefore deviate from other models of news markets such as Mullainathan and Shleifer [34], Baron [4] and Gentzkow and Shapiro [22]. Mullainathan and Shleifer analyze Hotelling-competition on prices and positioning. They assume that news consumers have an intrinsic preference for news that are slanted toward their own position whereas in our model such preferences are derived from a desire to obtain the best information. Baron models news organizations that set prices and choose how much discretion to allow their journalists. Journalists always fully exercise this discretion to misrepresent news (their desire to do so could for instance be out of career concerns). Gentzkow and Shapiro take a different approach in which media may be of high or low quality and low quality media slant news toward consumers' priors in order to appear to

be of high quality. Mullainathan and Shleifer find that duopoly competition can lead to higher bias than monopoly whereas Gentzkow and Shapiro find that competition tends to decrease bias.

We choose to model information transmission in the model as cheap talk although evidence such as a scientific report is normally freely available to the public. The rationale is that although the evidence is accessible it often requires significant effort and knowledge to read and understand. Therefore, for the vast majority of the public, messages about complicated subjects such as climate change are effectively unverifiable. All the above papers on news markets feature cheap talk although in Baron journalists have only limited discretion to distort news and in Gentzkow and Shapiro some news sources always truly report the news. In an extension of the model we allow for uncertainty about the journalist's prior, which can be related to information transmission when sender preferences are uncertain (Stein [44]; Wolinsky [48]; Morgan and Stocken [33]; Kydd [30]) and noisy information transmission (Blume, Board and Kawamura [5]). These papers generally find that a small amount of uncertainty or noise improves the informativeness of the equilibrium. In our model the effect may go either way as the voters' preferences move ambiguously with the level of uncertainty.

Our model has voters subscribing to news only to obtain information (which is different to for instance Mullainathan and Shleifer). A common objection is that people perceive news as entertainment and simply prefer to be confirmed in their beliefs. Again we appeal to data from the Pew Research Center's 2010 Media Consumption Survey, where regular consumers of a given news source are asked for their reason to turn to the source.⁷ The Colbert Report, the Daily Show, USA Today and the morning shows are the only shows for which at least 15% of the regular viewers state that they mainly watch the show for entertainment. This is hardly surprising given that these shows are staged as entertainment shows. We conclude that, at least according to consumers themselves, entertainment is not an important motivation for following the news.

The particulars of the model are set in terms of news and voters but we imagine the results to carry over to any situation in which a profit maximizing organization hires

⁷Question 82a-x of the survey.

agents to provide information to a third party who must take a (binary) action. These could be consultancy firms and companies, rating agencies and investors or political think tanks and politicians.

Section 1.2 introduces the model. Section 1.3 describes the demand for news and identifies voters' preferred journalist types. Section 1.4 analyzes market outcomes. First we focus on monopoly versus duopoly. Then it is shown that the results do not vary monotonically with the number of firms and finally the limit case of a large market with many firms is analyzed. Section 1.5 introduces uncertainty about journalist bias and heterogeneity between voters and journalists in the precision of their prior beliefs. Comparative statics are derived on how these extensions affect market bias. Section 3.6 concludes. All proofs are relegated to Section 1.7.

1.2 A Model of a News Market

1.2.1 Payoffs

Let θ be a random variable supported on \mathbb{R} which, for illustration, represents the state of the climate, such that if $\theta < 0$ climate change is important and we should invest in renewable energy whereas if $\theta > 0$ the extent of climate change is too small to affect us. The voter (he) and the journalist (she) have the same quadratic utility function

$$u(y, \theta) = -(y - \theta)^2,$$

where $y \in \{-\gamma, \gamma\}$ is the ballot cast by the voter. Here $y = \gamma > 0$ corresponds to voting against investing in renewable energy and $y = -\gamma$ to voting for investing. We want to make two observations about the utility function. (i) The quadratic utility function captures a situation in which there are convexly increasing costs to making the wrong decision. Thus we are implicitly making the assumption that it is equally costly to err on both sides. An alternative example to illustrate this is that climate change is occurring and voters must choose the best response (for instance, either emphasize investment in technology or carbon trading). (ii) This is not a voting model in the sense that voters' utility depends on winning a referendum: voters gain utility from doing

“the right thing”. That is to say, if they believe that climate change is indeed occurring, they gain utility from voting for investing in renewable energy regardless of whether they win the referendum or not.

The news market consists of a discrete and non-empty set \mathbf{N} of N news networks. Each network $n \in \mathbf{N}$ seeks only to maximize its market share $s_n \in [0, 1]$. As discussed above, this corresponds to a situation in which networks generate income only from advertising. We also assume that voters subscribe to at most one (and possibly zero) news shows at zero cost. Perhaps the best illustration of this setup is a person who likes to watch the news on TV each night. He only watches one show, for instance because he does not have time for more or because the shows overlap and he owns no recording device. He can also choose to turn the TV off. Hence, he will either watch the show that gives him the most utility or, if he does not expect to learn anything useful from any of the shows, he will spend his evening in other activities.

One important consequence of the setup is that networks are always focussed on capturing new subscribers and not on providing better content for existing subscribers. A change in strategy that increases market share but uniformly lowers the utility of all existing subscribers is always beneficial and thus networks search for the lowest common denominator.

1.2.2 Information Structure

Before observing any information players have normally distributed prior beliefs over the state of the climate. The beliefs of a type- b voter are given by

$$\theta \overset{b}{\sim} N(b, 1).$$

The voter’s bias is parametrized by b , his prior expectation of θ . Denote the density function of the beliefs by $f_b(\theta)$. To facilitate the interpretation we refer to voters with $b > 0$ as conservatives and voters with $b < 0$ as liberal.⁸ On the other hand, a type- β

⁸Although there seems to be a general perception that liberals are “tree-huggers” and conservatives are against climate change investments (believing for example that climate change is not happening or that the market will automatically solve the problem) this is of course a gross simplification but we hope that the reader will bear with it for the sake of exposition.

journalist has beliefs

$$\theta \stackrel{\beta}{\sim} N(\beta, 1).$$

Journalist bias is parametrized by β , her prior expectation of θ . We make two assumptions about beliefs that will be relaxed later. First we have assumed that journalist bias is known. In Section 1.5.1 we extend the model to allow the bias to be unobservable but drawn from a commonly known distribution. Second we have assumed that the precision of voters' and journalists' priors is the same. In Section 1.5.2 we allow this to differ and give precision the interpretation of confidence or stubbornness. These assumptions have no implications on the qualitative results of Section 1.4 but by introducing them we can do comparative statics on the market equilibrium to analyze how it moves with different levels of confidence and bias uncertainty. The private information of journalists (for instance a climate report) takes the form of an unbiased and normally distributed signal x , which is the same for all journalists.

$$x \sim N(\theta, 1).$$

The precision of the signal has no qualitative effect on the results and therefore we set it equal to one to save on notation. Let the density function of x conditional on θ be denoted by $\lambda(\cdot|\theta)$. A priori a type- b voter will believe x to be distributed with the following density function.

$$l_b(x) \equiv \int_{-\infty}^{\infty} \lambda(x|t) f_b(t) dt.$$

This is a normal distribution with expectation b and variance 2. Often it will be useful to express quantities in terms of $l(\cdot) \equiv l_0(\cdot)$ such that the beliefs of a type- b voter are given by $l_b(x) = l(x - b)$. Observe that this is one of the important implications of modeling bias in beliefs rather than in preferences. Bias in beliefs affects not only the preferred action conditional on x but also the expectation of the distribution of x . Bias in preferences, on the other hand, affects only the conditional preferred action. This will be important for several results since, the more biased a voter is, the more likely he thinks it is that he will receive a message that confirms his prior beliefs. Consequently, he will have a stronger preference for same-bias journalists than if beliefs were independent of bias.

1.2.3 Distribution of Voters

Voter types are characterized by their prior beliefs, which are summarized by b . We assume that bias is bounded such that $\underline{b} \leq b \leq \bar{b}$ and distributed according to $\pi(b)$, which may have continuous or discrete support. When b and b' have the same sign we say that *voter type b is more moderate than type b' if $|b| < |b'|$* . Furthermore, we always assume that there are voters on both sides of the market, i.e. $\pi(b) > 0$ for some $b > 0$ and $\pi(b') > 0$ for some $b' < 0$.

When the support is discrete then $b \in \{b_k\}_{k=\underline{K}}^{\bar{K}}$ and we define $\pi_k \equiv \pi(b_k)$ for $k \in \{\underline{K}, \dots, \bar{K}\}$. We use the index k to represent the direction of bias such that $b_k < 0$ for $k < 0$ and $b_k > 0$ for $k > 0$ with $b_0 = 0$ and order the groups such that $k' > k$ if $b_{k'} > b_k$. In this case, we define a *symmetric voter population* by $\bar{K} = \underline{K} = K$, $b_{-k} = b_k$ and $\pi_{-k} = \pi_k$.

1.2.4 Timing

The timing of the game is as follows.

Stage 1: The owner of network $n \in \mathbf{N}$ hires a journalist of bias $\beta_n \in \mathbb{R}$.⁹

Stage 2: The voter observes the bias of each network $n \in \mathbf{N}$ and subscribes to at most one. If a discrete voter group is indifferent between several networks we assume that it splits evenly between them.

Stage 3: Journalists observe the realization x of their private information, update beliefs and send a cheap talk message $m_n \in M$ to their subscribers.

Stage 4: Voters cast their ballot y .

All aspects of the game except x are common knowledge. As is shown below journalists will always send one of two messages and therefore we can assume without loss of generality that $M = \{-\gamma, \gamma\}$ such that the message is to be understood as a recommended action. Our setup corresponds to a model in which journalists know that they are better

⁹Other papers such as Baron [4] discuss journalists as being generally biased in one direction but to simplify we assume that news networks can choose freely any type of journalist.

informed than voters and once in Stage 3 the audience of each network is fixed such that the best a journalist can do is to try to persuade her audience to vote for her preferred action. The equilibrium concept employed is Perfect Bayesian Equilibrium. Since there is no uncertainty about voters' type, sequential rationality implies that the players cannot commit to any other strategy than playing the Bayesian Nash Equilibrium¹⁰ in Stages 3 and 4. If the set of voters that subscribe to a given network is non-empty the journalist of the network will have only one consistent strategy, which is to truthfully reveal his preferred action but not the exact value of the signal. This strategy does not depend on the size nor the composition of the audience.

1.3 The Information Transmission Game

In this section we analyze Stages 3 and 4 of the game, which will pin down the demand side of the market. For each voter group we identify two journalist types that will play a role in the market equilibrium. First we identify the optimal journalist type for each voter type. This journalist type will be important when a voter group is sufficiently large such that networks are willing to choose their bias solely to attract this group and disregard other voters. Second we define the least informative journalist type of a voter as the journalist type that makes the voter exactly indifferent between following the message of the journalist or ignoring her. This type will help us to bound the market bias in monopoly and duopoly.

1.3.1 Equilibrium of the Information Transmission Game

Conditional on x , a type- β journalist has posterior expectation $(\beta + x)/2$. Using the bias-variance decomposition and the updating formula for normally distributed priors with normally distributed data, we can calculate a type- β journalist's expected utility

¹⁰Bayesian Nash Equilibrium in this case consists of a signaling rule for the journalist and an action rule for each voter with the properties that (i) conditional on the action rules of their subscribers, the journalists' signaling rules maximize their expected utility, and (ii) conditional on the signaling rule of the journalist of the network to which a voter subscribes, his action rule maximizes his expected utility.

conditional on x , y and her type as

$$U_{\beta}^J(y, x) = - \left(y - \frac{\beta + x}{2} \right)^2 - 1.$$

Having observed x , the journalist wants y to be chosen by the voter so as to maximize $U_{\beta}^J(y, x)$, which corresponds to minimizing the absolute value of $(y - (\beta + x)/2)$. In equilibrium at most two actions can be induced and therefore the optimal strategy $m_{\beta}(x)$ of a utility-optimizing type- β journalist reduces to

$$m_{\beta}(x) = \begin{cases} -\gamma & \text{whenever } x < -\beta \\ \gamma & \text{whenever } x \geq -\beta. \end{cases}$$

This is always an equilibrium strategy regardless of the audience of the network. If the audience is empty the journalist is indifferent between any two strategies. If the audience is non-empty the only strategies that maximize her expected utility are strategies that induce $y = m_{\beta}(x)$ conditional on x and therefore this is essentially the unique equilibrium strategy. Observe that in equilibrium the journalist truthfully reveals her own preference over the actions. What she does not reveal is the intensity of her preference. Conditional on x and y , the expected utility of a type- b voter is

$$U_b^V(y, x) = - \left(y - \frac{b + x}{2} \right)^2 - 1.$$

Letting y_b be the default action for a type b -voter (i.e. $y_b = -\gamma$ if $b < 0$ and $y_b = \gamma$ if $b \geq 0$) the utility gain from following the message of the news source is defined by

$$G(b, \beta) \equiv \int_{-\infty}^{-\infty} [U_b^V(m_{\beta}(x), x) - U_b^V(y_b, x)] l_b(x) dx.$$

This function is symmetric in b around 0 and therefore we can focus on the case where $b \geq 0$ (i.e. the voter is “conservative”) without loss of generality. In this case $y_b = \gamma$ and therefore the voter only potentially changes his action whenever $m = -\gamma$. Hence, we can write his gain as $G(b, \beta) = \int_{-\infty}^{-\beta} [U_b^V(-\gamma, x) - U_b^V(\gamma, x)] l_b(x) dx$. Calculating this we arrive at

$$\begin{aligned} G(b, \beta) &= -4\gamma \int_{-\infty}^{-\beta} \frac{b + x}{2} l_b(x) dx. \\ &= -2\gamma \int_{-\infty}^{-\beta-b} (2b + z) l(z) dz. \end{aligned} \tag{1.1}$$

The voter should follow the message whenever the gain is positive and will prefer the news source with the highest gain. Hence, this function is a sufficient statistic for determining voter behavior. Observe that the integrant in the first line is the voter's posterior expectation of the state given x weighted by the density of x , and that this density is evaluated using the voter's prior beliefs. Thus priors affect expected gains not only through the expectation of the state conditional on the signal but also in through the assessment of how likely each signal is.

1.3.2 Some Comparative Statics of News Demand

In this subsection we make a number of observations that will be useful in the further analysis and which are interesting in their own right. It is straightforward to identify the voter's preferred journalist type and show that gains are monotonic in their distance to the optimal type. We can then show that (i) moderates gain more in optimum, (ii) for each voter there is a unique journalist type β that gives zero gains and this type is more moderate than the voter, and (iii) if two journalist types are spaced equally far from a voter the voter prefers the most extreme type. Let the voter's preferred journalist-type be denoted

$$\beta^*(b) \equiv \arg \max_{\beta \in \mathbb{R}} G(b, \beta).$$

Take the derivative of (1.1) with respect to β to obtain $G_2(b, \beta) = -2\gamma(b - \beta)l(-\beta - b)$. It follows from the first-order condition $G_2(\cdot) = 0$ and by checking the second-order condition that the voter's preferred journalist type is the type that has the same bias as the voter, which is not surprising given the symmetry of the setup.

$$\beta^*(b) = b.$$

It follows that in the baseline version of the model there is no demand-side incentive for bias to be extreme. The result seems almost self-evident but we shall see in Section 1.5 how the preferred type changes when we introduce different levels of prior confidence and uncertainty about journalist bias. From $G_2(\cdot)$ it is also easy to see that the gain is monotonic around the optimal type.

$$G_2(b, \beta) > 0 \text{ for } \beta < \beta^*(b) \text{ and } G_2(b, \beta) < 0 \text{ for } \beta > \beta^*(b). \quad (\text{M})$$

This is important since competition will often take place at the center ($b = 0$). Given that $G(0, \beta)$ is symmetric around $\beta = 0$ we can deduce from (M) that the network with the smallest absolute bias will win the central group. Denote the gain at the optimum for a type- b voter by $G_*(b) \equiv G(b, \beta^*(b))$. Substituting we get $G_*(b) = -2\gamma \int_{-\infty}^{-2b} (2b + z)l(z) dz$ and taking the derivative with respect to b and applying the Envelope Theorem we arrive at

$$G'_*(b) = -4\gamma \int_{-\infty}^{-2b} l(z) dz.$$

This is clearly decreasing in b for $b > 0$ and increasing in b for $b < 0$, which leads to the following observation.

Observation 1. *Moderates gain more in optimum. I.e., $G'_*(b) > 0$ for $b < 0$ and $G'_*(b) < 0$ for $b > 0$.*

Information is more valuable to moderates because they are more likely to make use of it, that is to say, to change their mind. An extremist will consider it unlikely that his preferred news source will receive a signal strong enough to send a message that recommends him to switch away from his default action. Therefore he expects to gain less than a moderate.

Notice that the gain can also be negative. However, negative gain occurs only if the journalist is more moderate than the voter. To see this, notice that the limit of $G(b, \beta)$ as β goes to infinity is $-4b\gamma$ and as β goes to negative infinity is zero. It then follows from (M) that there can at most be one crossing and that for $b > 0$ ($b < 0$) this occurs for some $\beta < \beta^*(b)$ ($\beta > \beta^*(b)$) whereas for $b = 0$ there is no crossing. Thus for each $b \neq 0$ there is a unique journalist type which yields exactly zero gain. This type will play an important role in the determination of the market equilibrium and therefore we denote it by $\beta^0(b)$ such that

$$G(b, \beta^0(b)) = 0.$$

We call this journalist type the *least informative journalist* for a type- b voter and applying the above analysis we can state the following observation.

Observation 2. *For each voter type, the least informative journalist is always more moderate than the voter. I.e., $\beta^0(b) < b$ for $b > 0$ and $\beta^0(b) > b$ for $b < 0$.*

It is interesting to consider this for a moment. A biased voter will always want to follow the message sent by an even more biased journalist but he may not want to follow a less biased journalist. The reason is that an extremely conservative/liberal journalist is always informative to a conservative/liberal voter when she sends a message that goes against her prior, since in this case the journalist reveals a great deal of information (the set of signals that could have induced this message is small). The same is not necessarily true for a more moderate journalist or a journalist who is biased in the other direction. Yet, even though an extremely biased journalist always gives positive gain to voters of the same direction of bias, the overall expected gain from subscribing to her may be very small since there is only a very small probability that she sends a message that goes against her prior. The majority of the time she will send a message that confirms her prior and thus be uninformative.

If two news networks are available, how will voters choose between them? When the two news sources are either more moderate or more extreme than the voter we know from the monotonicity of the gain function that he will choose the one that is closest to his optimal type. When one is more moderate and the other is more extreme the answer is less straightforward. Suppose there are two networks, L and R , with biases β_L and β_R such that $\beta_R > \beta_L$. If voters of type b are indifferent between subscribing to L and R then they must have the same gain from the two networks.

$$G(b, \beta_L) - G(b, \beta_R) = 0. \quad (1.2)$$

It is shown in the proof of Corollary 2 that (1.2) always has a unique solution, which we denote by the function $b(\beta_L, \beta_R)$. It can be checked that $b(\cdot)$ is a continuous and monotonically increasing in both arguments. Recall that the derivative of the gain with respect to the journalist type when $b > 0$ is $G_2(b, \beta) = -2\gamma(b - \beta)l(-\beta - b)$. Since $l(\cdot)$ is a normal distribution with zero mean and hence symmetric around $\beta = -b$ whereas the utility function is symmetric around $\beta = b$, the derivative is asymmetric around $\beta = b$. In particular for any $\epsilon > 0$ we have $G_2(b, b - \epsilon) + G_2(b, b + \epsilon) > 0$. Combining this with $\beta_L < b(\beta_L, \beta_R) < \beta_R$ and (M) we arrive at the following observation.

Observation 3. *The indifferent voter is always closer to the most biased journalist. I.e., suppose $|\beta_R| > |\beta_L|$. Then $|\beta_R - b(\beta_L, \beta_R)| > |\beta_L - b(\beta_L, \beta_R)|$.*

Why? Effectively, voters are concerned about making mistakes. Take the example of a conservative ($b > 0$) voter. Without any message he will vote against investing in renewable energy. If a journalist is less conservative than the voter there will be signals for which the journalist recommends investing but the voter would have preferred not investing. If the journalist is more conservative there will be signals for which the journalist recommends not investing but the voter would have preferred investing. Given the symmetry of the utility function both types of mistakes are equally costly to the voter. But because his beliefs about the distribution of the signal x are derived from his prior he believes the first case to be more likely, and therefore he prefers the more conservative journalist.

1.4 The Market for News

In this section we analyze Stages 1 and 2. From the previous section we know the strategy of journalists and we have identified two journalist types for each voter type - the optimal journalist and the least informative journalist. Particularly the least informative journalist type will be important for determining the market equilibrium. We will be interested in analyzing which bias is chosen by networks as well as the welfare that voters obtain: previous work suggests that competition can both augment and diminish bias and that it is not necessarily conducive to welfare.

The main analysis of the section is a comparison of monopoly and duopoly competition. To make the comparison more interesting, we will assume that the monopolist manages two news shows whereas each duopolist network manages one show. This is to capture the idea of a large market in which a monopolist would ideally want to cover both sides of the spectrum to maximize his profits. We analyze examples of a triopoly to show that the relationship between bias and market size is not monotonic and suggest why this might be so. Lastly we look at the limit case of a fully competitive news market in which there are enough media such that at least one network caters directly for the bias of each voter group.

Since the gain function gives us the expected utility gain from following a message a natural measure of a type- b voter's welfare is the following.

Definition The welfare of a type- b voter in the news market \mathbf{N} is

$$W_b = \max_{n \in \mathbf{N}} \{G(b, \beta_n), 0\}.$$

The voter's welfare in a news market \mathbf{N} is the maximum gain he can achieve if this is positive and zero if this is negative, since he can always choose not to pay attention to the news. Any voter can potentially obtain strictly positive welfare a priori since $G(b, \beta^*(b)) > 0$ for all b but in a given market this may not be possible if the media are too biased away from the voter's own bias.

Although we have assumed that news can be acquired at zero cost we want to rule out equilibria in which voters acquire news although they do not intend to follow them. We therefore make the following assumption.

Assumption 1. *If type- b voters have negative gain from following the advice of a type- β news network ($G(b, \beta) < 0$) they do not subscribe to it.*

This assumption implies that whenever a voter cannot obtain a positive expected gain from subscribing to a network he chooses not to do so (even though subscription is costless in our model). We can think of this in terms of a model where the voter has a subscription cost c . The above assumption then corresponds to the limit case as $c \rightarrow 0$. Before starting the analysis we establish a lemma which will be the backbone of the results in this section.

Lemma 1. *If a given voter has positive gain from a network then so do all voters who are more moderate. I.e., if $b' > 0$ and $G(b', \beta) > 0$ then $G(b, \beta) > 0$ for all $0 \leq b < b'$. Conversely for $b' < 0$.*

Lemma 1 implies that in the absence of competition networks will choose their bias to attract extremists. Whenever they have the extremists on board, the moderates will follow. This is important because it implies that behavior will not resemble that of, for instance, a Hotelling model in which attracting extremists comes with the potential cost of losing moderates. *In our model moderate voters cannot discipline a monopolist and drive him toward the center.* This observation is essential for our comparison of monopoly and duopoly. The lemma is somewhat stark and stems from the fact that

we assume that all voters subscribe as long as they have positive gain. If for instance subscription is probabilistic and this probability depends on the size of the gain, it is possible that the incentive to increase the subscription probability of moderate voters would be greater than the incentive to attract extremists and in this setting networks could be disciplined by moderate voters even in the absence of competition.

To have a reference point for market bias we identify the most moderate journalist type that is informative to the most extreme voter on both sides. In particular, $\underline{\beta} \equiv \beta^0(\underline{b})$ and $\bar{\beta} \equiv \beta^0(\bar{b})$. As we shall see, these are the smallest biases a monopolist will choose and the largest biases that a duopolist will choose. However, when we increase competition further and move to triopoly, it is possible that some network will strictly prefer $\beta < \underline{\beta}$ or $\beta > \bar{\beta}$. As competition increases networks abandon the small-market strategy of attracting a large audience and begin to focus on smaller groups and on choosing a news bias that is closer to these voters' bias.

1.4.1 Monopoly versus Duopoly

First we will compare monopoly and duopoly to show that market bias clearly diminishes under duopoly. Then we will move on to triopoly to show that generally speaking the case is not so clear and that under reasonable assumptions about the voter distribution, triopoly will lead to more bias at the extremes and less at the center. To make the comparison between the monopoly and the duopoly cases fair, we assume that the monopolist owns two news shows, L and R , and hires one journalist for each show, with respective bias β_L^M and β_R^M . Without loss of generality let $\beta_L^M \leq \beta_R^M$. We can then show the following.

Proposition 1. *The monopolist chooses any β_L^M and β_R^M to cover the whole market. I.e. $\beta_L^M \leq \underline{\beta}$ and $\beta_R^M \geq \bar{\beta}$.*

The key to this result is Lemma 1: moderate voters do not discipline the monopolist. Rather, the situation is the opposite - in monopoly the extremists dictate the market. The assumption of the monopolist owning two networks is also important. If he only owns one then it might be optimal to locate at the center and the comparative statics with duopoly case would be very different. If we assume that there is a fixed cost of

running a network then the case we consider is the case in which there are sufficiently many extreme voters such that a monopolist cannot cover a large part of the market with just one news show. In this case it will be profitable to set up a second show but never more, since with two shows the monopolist can cover the entire market. As an example of organizations with multiple news outlets, Rupert Murdoch's News Corporation owns a variety of British papers from the serious *The Times* and *The Sunday Times* to tabloids *News of the World* (before it was dismantled) and *The Sun*, all of a different profile and supporting both the Labour Party and the Conservative Party.¹¹

We now analyze the situation of two independent networks, L and R , who each choose one journalist to maximize their own market share. Denote the biases chosen in equilibrium by β_L^D and β_R^D . To avoid less interesting cases where the voter population is so skewed that networks compete only on one side of the spectrum, we suppose that $\beta_L^D \leq 0 \leq \beta_R^D$. If the voter distribution is continuous then the two networks will be in direct competition, in the sense that they can steal market shares from one another by moving closer to the center. This will always be true because of Lemma 1 but may not be true in other models. Returning to the example of Hotelling competition, two vendors may be sufficiently far from each other such that their pools of potential buyers do not overlap. In our model there is always overlap at the center and therefore there is always competition. It follows that it will never be optimal to set $\beta < \underline{\beta}$ or $\beta > \bar{\beta}$. We summarize this in the following proposition.

Proposition 2. *If voter distribution is continuous, market bias is less extreme in duopoly than in monopoly. I.e.*

$$\underline{\beta} \leq \beta_L^D \text{ and } \beta_R^D \leq \bar{\beta}. \quad (1.3)$$

If the voter distribution is discrete and symmetric with $\pi_0 > 0$ the same result holds.

In the present model, competition has the potential to make news less biased since networks will be competing for the moderate voters. If instead we were in a situation where the monopolist chooses only one outlet and places himself at the middle the

¹¹See for instance http://en.wikipedia.org/wiki/Rupert_Murdoch.

proposition would have been the other way around, i.e. competition would increase or maintain bias. But this would have been an effect of the quantity of news shows rather than competition. We can make a comparison here to two other papers that consider monopoly versus duopoly competition in the news market. Mullainathan and Shleifer [34] find that duopoly is more biased than monopoly. There are two reasons for the difference between our result and theirs. First, in Mullainathan and Shleifer duopolists differentiate themselves to charge a higher price. In our model, only market shares are important and therefore there is no incentive to differentiate. Second, they use a Hotelling model with intrinsic preference for bias in which the transport cost implies that if a news network is sufficiently biased it may lose the moderate subscribers. This will never happen in our model as shown in Lemma 1: given that we see messages as the result of an information transmission game an extreme news network will always be informative when it sends a signal opposite to its prior.

The following corollary to Proposition 2 states a condition that allows us to identify equilibrium bias in a special case. Let $\beta_k^0 \equiv \beta^0(b_k)$.

Corollary 1. *Suppose the voter population is discrete and symmetric with $k^* \equiv \max\{k \in \mathbb{N} : \pi_k \geq \pi_0/2\}$. In any duopoly equilibrium,*

$$\beta_L^D = -\beta_{k^*}^0 \text{ and } \beta_R^D = \beta_{k^*}^0.$$

For symmetric distributions this corollary tells us when the duopoly news market is maximally biased, i.e. when (1.3) holds with equality. This will occur whenever

$$\pi_K \geq \frac{\pi_0}{2}.$$

The corollary is interesting in that it tells us that in *symmetric equilibria the market bias is determined by the relation of neutrals to extremists*. Even if an extremely large group of moderates exists it will not influence equilibrium bias if only the extremists are sufficiently many compared to the neutrals. To see why the corollary holds, observe that when the voter distribution is symmetric networks will be competing over the central group of voters ($b = 0$) while trying to hold on to as many of the biased voters as possible. Suppose we start from a symmetric position and let \bar{k} and \underline{k} be the most extreme group of voters that buy from R and L , respectively. In this case each network

has half of the neutral group, i.e. $\pi_0/2$, and a small deviation toward the center would allow them to gain the whole group, i.e. a gain of $\pi_0/2$ in market share. Such a deviation is costless for R if $\beta_R > \beta_k^0$ and carries the cost of $\pi_{\bar{k}}$ if $\beta_R = \beta_k^0$. Clearly, in the first case R will always make the deviation and in the second case he will make the deviation if $\pi_0/2 > \pi_{\bar{k}}$.

Until now we have been concerned with the bias of the news market. The goal of a social planner may be to diminish news market bias, either because bias is viewed as intrinsically bad or to incur voters to make better collective decisions. But the social planner may also be concerned with the welfare of voters, as measured by their gain. The following corollary to Proposition 2 tells us how welfare is affected by competition.

Corollary 2. *Suppose the monopoly and duopoly equilibria are not identical. Then in duopoly welfare is lower for extreme voters and higher for moderate voters than in monopoly. In particular, there exists $b' \geq 0$ and $b'' \leq 0$ such that $W_b^D > W_b^M$ for $b' < b < b''$ and $W_b^D < W_b^M$ for $b < b'$ and $b > b''$.*

The result is intuitive. We know from (M) that a voter's gain is monotonically decreasing in the distance between her bias and that of the network. Therefore the duopoly is necessarily better for voters at the center because it is less biased. Furthermore, we know that the monopolist bias equals the minimally biased network for the most extreme type. Hence, for this type the monopoly is already too moderate and therefore he loses out further in duopoly when network bias becomes more moderate.

1.4.2 Increasing Competition

In this section we will first provide a class of voter distributions for which we can say that increasing competition increases the bias at the extremes. We will argue that the type of distributions that we are considering here are close to the empirical distribution of political bias in the United States. Afterwards we will show a counterexample with a discrete voter distribution and consider why the results are different.

Competition Increases Bias When the Voter Distribution Is Single-Peaked

In this section we will consider a class of voter distributions that satisfy $|b|, |\bar{b}| \leq 1$ and that $\pi(\cdot)$ is (weakly) single-peaked with peak at \tilde{b} such that $\pi'(b) \geq 0$ for $b \leq \tilde{b}$ and $\pi'(b) \leq 0$ for $b \geq \tilde{b}$. The first assumption says that we need bias to be at most one standard deviation away from neutrality - recall that voters' beliefs have variance 1. This assures that the normal distribution is concave at the points where we need to evaluate it to prove the lemma below. The second assumption implies that networks have more to gain the closer they move to the peak of the voter distribution. As Table 1.1 illustrates the assumption of a single-peaked population matches (coarsely) the distribution of ideology in the US, which has a small peak at "moderates" and is skewed towards conservatism. We first consider what happens when a network moves

Table 1.1: US Political Ideology

Very Liberal	Liberal	Moderate	Conservative	Very Conservative
4.4%	14.3%	40.1%	33.3 %	7.9%

Source: Pew Research Centers 2010 Media Consumption Survey.

closer to another. It is not clear from the setup of the model how networks react to increased competition: do they stand their ground or do they withdraw? There is no clear intuition here. If L moves closer to R then R loses market share at the center and this market share could be won back by R if he moves closer to the center but this comes at price in terms of losing market shares on the extreme. If we suppose that L and R are initially in an equilibrium then R will only compete and move towards the center if in the new situation $b(\cdot)$ is more responsive than before, i.e. if R can gain more market share by moving toward the center than before. Will this be the case? The following lemma derives the comparative static.

Lemma 2. *Suppose $b \in [-1, 1]$ and let $\beta_R > |\beta_L| \geq 0$. Then competition decreases the closer L and R are together in the sense that*

$$b_{21}(\beta_L, \beta_R) < 0.$$

Since $b_2(\cdot) > 0$ the lemma states that network R gains less market share from moving towards the center the closer L is to him and hence if the voter distribution is uniform networks will always move in the same direction. I.e. if L moves to the right R , rather than moving to the left to compete for his market, will resign and also move to the right. Hence, what keeps the networks from both moving towards the middle is the discipline of extreme voters who might cancel their subscription if networks become too moderate. The lemma is driven by the fact that the closer the networks are to each other the less responsive is $b(\cdot)$. However, as L moves closer to R then $b(\cdot)$ increases and this has an ambiguous affect on the derivative which is, however, dominated by the first effect.

Suppose the two networks L and R are in an equilibrium in which they choose biases β_L^D and β_R^D . Now a third network M enters the market and we assume, without loss of generality, that $\beta_R^T \geq \beta_M^T \geq \beta_L^T$ and $|\beta_R^D| \geq |\beta_L^D|$. I.e. the new network enters between the two old networks and in duopoly, R is the most biased network. There are two cases. In the first case $\beta_R^T \geq \beta_M^T \geq \beta_R^D$ or $\beta_L^T \leq \beta_M^T \leq \beta_L^D$. That is to say, in triopoly the network M is more extreme than R or L was in duopoly. In this case, by assumption one side of the market becomes more biased. In the other case $\beta_R^D \geq \beta_M^T \geq \beta_L^D$. Then, by Lemma 2, R will have less incentive to compete toward the center and unless it was already maximally biased ($\beta_R = \bar{\beta}$) it will move toward the extreme. This establishes that at least one of the two original networks L and R becomes more biased in triopoly. To see that it is possible that both L and R become more extreme suppose that the duopoly equilibrium is symmetric such that $\beta_R^D = -\beta_L^D$ and that $\beta_M^T = 0$. Then by Lemma 2, $\beta_L^T \leq \beta_L^D$ and $\beta_R^D \leq \beta_R^T$, with the inequalities holding strict if not already L and R are maximally biased ($\beta_R = \bar{\beta}$ and $\beta_L = \underline{\beta}$). We summarize this in the following proposition.

Proposition 3. *Suppose $|b|, |\bar{b}| \leq 1$ and that $\pi(\cdot)$ is (weakly) single-peaked. Without loss of generality, let $\beta_M^T \geq 0$. Then either*

$$\beta_L^T \leq \beta_L^D \leq \beta_M^T \leq \beta_R^D \leq \beta_R^T,$$

or

$$\beta_L^T \leq \beta_M^T \leq \beta_L^D \leq \beta_R^T \leq \beta_R^D.$$

The proposition states that a new entrant always creates more bias on at least one side of the market although he will also create less bias towards the center. It follows that competition (in the sense of moving from duopoly to triopoly) can not uniformly diminish bias - there will always be at least one side of the market that becomes more biased. We can relate this finding to some data. DellaVigna and Kaplan [16] estimate that after its introduction to the American news market Fox News persuaded 3 to 8 percent of its non-Republican voters to vote Republican. This appears to be an example of a new player entering a news market, which was to some extent disciplined by moderate voters initially, and “scooping up” the more biased consumers on one side of the spectrum, thus adding bias to the market.

The effect observed in triopoly suggests it will always be impossible to add a new entrant without making some networks more biased (although others may become less biased). Thus, a social planner interested in diminishing news market bias will need to specify an objective function to make decisions. In the above case, the planner must weigh two effects: (i) the gain/loss from those who subscribe to networks that become less/more biased, (ii) the gain/loss from those who did not/ did subscribe in duopoly who do/do not subscribe in triopoly. The second effect is not considered by the proposition: making news markets more desirable in some cases since it means that extremists who, in the absence of subscribing to any news, are completely rigid in their opinions may actually respond to the signal.

A Counterexample

In this section we give a counterexample to the result of the previous section and try to explain why this occurs. We furthermore establish a simple limit result for large voter populations.

Suppose that the voter distribution is symmetric and that, apart from the neutral group, on each side of the spectrum there is a group of very biased voters and a group of moderate voters. The key is that the neutral group, although being the biggest voter group, is not large enough to discipline voters neither in duopoly nor in triopoly. In duopoly the two networks choose each their side of the market which they dominate whereas in triopoly one network is neutral and the two biased networks must choose

between targeting the moderate or the extreme group. Since the former group is larger they choose a moderate bias. This results in the following counterexample to the result of the previous section, which is proved in the appendix.

Example 1. *Suppose $K = 2$ and $G(b_1, 0) > G(b_1, \beta^0(b_2))$, $G(b_2, \beta^*(b_1)) < 0$ and $\pi_0/2 < \pi_2 < \pi_1 < \pi_0$. Then if a biased voter subscribes to news he always finds the media more extreme in duopoly than in triopoly. In particular,*

$$\beta^0(b_{-2}) = \beta_L^D < \beta_L^T < \beta_M^T = 0 < \beta_R^T < \beta_R^D = \beta^0(b_2).$$

In this case, the smaller news market bias in triopoly comes at a price: extremists will not subscribe to news. Thus they will always vote according to their prior. As mentioned in the previous section, a social planner comparing the two situations must weigh the bias of the moderate groups against the participation of the extreme groups. The difference to the result of the previous section is due to the following: as soon as one of the networks choose the central position ($\beta = 0$) the other networks must choose between attracting moderates or extremists and hence the marginal argument of Lemma 2 has no effect. Two arguments make the result of the previous section more realistic in our view. First, the result of this section relies very much on the discreteness of the voter distribution. In reality, we believe that voter distributions are more resemblant of continuous distributions. Second, the distributional assumptions of the example do not match the data of the previous section. In particular, the extreme groups are smaller in the data than required for the example. Using distributional assumptions that reflect the data the duopolists would move toward the center to fight for the large neutral group and the result of the example would not hold.

Having established that the bias of the news market is not monotonically changing in the number of networks we investigate the limit case of a large market with many news providers. If the number of networks is large enough the only equilibrium is that each voter group is covered by at least one network that has exactly the preferred bias of that group. The reason is simple. When there are sufficiently many networks the best each one can do is to focus on one voter group and the optimal bias level for each network is then the preferred bias of its target group. If not the preferred bias was chosen, another

network could steal this group entirely. We formalize this in the following proposition. Let $\underline{\pi} \equiv \min \pi_k$ and $n_k = \min\{n \in \mathbb{N} : \pi_k/n \leq \underline{\pi}\}$.

Proposition 4. *Suppose the voter distribution is discrete with $N \geq \sum n_k$. Then the news market is completely covered. I.e. in equilibrium, for any k there exists $n \in \mathbb{N}$ such that $\beta_n = \beta_k^*$.*

In the limit the bias of the news market is exactly equal to the bias of the voters. Recall that with fewer networks it is possible that the news market is completely unbiased. Thus it is not obvious that the preferred market structure from the point of view of a social planner is more competition. However, voter welfare is clearly maximized in the limit case as the following corollary shows.

Corollary 3. *Suppose the voter distribution is discrete with $N \geq \sum n_k$. Then the welfare of all voter groups is maximized. I.e. for all k then $W_k = \max_{\beta} G_k(\beta) > 0$.*

Since each voter group has access to their preferred network-type they choose to subscribe and obtain their maximum expected gain. Hence, voters would always prefer the limit case of the very competitive news market. We could think of the very competitive market as an approximation to the US news market in which a large number of nationally available networks provide news coverage from a variety of different political points of view. The smaller markets are more akin of most European countries in which news markets are dominated by a large public network and a smaller number of private networks.

1.5 Demand-Side Bias: Adding Journalist Uncertainty and Voter Stubbornness to the Model

Until now we have considered competition as the driving force of news bias and assumed that voters and journalists differ only in their prior expectation of the state of the world. However, it is also possible that news bias is created on the demand side. In this section we explore two channels through which this can happen and investigate whether increased demand for biased news also leads to more biased news in the market

equilibrium. We focus on discrete distributions and let $G_k(\beta) \equiv G(b_k, \beta)$ and $l_k(\cdot) \equiv l_{b_k}(\cdot)$.

First we consider uncertainty about journalist bias. Voters may not always know the exact bias of journalists but rather have a more or less precise idea about it. There could be many reasons for this. For instance, The New York Times employs many journalists and each of them have a different bias. Readers will probably not know the bias of each journalist but have a good idea about the average. Alternatively, each journalist may have a different bias on different issues and only the average bias is observable. The precise technology of the uncertainty is not important, we merely assume that such uncertainty exists and show that this leads voters to demand more extreme news networks. The intuition is that the uncertainty adds more noise to the messages that networks send and therefore voters place less weight on them and need a stronger signal in order to switch away from their a priori preferred action.

Second we increase the heterogeneity between voters and journalists by allowing them to differ in the precision of their prior, which we can interpret as their confidence or stubbornness. The more stubborn voters are, the less prone they are to changing their opinion. Therefore their optimal network bias is more extreme since more biased networks provide stronger signals.

1.5.1 Journalist Bias Uncertainty

In this section we allow for uncertainty about the bias of the journalist. Voters and networks must make decisions based on the observable average bias of the journalist and their knowledge of the distribution of her true bias. In particular, we assume that the bias distribution of a type- $\bar{\beta}$ journalist is

$$\beta \sim N(\bar{\beta}, 1/\phi). \quad (1.4)$$

The parameter $\bar{\beta}$ is the expected bias of the journalist and ϕ is the precision of the distribution from which β is drawn. This is common knowledge between all players. To analyze news demand in this setting we attack the problem a bit differently. Conditional on x the journalist has posterior expectation $\mu = (\beta + x)/2$ and hence $x = 2\mu - \beta$. Let $z \equiv 2\mu - \bar{\beta} = x - (\bar{\beta} - \beta)$. Then z is normally distributed with expectation

θ and precision $\phi/(1 + \phi)$. Hence, type- k voters' posterior beliefs over θ conditional on z are normally distributed with expectation $(b_k + \phi/(1 + \phi)z)/(1 + \phi/(1 + \phi)) = ((1 + \phi)b_k + \phi(\mu - \bar{\beta}))/ (1 + 2\phi)$ and precision $1 + \phi/(1 + \phi)$. Furthermore, a type- k voter a priori expects z to be distributed normally with expectation b_k and precision $\phi/(2\phi + 1)$. Denote the distribution of type- k voters' beliefs about z by $q_k(\cdot)$. Since the journalist sends the signal $m = -\gamma$ whenever $x < \beta$, which corresponds to $z < -\bar{\beta}$, we can write the gain function as

$$G_k(\bar{\beta}) = -4\gamma \int_{-\infty}^{-\bar{\beta}} \frac{(1 + \phi)b_k + \phi z}{1 + 2\phi} q_k(z) dz. \quad (1.5)$$

This function is very similar to the gain function of the simple version of the model except it depends only on $\bar{\beta}$ and the uncertainty parameter ϕ , and not the true value of β . Notice that the weight of the voter's prior is increasing in uncertainty. We then derive the following observation.

Observation 4. *Suppose journalist bias is unobserved but commonly known to be distributed according to (1.4). The optimal expected journalist type for a type- k voter is*

$$\beta_k^* = \frac{1 + \phi}{\phi} b_k.$$

Notice from the observation that *whenever there is uncertainty voters prefer journalists that are more biased than themselves*. The intuition is the following. In the simple version of the model voters could make a direct inference about x conditional on the message received from the journalist. However, when uncertainty is present voters must instead make a similar inference about the variable z . But as we have shown above z is a more noisy signal about the state of the world than x (z has precision $\phi/(1 + \phi) < 1$) and therefore the voter assigns lower weight to this message. Consequently he needs a stronger signal to switch away from his a priori preferred action and for this he needs a journalist who he thinks is more extremely biased on average. As an example, if the voter is conservative but finds it hard to determine whether journalists are too liberal or too conservative he prefers a journalist that he expects to be more conservative than himself such that if he receives the message “invest in renewable energy” he feels confident that this message is due to the information contained in the scientific report x and

not due to the journalist's bias. We next establish a lemma that will aid in determining the comparative statics.

Lemma 3. *Suppose journalist bias is unobserved but commonly known to be distributed according to (1.4). Then for all voter types, the least informative journalist-type is increasing in uncertainty for moderate voters and decreasing in uncertainty for extreme voters.*

In particular, suppose $\tilde{\phi} < \phi$ (more uncertainty under $\tilde{\phi}$) and denote by $\tilde{\beta}_k^0$ and β_k^0 the respective least informative journalist types of a type- k voter. Suppose $b_k \geq 0$. Then there exists $\underline{b} > 0$ such that for $b_k < \underline{b}$

$$\tilde{\beta}_k^0 < \beta_k^0.$$

Furthermore, there exists $\bar{b} > 0$ such that for $b_k > \bar{b}$ then

$$\tilde{\beta}_k^0 > \beta_k^0.$$

Conversely for $b_k \leq 0$.

The lemma follows as increasing uncertainty produces two effects of opposite sign. First the voter places more weight on his prior and less on the information derived from the news message which diminishes his expected gain. Second his posterior after receiving the message moves away from his prior, which augments his expected gain. When the voter is extreme the first effect dominates and it is less likely that he subscribes to news. When the voter is moderate the second effect dominates and it is more likely that he subscribes. Recall that $k^* \equiv \max\{k \in \mathbb{N} : \pi_k \geq \pi_0/2\}$. We can then directly derive the following proposition.

Proposition 5. *Suppose journalist bias is unobserved but commonly known to be distributed according to (1.4) and that the population is symmetric. Let $\tilde{\phi} < \phi$, such that there is more uncertainty under $\tilde{\phi}$, and denote by $-\tilde{\beta}_L = \tilde{\beta}_R$ and $-\beta_L = \beta_R = \beta_{k^*}^0$ the respective symmetric equilibria. If $b_{k^*} < \underline{b}$ then equilibrium bias is decreasing in uncertainty. I.e.*

$$\beta_L^D < \tilde{\beta}_L^D \text{ and } \tilde{\beta}_R^D < \beta_R^D.$$

On the other hand, if $b_{k^} > \bar{b}$ then equilibrium bias is increasing in uncertainty. I.e.*

$$\tilde{\beta}_L^D < \beta_L^D \text{ and } \beta_R^D < \tilde{\beta}_R^D.$$

Despite the unambiguous effect of journalist uncertainty on the optimal news bias demanded by voters the effect on the equilibrium may go either way. The reason for this is to be found in Lemma 3, which shows that the unambiguousness of the optimal journalist type does not carry over to the effect on the least informative journalist type. This may become more or less extreme as uncertainty grows and since the least informative journalist-type (rather than the optimal journalist-type) determines the market equilibrium there is ambiguity. If the marginal voter (the most biased voter who subscribes to news) is relatively moderate an increase in journalist bias uncertainty will imply that the marginal voter pays more attention to news in general and hence networks can move closer to the center without losing their marginal audience, which they will do in order to compete for the neutral voters. On the other hand, if the marginal voter is relatively extreme an increase in journalist bias uncertainty implies that the marginal voter pays less attention to news and hence networks become more biased in order to attract him.

1.5.2 Voters of Different Prior Precision

So far voters and journalists have differed only in their expectation of the state but in this section we want to explore how the situation changes if they also have different levels of prior precision. Imagine two different voters who both think that on average there is some degree of climate change occurring but not enough for us to need to invest in renewable energy. One of them may feel very unsure about his beliefs and in reality places equal probability on the events of large, moderate or small climate change. We can think of this voter as unconfident or open-minded. The other voter on the contrary feels very sure about his beliefs and places very low probabilities on any other interpretation of events. This voter may be thought of as confident or stubborn. The question is whether these two voters have the same preferences for news and how the composition of confidence in the voter population affects the competitive equilibrium. To investigate

this we assume that the prior beliefs of a type- k voter are distributed according to

$$\theta \stackrel{k}{\sim} N(b_k, 1/p_k). \quad (1.6)$$

Hence, b_k remains k 's prior expectation of θ and $p_k \in \mathbb{R}_+$ is the precision of his prior beliefs, which measures his confidence or stubbornness. The precision of journalists' beliefs remains one and hence if $p_k < 1$ the voter is open-minded and if $p_k > 1$ he is stubborn. Conditional on x a type- k voter has posterior expectation

$$\mu_k(x) = \frac{p_k b_k + x}{p_k + 1}. \quad (1.7)$$

The precision parameter acts as a weight on his prior expectation. Following the same steps as previously we can derive the gain function of a type- k voter.

$$G_k(\beta) = -4\gamma \int_{-\infty}^{-\beta} \frac{p_k b_k + x}{p_k + 1} l_k(x) dx. \quad (1.8)$$

Notice that p_k will also affect $l_k(\cdot)$, voter k 's beliefs about the distribution of x , but for the purpose of identifying the optimal journalist bias this does not matter. As can be seen directly from (1.8) the gain is increasing for $x < -p_k b_k$, which leads to the following observation.

Observation 5. *Suppose voter beliefs are given by (1.6). The optimal journalist type for a type- k voter is*

$$\beta_k^* = p_k b_k.$$

The observation identifies the optimal journalist bias as being the voter's own bias modified by his degree of stubbornness and hence he always prefers to get his news from a source that has the same direction of bias as himself. If the voter is stubborn he prefers a journalist who is more biased than himself. Conversely, if he is open-minded he prefers the journalist to be more moderate. The following lemma shows that this comparative static also carries over to the least informative journalist type.

Lemma 4. *Suppose that voters potentially have different prior precisions such that their beliefs are distributed according to (1.6). Then the bias of the least informative journalist type is increasing in voter stubbornness.*

In particular, let $\tilde{p}_{k'} > p_{k'}$ for some k' with everything else equal and denote by $\tilde{\beta}_k^0$ and β_k^0 the respective least informative journalist types of a type- k voter. Suppose $b_k > 0$. Then

$$\tilde{\beta}_k^0 > \beta_k^0.$$

Conversely for $b_k < 0$.

We observed earlier that the prior precision affects the voter's beliefs about the distribution of x but whereas this was not important for the optimal journalist type it plays a role for the least informative journalist type. The more confident k is the more tightly he expects x to be distributed around b_k . At the same time he will place more weight on his own prior beliefs and these two effects will both decrease his gain. This together with (M) implies that news sources will have to move closer to the preferences of high confidence voters in order to attract them, which yields the result. As we know that the least informative journalist type conditions the duopoly market equilibrium this lemma leads the following proposition.

Proposition 6. *Suppose that voters potentially have different prior precisions such that their beliefs are distributed according to (1.6) and that the population is symmetric. Let $\tilde{p}_{k'} > p_{k'}$ for some k' with everything else equal and denote by $-\tilde{\beta}_L = \tilde{\beta}_R$ and $-\beta_L = \beta_R$ the respective symmetric equilibria. In equilibrium media bias is more extreme the more stubborn voters are. I.e.*

$$\tilde{\beta}_L^D < \beta_L^D \text{ and } \beta_R^D < \tilde{\beta}_R^D.$$

The effect of increasing voter stubbornness is to unambiguously increase the news market bias in duopoly. This is consistent with the intuition that stubborn voters are more likely to seek self-confirming news sources only that here the result is derived as an optimal choice on part of the voters rather than being assumed as an intrinsic preference. In duopoly it is the most extreme audience group who determines the equilibrium and therefore if bias and stubbornness are correlated moderate voters suffer even further.

1.6 Conclusion

This article has sought explanations for why moderately biased voters often subscribe to very biased news shows. Our results suggest that this is natural in a setting where news are purely observed for their informational value (not for entertainment) and where journalists are conscientious in the sense that they report the news honestly according to their own beliefs. The key mechanism that brings about this link is that moderates have little ability discipline networks. Hence, networks will stage their news shows to cater to the tastes of the most biased voters and in doing so they will also capture more moderate voters. The result is that most voters will tend to observe a news source that is more biased than themselves. Although market bias does not change monotonically with the number of networks we observe that in the limit as the number of networks grows large the market bias converges to the prior bias of the voters. We also suggest possible demand-side factors that can drive the disconnect between voter and news show bias. When there is uncertainty about journalist bias or when voters are very stubborn we observe that moderates prefer more biased news shows. In the first case however, even though voters' optimal news source bias is increasing in uncertainty the effect on the market is ambiguous.

1.7 Proofs

Proof of Lemma 1. Let $L_b(a) \equiv \int_{-\infty}^a l_b(s) ds$ and $l_b(x|a) \equiv l_b(x)/L_b(a)$ such that we can write $G(\cdot)$ as

$$G(b, \beta) = -2\gamma L_b(-\beta) \left[b + \int_{-\infty}^{-\beta} x l_b(x|-\beta) dx \right].$$

Suppose $G(b, \beta) = 0$ and choose $b' < b$. Then

$$\begin{aligned} G(b, \beta) &= -2\gamma L_b(-\beta) \left[b + \int_{-\infty}^{-\beta} x l_b(x|-\beta) dx \right] \\ &= -2\gamma L_{b'}(-\beta) \left[b + \int_{-\infty}^{-\beta} x l_b(x|-\beta) dx \right] \\ &< -2\gamma L_{b'}(-\beta) \left[b' + \int_{-\infty}^{-\beta} x l_b(x|-\beta) dx \right] \\ &< -2\gamma L_{b'}(-\beta) \left[b' + \int_{-\infty}^{-\beta} x l_{b'}(x|-\beta) dx \right] \\ &= G(b', \beta). \end{aligned}$$

The second equality follows because $G(b, \beta) = 0$ implies that the term in the bracket is zero. The first inequality follows since $b' < b$ and the second since the integral is the conditional expectation of x given $x < -\beta$ and this is increasing in b . \square

Proof of Proposition 1. Suppose $\beta_R^M = \bar{\beta}$. If $\beta_L^M > \underline{\beta}$, then by definition of $\underline{\beta}$ the monopolist has market share $s < 1$. On the other hand, as a direct consequence of Lemma 1, if he sets $\beta_L^M = \underline{\beta}$ he will have market share $s = 1$. The same argument applies for $\beta_R^M < \bar{\beta}$. \square

Proof of Proposition 2. As is argued in the text, when the voter distribution is continuous it will never be optimal to set $\beta_L^D < \underline{\beta}$ since in this case L could move towards the center and gain market share, since $b_1(\beta_L, \beta_R) > 0$, without loosing any market share in the extreme as long as $\beta_L^D \leq \underline{\beta}$. Hence, in an equilibrium we must always

have $\beta_L^D \geq \underline{\beta}$. Similarly for R . To see the discrete case, notice that $\underline{\beta} = -\bar{\beta}$ for symmetric voter distributions. Suppose $\beta_R^D = -\beta_L^D > \bar{\beta}$. Then L and R split the central group. But by moving slightly toward the center, either L or R could gain the entire central group without loosing any other groups. Hence, this cannot be an equilibrium. If $\beta_R^D > -\beta_L^D \geq \bar{\beta}$ then L has the entire central group and R could gain $\pi_0/2$ by setting $\beta_R^D = -\beta_L^D$. Hence, this cannot be an equilibrium either. The same argument holds if $-\beta_L^D > \beta_R^D \geq \bar{\beta}$. Thus, we must have $\underline{\beta} \leq \beta_L^D \leq \beta_R^D \leq \bar{\beta}$. This yields the result. \square

Proof of Corollary 1. Observe that if $\beta_L^D = -\beta_{k^*}^0$ and $\beta_R^D = \beta_L^D = -\beta_{k^*}^0$, a deviation towards the center can at most give them $\pi_{k^*} - \pi_0/2$. This yields the result. \square

Proof of Corollary 2. First we state and prove the following lemma.

Lemma 5. *The indifferent voter function $b(\beta_L, \beta_R)$ always exists and is increasing in both arguments. I.e. the indifferent voter is unique for all β_L and β_R and $b_1(\cdot) > 0$ and $b_2(\cdot) > 0$.*

Proof. Notice that

$$G(b, \beta_L) - G(b, \beta_R) = -2\gamma L_b(-\beta_R, -\beta_L) \left[b + \int_{-\beta_R}^{-\beta_L} x l_b(x | -\beta_R, -\beta_L) dx \right].$$

The proof of uniqueness of b can then be completed by considering b such that $G(b, \beta_L) - G(b, \beta_R) = 0$ and following the same steps as in Lemma 1 to show that $G(b', \beta_L) < G(b, \beta_R)$ for $b' > b$ and $G(b'', \beta_L) > G(b, \beta_R)$ for $b'' < b$. Furthermore, since $\beta_L \leq b(\beta_L, \beta_R) \leq \beta_R$ then $G_2(b, \beta_L) > 0$ and $G_2(b, \beta_R) < 0$. This together with the above result implies that if $\beta'_L \geq \beta_L$ and $\beta'_R \geq \beta_R$ with at least one of the inequalities holding strict, then $b(\beta'_L, \beta'_R) > b(\beta_L, \beta_R)$. \square

If the equilibria are not identical then by Proposition 2 either $\beta_L^M \leq \beta_L^D$ or $\beta_R^D \leq \beta_R^M$ or both hold with strict inequality. Consider β_R^M and β_R^D . From Lemma 5 we know that $\beta_R^M \geq b(\beta_R^D, \beta_R^M) \geq \beta_R^D \geq 0$. Hence, $b'' = b(\beta_R^D, \beta_R^M)$. Similar for β_L^M and β_L^D . \square

Proof of Lemma 2. Let $D(\beta_L, \beta_R, b) \equiv G(b, \beta_L) - G(b, \beta_R)$ and $b_0 = b(\beta_L, \beta_R)$. Then by definition $D(\beta_L, \beta_R, b_0) = 0$. We denote the partial derivatives of $D(\cdot)$ by $D_i \equiv D_i(\beta_L, \beta_R, b_0)$ for $i = 1, 2, 3$. Using the notation $l_L = l(-\beta_L - b_0)$, $l_R = l(-\beta_R - b_0)$ and $L = \int_{-\beta_R - b_0}^{-\beta_L - b_0} l(z) dz$ we can write these derivatives as follows.

$$\begin{aligned} D_1 &= 2\gamma(b_0 - \beta_L)l_L \\ D_2 &= 2\gamma(b_0 - \beta_R)l_R \\ D_3 &= D_1 - D_2 - 4\gamma L. \end{aligned}$$

The first derivative of $b(\cdot)$ with respect to β_R can then be obtained by implicit differentiation.

$$b_2(\beta_L, \beta_R) = -\frac{D_2}{D_3} = \frac{(\beta_R - b_0)l_R}{2L - [(b_0 - \beta_L)l_L + (\beta_R - b_0)l_R]}.$$

Since the factor 2γ cancels out it is convenient to use the shorthand $d_i = D_i/(2\gamma)$, such that the derivative above can be written $-d_1/d_3$. The second derivative can be computed by the chain rule.

$$b_{21}(\beta_L, \beta_R) = b_{21}(\beta_L, \beta_R)|_{b=b_0} + b_1(\beta_L, \beta_R) \times \frac{\partial^2 b(\beta_L, \beta_R)}{\partial \beta_R \partial b_0}.$$

As should be clear, the first term measures the effect of changing β_L when keeping b_0 constant and the second term captures the change in b_0 and its effect. First we notice that if h is the density function of a normal distribution with variance σ^2 its derivative has the following property $h'(z) = -h(z)z/\sigma^2$. It follows that $(\beta_R - b_0)l'_R = (\beta_R^2 - b_0^2)l_R/2$ and $(\beta_R - b_0)l'_R = (\beta_R^2 - b_0^2)l_R/2$. We then calculate the derivatives.

$$\begin{aligned} b_1(\beta_L, \beta_R) &= -\frac{d_1}{d_3} = -\frac{(b_0 - \beta_L)l_L}{d_3} \\ b_{21}(\beta_L, \beta_R)|_{b=b_0} &= -\frac{l_L l_R}{d_3^2} \left(1 + \frac{b_0^2 - \beta_L^2}{2} \right) (\beta_R - b_0) \\ \frac{\partial^2 b(\beta_L, \beta_R)}{\partial \beta_R \partial b_0} &= \frac{l_L}{d_3^2} \left[\left(1 + \frac{\beta_R^2 - b_0^2}{2} \right) d_3 - (\beta_R - b_0)d_{33} \right], \end{aligned}$$

where d_{33} is the derivative of d_3 with respect to b_0 , i.e. $d_{33} = D_{33}/(2\gamma)$. This is equal to

$$d_{33} = l_R \cdot \left(3 + \frac{\beta_R^2 - b_0^2}{2}\right) - l_L \cdot \left(3 - \frac{b_0^2 - \beta_L^2}{2}\right).$$

Putting this together we arrive at the following.

$$b_{21}(\beta_L, \beta_R) = \frac{l_L l_R}{d_3^3} [(\beta_R - b_0)(b_0 - \beta_L)d_{33} - B d_3],$$

where

$$B = \left(1 + \frac{b_0^2 - \beta_L^2}{2}\right)(\beta_R - b_0) + \left(1 + \frac{\beta_R^2 - b_0^2}{2}\right)(b_0 - \beta_L).$$

First, notice that since $b(-\beta_R, \beta_R) = 0$, $b_1(\cdot) > 0$ and $\beta_R \geq |\beta_L|$ then we must have $b_0 \geq 0$. Hence, $\beta_R - b_0 \geq b_0 - \beta_L \geq 0$ by Observation 3 and $(\beta_R - b_0)/(\beta_R - \beta_L) \geq 1/2 \geq (b_0 - \beta_L)/(\beta_R - \beta_L)$.

To determine the sign of the derivative we need to use the property that if h is the density function of a normal distribution with variance σ^2 its second derivative is $h''(z) = [(y/\sigma^2)^2 - 1]h(z)$. Hence, the distribution function is concave whenever $|y| \leq \sigma^2$. Hence, given the assumption that $b \in [-1, 1]$ we must have $b_0, \beta \in [-1, 1]$ and therefore $l(\cdot)$ is concave over the interval $[-\beta_R - b_0, -\beta_L - b_0]$.

Using the concavity property we can establish that $L > (\beta_R - \beta_L)(\frac{1}{2}l_L + \frac{1}{2}l_R) \geq (b_0 - \beta_L)l_L + (\beta_R - b_0)l_R$. This allows us to upper bound d_3 .

$$d_3 = [(b_0 - \beta_L)l_L + (\beta_R - b_0)] - 2L < -L < -(\beta_R - \beta_L)l_L/2 < 0.$$

Furthermore, since $\beta_R \geq |\beta_L|$ then $\beta_R^2 \geq \beta_L^2$ and therefore

$$d_{33} \geq -\left(3 + \frac{b_0^2 - \beta_L^2}{2}\right)l_L.$$

Combining this we can upper bound $b_{21}(\cdot)$ by $l_L^2 l_R / d_3^3 (\beta_R - \beta_L)$ times

$$\begin{aligned}
& B - \frac{(\beta_R - b_0)(b_0 - \beta_L)}{\beta_R - \beta_L} \left(3 + \frac{\beta_R^2 - b_0^2}{2} \right) \\
&= (\beta_R - b_0) \left(1 + \frac{b_0^2 - \beta_L^2}{2} - 2 \frac{b_0 - \beta_L}{\beta_R - \beta_L} \right) \\
&\quad + (b_0 - \beta_L) \left(1 + \frac{\beta_R^2 - b_0^2}{2} - \frac{\beta_R - b_0}{\beta_R - \beta_L} \left(1 + \frac{b_0^2 - \beta_L^2}{2} \right) \right) \\
&\geq (\beta_R - b_0) \left(\frac{b_0^2 - \beta_L^2}{2} \right) + (b_0 - \beta_L) \left(\frac{\beta_R^2 - b_0^2}{2} - \frac{\beta_R - b_0}{\beta_R - \beta_L} \frac{b_0^2 - \beta_L^2}{2} \right) \\
&\geq 0.
\end{aligned}$$

The first inequality holds since $1 \geq (\beta_R - b_0)/(\beta_R - \beta_L) \geq 1/2 \geq (b_0 - \beta_L)/(\beta_R - \beta_L)$. The second holds, since either $b_0^2 - \beta_L^2 \geq 0$ and in this case $\beta_R^2 - b_0^2 \geq b_0^2 - \beta_L^2$ and both terms are positive. Or $b_0^2 - \beta_L^2 < 0$, in which case $(\beta_R - b_0)(b_0^2 - \beta_L^2) + (b_0 - \beta_L)(\beta_R^2 - b_0^2) = (\beta_R - b_0)(b_0 - \beta_L)(2b_0 + \beta_L + \beta_R) \geq 0$. \square

Proof of Example 1. We need to check duopoly against triopoly.

Duopoly. In duopoly $\beta_L^D = \beta_{-2}^0$ and $\beta_R^D = \beta_2^0$, which yields $s_L = s_R = \pi_2 + \pi_1 + \pi_0/2$. There is no incentive to move towards the center to capture the entire neutral group since $\pi_0/2 < \pi_2$.

Triopoly. On the other hand, in triopoly $\beta_L^T = \beta_{-1}^*$, $\beta_M = 0$ and $\beta_R^T = \beta_1^*$ which gives $s_L = s_R = \pi_1$ and $s_M = \pi_0$. By deviating to the left L can obtain π_2 and by deviating to the right he can obtain $\pi_0/2$. Since $\pi_1 > \pi_2 > \pi_0/2$ none of these deviations are profitable. Idem for R . M can either obtain $\pi_1/2$ or π_2 . Neither of these are greater than π_0 and therefore M has no profitable deviation. \square

Proof of Proposition 4. Suppose that for some k' there is no $n \in \mathbf{N}$ such that $\beta_n = \beta_{k'}^*$. By assumption, some network n' must have market share $s_{n'} < \underline{\pi} \leq \pi_k$ and hence $\beta_{n'} = \beta_{k'}^*$ constitutes a profitable deviation. This is a contradiction. \square

Proof of Corollary 3. The result follows since for all k , $\beta_n = \beta_k^*$ for some $n \in \mathbf{N}$ and β_k^* maximizes the gain of group k voters. \square

Proof of Observation 4. Differentiate (1.5) with respect to $\bar{\beta}$.

$$G'_k(\bar{\beta}) = 4\gamma \left[\frac{(1 + \phi)b_k - \phi\bar{\beta}}{1 + 2\phi} \right] q_k(-\bar{\beta}).$$

The observation follows directly. \square

Before proving the results of Section 1.5 let us first establish a useful lemma.

Lemma 6. *Suppose $H(t)$ and $\tilde{H}(t)$ are the probability distribution functions of two normal distributions with the same expectation μ and precision τ and $\tilde{\tau}$, respectively, where $\tau < \tilde{\tau}$.*

$$\tilde{H}(t|c, d) < H(t|c, d) \quad \text{whenever } c < t \leq d < \mu \quad (1.9)$$

$$\tilde{H}(c) < H(c) \quad \text{for all } c < \mu \quad (1.10)$$

$$\tilde{H}(\mu - \epsilon, \mu + \epsilon) > H(\mu - \epsilon, \mu + \epsilon) \quad \text{for all } \epsilon > 0 \quad (1.11)$$

Proof. Let $h_\tau(\cdot)$ denote the respective density functions and notice that

$$\frac{h_\tau(t_2)}{h_\tau(t_1)} = e^{-\tau[(t_2 - \mu)^2 - (t_1 - \mu)^2]/2}. \quad (1.12)$$

If $t_1 < t_2 < \mu$ it follows from (1.12) that $h_\tau(t_2)/h_\tau(t_1)$ is increasing in τ , implying that $h_\tau(\cdot)$ has the MLR-property in τ for $t < \mu$. This property implies (1.9) and (1.10). Last, (1.11) follows from (1.10). \square

Proof of Lemma 3. Suppose $b_k \geq 0$ and $\tilde{\phi} < \phi$. Let $q_k(\cdot | -\bar{\beta})$ denote the density of k 's beliefs over z conditional on $z < -\bar{\beta}$ on ϕ and let $\tilde{q}_k(\cdot | -\bar{\beta})$ denote the same density conditional on $\tilde{\phi}$. Similarly let $\tilde{G}_k(\cdot)$ and $G_k(\cdot)$ be evaluated under $\tilde{\phi}$ and ϕ . Let $\bar{\beta}'$ satisfy $G_k(\bar{\beta}') = 0$. Multiplying this condition by $-4\gamma(1+2\phi)/\phi$ on both sides we arrive at

$$\frac{1+\phi}{\phi}b_k + \int_{-\infty}^{-\bar{\beta}'} z q_k(z | -\bar{\beta}') dz = 0. \quad (1.13)$$

Since $b_k \geq 0$ the first term in (1.13) is non-negative and hence the integral must be non-positive. As $b_k \rightarrow 0$ then

$$\int_{-\infty}^{-\bar{\beta}'} z \tilde{q}_k(z | -\bar{\beta}') - \int_{-\infty}^{-\bar{\beta}'} z q_k(z | -\bar{\beta}') \rightarrow \epsilon.$$

By Lemma 6 then $\epsilon < 0$. On the other hand, as $b_k \rightarrow 0$ then $[(1+\tilde{\phi})/\tilde{\phi} - (1+\phi)/\phi]b_k \rightarrow 0$. Hence, by the continuity of all the terms there exists \underline{b} such that the left-hand side of (1.13) is negative for $b_k < \underline{b}$ and hence $\tilde{G}_k(\bar{\beta}') > 0$. Let $\bar{\beta}''$ satisfy $\tilde{G}_k(\bar{\beta}'') = 0$ and notice that $\bar{\beta}' < \beta_k^* < \tilde{\beta}_k^*$ and $\bar{\beta}'' < \tilde{\beta}_k^*$. Therefore, by (M), then $\tilde{G}'_k(\beta) > 0$ for all β between $\bar{\beta}'$ and $\bar{\beta}''$. Thus $\bar{\beta}'' < \bar{\beta}'$. This establishes the first part of the lemma.

As $b_k \rightarrow \infty$ then for any $z_1 < z_2 < b_k$ or $b_k < z_2 < z_1$ we have $q_k(z_2)/q_k(z_1) \rightarrow \infty$ and thus $q_k(\cdot | -\bar{\beta}')$ converges to a distribution which is degenerate at $z = -\bar{\beta}'$. Hence, as $b_k \rightarrow \infty$ then

$$\int_{-\infty}^{-\bar{\beta}'} z \tilde{q}_k(z | -\bar{\beta}') - \int_{-\infty}^{-\bar{\beta}'} z q_k(z | -\bar{\beta}') \rightarrow -\bar{\beta}' - (-\bar{\beta}') = 0.$$

On the other hand, as $b_k \rightarrow \infty$ then $[(1+\tilde{\phi})/\tilde{\phi} - (1+\phi)/\phi]b_k \rightarrow \infty$. Hence, by the continuity of all the terms there exists \bar{b} such that the left-hand side of (1.13) is positive for $b_k > \bar{b}$ and hence $\tilde{G}_k(\bar{\beta}') < 0$. Thus if $\bar{\beta}''$ satisfies $\tilde{G}_k(\bar{\beta}'') = 0$ then $\bar{\beta}'' > \bar{\beta}'$ by the same argument as above. This establishes the second part of the lemma. The proof is identical for $b < 0$. \square

Proof of Proposition 5. The proposition follows directly from Corollary 1 and Lemma 3. \square

Proof of Observation 5. Differentiate (1.8) with respect to β .

$$G'_k(\beta) = 4\gamma\mu_k(-\beta)l_k(-\beta) = 4\gamma\frac{pb_k - \beta}{p+1}l_k(-\beta). \quad (1.14)$$

The observation follows directly. \square

Proof of Lemma 4. Let $L_k(c, d)$ be the probability assigned by voter k to the event that $x \in [c, d]$ and let $l_k(x|c, d)$ for $x \in [c, d]$ denote the density of the corresponding conditional distribution of the voter's beliefs over x . Furthermore, let $L_k(d) \equiv \lim_{c \rightarrow -\infty} L_k(c, d)$ and $l_k(x|d) \equiv \lim_{c \rightarrow -\infty} l_k(x|c, d)$. Suppose $G(\beta) \geq 0$. The two first properties of Lemma 6 tell us that the conditional distributions can be ranked by first order stochastic dominance in p . Let $\tilde{l}(\cdot)$ be equal to $l_k(\cdot)$ evaluated under \tilde{p}_k and $l(\cdot)$ equal to $l_k(\cdot)$ evaluated under p_k , and similarly for $L(\cdot)$ and $\mu(\cdot)$. Suppose $b > 0$. There are two cases. First, suppose $-\beta \leq b$ and drop the k subscript. Then

$$\begin{aligned} \tilde{G}(\beta) &= -4\gamma\tilde{L}(-\beta) \int_{-\infty}^{-\beta} \tilde{\mu}(x)\tilde{l}(x|-\beta) dx \\ &< -4\gamma\tilde{L}(-\beta) \int_{-\infty}^{-\beta} \tilde{\mu}(x)l(x|-\beta) dx \\ &< -4\gamma\tilde{L}(-\beta) \int_{-\infty}^{-\beta} \mu(x)l(x|-\beta) dx \\ &\leq -4\gamma L(-\beta) \int_{-\infty}^{-\beta} \mu(x)l(x|-\beta) dx \\ &= G(\beta) = 0. \end{aligned}$$

The first inequality follows from (1.10) and the fact that $\tilde{\mu}(x)$ is an increasing function, the second since $\mu(x)$ is increasing in p for $b \geq x$ and the equality from (1.10) and $G(\beta) \geq 0$. Hence, for any news source with $-\beta \leq b$ and a voter who has non-negative expected gain, the expected gain of the voter is decreasing in his prior precision, p . To take care of the second case, let $-\beta > b > 0$. Define $\beta' \equiv 2b - \beta > 0$ and notice that β and β' are symmetric around b , i.e.

$$-(-\beta' - b) = -\beta - b > 0. \quad (1.15)$$

Then

$$\begin{aligned}
\tilde{G}(\beta) &= \tilde{G}(\beta') - 4\gamma\tilde{L}(-\beta', -\beta) \int_{-\beta'}^{-\beta} \tilde{\mu}(x)\tilde{l}(x|-\beta', -\beta) dx \\
&= \tilde{G}(\beta') - 4\gamma\tilde{L}(-\beta', -\beta)b \\
&< G(\beta') - 4\gamma\tilde{L}(-\beta', -\beta)b \\
&< G(\beta') - 4\gamma L(-\beta', -\beta)b \\
&= G(\beta).
\end{aligned}$$

The integral in the first line is equal to b by (1.15) and the fact that $\tilde{l}(\cdot|-\beta', -\beta)$ is symmetric around b . The first inequality follows from the first case (since $-\beta' \leq b$), and the second from (1.11) and (1.15). Thus we have established that for any b such that the voter subscribes to news, the gain $G(\beta)$ is decreasing in p . Hence, if β' and β'' satisfy $G(\beta') = 0$ and $\tilde{G}(\beta'') = 0$ then $\beta'' > \beta'$. The proof is identical for $b < 0$. \square

Proof of Proposition 6. There are three cases. First suppose $k' > k^*$. In this case we know from Corollary 1 that $\pi_{k'} < \pi_0/2$ and hence changing $p_{k'}$ does not influence the equilibrium. Second, suppose $k' = k^*$. From Lemma 4 we know that for $k' \geq 0$ then $\tilde{\beta}_{k'}^0 > \beta_{k'}^0$ and hence $\tilde{\beta}_L^D < \beta_L^D$ and $\tilde{\beta}_R^D > \beta_R^D$. Last, suppose $k' < k^*$. If $\pi_{k'} < \pi_0/2$ then the equilibrium always remains the same by Corollary 1. Suppose on the contrary $\pi_{k'} \geq \pi_0/2$. Denote by \tilde{k} the reordering of the groups under \tilde{p}_k with \tilde{k}' corresponding to the new position of category k' . If $\tilde{k}' < \tilde{k}^*$ the equilibrium remains the same. If $\tilde{k}' \geq \tilde{k}^*$ then $\tilde{\beta}_L^D < \beta_L^D$ and $\tilde{\beta}_R^D > \beta_R^D$ by Corollary 1 and Lemma 4. This completes the proof. \square

Chapter 2

Using Information with a Known Bias in the Lab

2.1 Introduction

We all receive copious amounts of information from other people in the form of advice. Most often this advice, even in the absence of strategic motives, is biased. A friend who recently crashed his car is probably more likely than average to advise you to buy insurance. And receiving a bad review of an airline from a relative who is afraid of flying might not tell you so much about the quality of the airline. In this paper we ask the following question:

if advisors are biased but bias is publicly known and easily adjusted for, will people give higher weight to similarly biased advisors?

We investigate this by setting up an experiment where subjects estimate probabilities. They take either the role of Receivers or Senders. We will use the pronominal convention that Receivers are denoted by *He* and Senders by *She*. Receivers observe the advice of a Sender who has private information. The Sender is only rewarded according to the accuracy of her own estimate and hence she has no incentive to lie. Sender and Receiver may have different biases but this is publicly known and easily adjusted for.

After receiving the Sender’s advice the Receiver makes his own estimate. The following example illustrates what we mean by easily adjustable bias.

Example. A manager has two project evaluators. Suppose that, from the manager’s point of view, one evaluator is very optimistic (say that he always overestimates the probability of success by 20% compared to the manager) and the other is slightly optimistic (say he overestimates by 10% compared to the manager) but otherwise they are equally good. The manager’s job should then be easy: he just divides the estimate of the former by 1.2 and the estimate of the latter by 1.1. There is no objective reason for him to give more weight to the opinion of the more similar evaluator (the slightly optimistic).

In the paper we estimate the weight that Receivers give to Senders’ advice and refer to this as *trust*. We then analyze how trust depends on the difference in the bias of the Sender and the Receiver and show that the answer to the question posed above is affirmative: Receivers trust similarly biased Senders more. Furthermore, we provide evidence on some of the contextual factors that affect trust: the experiment allows us to investigate how trust varies when the Sender makes “extreme estimates” and when we allow Receivers to choose the Sender, knowing only her bias.

Preview of experiment. The experiment is illustrated in Figure 2.1. Subjects estimate the probability of drawing a black ball from a cage containing only black and white balls.¹ In particular, 6 balls are known and 4 balls are unknown and either all black or all white (i.e. $X_1 = X_2 = X_3 = X_4$) with known probability. Subjects play two roles in the experiment. As *Senders* they observe a sample drawn with replacement from the 10 balls in the cage \mathbf{Y}_S before estimating the probability of drawing a black ball from \mathbf{Y}_S . We call these estimates *SenderEst*. Senders are only rewarded based on their own estimate and therefore have no incentives to lie. As *Receivers* they have no sample but instead learn the estimate of a sender (*SenderEst*). This estimate is transmitted

¹They are rewarded according to a scoring rule which depends only on the estimate and the color of a randomly drawn ball.

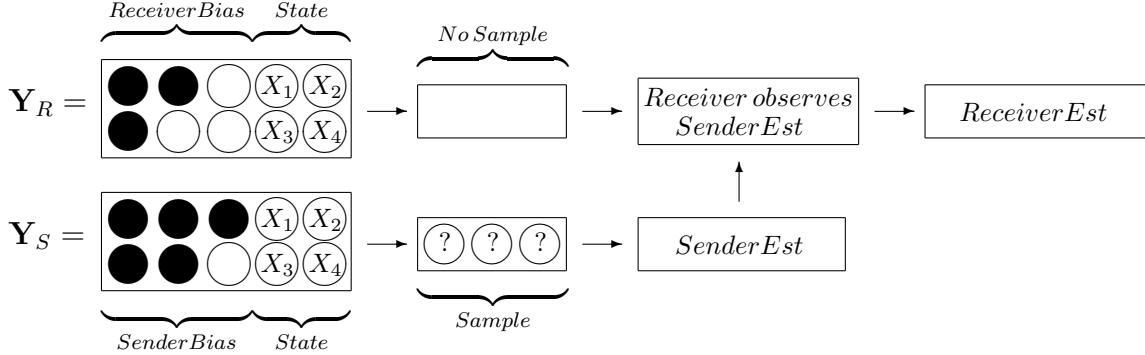


Figure 2.1: Experimental Setup

mechanically. The Receiver then estimates the probability of drawing a black ball from Y_R . We denote this estimate by *ReceiverEst*. Importantly,

- i. Sender and Receiver are in the same *State* (i.e. $\{X_1, \dots, X_4\}$ are the same for both) such that the Sender's estimate is relevant to the Receiver, and
- ii. Sender and Receiver may have different *Bias*. I.e. the 6 balls of known color may be different (in Figure 2.1, *ReceiverBias* = 3 and *SenderBias* = 5).

All this is public information such that if the Receiver believes the Sender to have made a correct estimate his own best estimate is

$$SenderEst - (SenderBias - ReceiverBias)/10.$$

Research questions. We are interested in measuring the effect of the distance between the two biases.

$$BiasDist = |SenderBias - ReceiverBias|.$$

We then ask the following questions: First, *does BiasDist matter?* In particular, will Receivers realize that they can easily correct for the Sender's bias or will they be less likely to use her estimate the greater is *BiasDist*. Second, *what if the Sender makes an unreasonable estimate?* In Figure 2.1 the Sender has at least 5 and at most 9 black balls

out of the 10 balls in \mathbf{Y}_S . Hence, if *SenderEst* is smaller than 0.5 or greater than 0.9, the estimate is logically impossible given any type of generalized Bayesian updating. We analyze how Receivers respond to such estimates. Lastly, we introduce a new stage (stage 3) in the experiment where Receivers can choose between two Senders, knowing only *SenderBias*. *Will Receivers place more or less weight on Sender estimates when they can choose the Sender?*

Results. The analysis is carried out at two levels: aggregate and disaggregated (subject) level. The results indicate that:

- i. *Bias matters.* The evidence that the weight given by Receivers to Senders is decreasing in *BiasDist* is strong at the aggregate level (Section 2.3) although less so at the disaggregated level (Section 2.5.1).
- ii. *Extreme estimates are generally not trusted.* Receivers trust extreme estimates much less but still give them positive weight. In 60 – 70% of the cases an extreme estimate contains qualitative information about the sample that has been observed by the Sender so Receivers seem to be right in not discarding these estimates completely. (Section 2.5.2).
- iii. *Choosing the Sender.* Receivers trust Senders more in stage 3 (where they have chosen the Sender knowing only his bias) than in stage 2 (where the Sender is exogenously given). We show that this effect is only present when the Receiver chooses the most similar Sender. (Section 2.5.3).

The experiment is motivated by theoretical work in Rudiger [39], where a Sender-Receiver game with non-common priors is analyzed. Different biases induce different ex post preferences but under common knowledge and truthful information transmission bias should not matter. We intend to test this hypothesis.

Applications. The setup is relevant to a number of situations apart from the above example, for instance media and news consumers, stock market reports and investors, consultants and firms, job market references and employers. More generally,

the experiment will be relevant to any type of relationship where an agent with private information but a different bias transmits information to a principal. In the discussion in Section 3.6 we will return to some of these applications and compare them to our results.

Probability updating in experiments. The literature on bias and heuristics in probability updating is extensive. We will not review all of these articles here but only mention some of the most related findings. *When estimating objective probabilities many subjects suffer from baserate neglect*²: Kahneman and Tversky [46] identify several heuristics including a type of baserate neglect; Griffin and Tversky [26] analyze the difference between weight and strength of evidence and use it to explain over- and underconfidence in probability estimates; Grether and El-Gamal [25, 18] investigate generalized Bayesian updating and find that many subjects use updating rules consistent with baserate neglect. *When predicting the behavior of others, subjects tend to be “conservative”*: Huck and Weizsacker [28] test whether subjects are able to correctly predict the actions of others and find that predictions are good on average but are systematically conservative, in the sense of beliefs being distorted toward a uniform prior. *Context affects updating*: Charness, Karni and Levin [10, 9] test updating in a setting which resembles a two-armed bandit problem and find that subjects are much less likely to make updating mistakes when: (i) the Bayesian choice is reinforced by a good payoff history, (ii) payments do not depend on outcomes and (iii) references to success and failure are eliminated. In the discussion of Section 3.6 we return to comment on how our results compare to the literature.

Information transmission in experiments. The literature on strategic information transmission experiments has been surveyed by Crawford [12] but our setting will be non-strategic and therefore the most related area is the literature on advice-giving. *Evidence shows that advice is a strong force in the creation of norms and conventions in intergenerational games*: “word-of-mouth”-learning is often far more influential in

²Baserate neglect refers to ignoring “base” or prior probabilities. I.e., placing too much weight on the evidence of data in updating.

creating social norms than learning from history (Schotter and Sopher [42]); in a trust game, allowing subjects to pass on advice tends to decrease trust (amount sent) but increase trustworthiness (amount returned) (Schotter and Sopher [40]); advice induces the creation of conventions in the ultimatum game (Schotter and Sopher [41]). *Advice tends to improve learning and decision making:* Çelen, Kariv and Schotter [7] find that receiving advice augments the tendency to herding in a sequential choice experiment but also that payoffs are higher with advice and approach those achieved under perfect information; Progrebna [37] presents an interesting empirical side-kick to the experimental evidence. Contestants in a particular Italian game show generally make better decisions when they follow the audience’s advice but do not choose to do so very often. *Higher weight given to advice of similar advisors:* Nyarko, Schotter and Sopher [35] set up an experiment with a market for data and advice, where the background (gender and major) of the advisor is known. They find a tendency for subjects to bid more for the information of advisors with the same background as themselves, even when this information consists of unprocessed data, strongly suggesting the presence of some sort of homophily.

Our experiment is closest to Nyarko, Schotter and Sopher’s experiment on advice giving but is different in that we induce a bias which is completely exogenous and then test how subjects deal with it. To our knowledge this has not been done before. The setup is closely related to many of the above experiments (e.g. outcomes determined by draws of balls from a “bingo cage”) but we ask subjects to estimate the probability of an outcome (drawing a black ball) rather than soliciting a state probability, as for instance in Grether’s experiment.

The paper proceeds as follows. Section 2.2 describes the setup of the experiment. In Section 2.3 we carry out an initial analysis of the aggregate data to gain intuition. Section 2.4 describes in detail the econometrics used to disentangle risk attitudes and probability estimates. Section 2.5 presents the results at subject level. In Section 3.6 we discuss the results and their relation to the findings of the existing literature.

2.2 An Experiment on Probability Estimation

2.2.1 Setup

The objective of subjects is to estimate the probability of drawing a black ball from a cage which contains 10 balls: 6 balls (referred to as **Bias**) are of known color and 4 balls are unknown (referred to as **State**). With probability π all the 4 balls in **State** are black and with probability $1 - \pi$ they are all white. The state probabilities alternate between $\pi = 1/2$ and $\pi = 2/3$.

Stage 1. In stage 1 the subject observes a **Sample** of 3 balls drawn with replacement from the 10 balls in **Bias** \cup **State**. Another ball y is then drawn from the cage with all balls replaced. The subject is asked to estimate the probability that y is black. Let (b, w) denote a cage with b black balls and w white balls and let $S \equiv \#\{\text{black balls in Sample}\}$ and $Bias \equiv \#\{\text{black balls in Bias}\}$. Table 2.1 summarizes stage 1.

Table 2.1: Setup Stage 1

1. Bias known State unknown		2. Mixed Cage		3. Draw Sample and y with replacement
Bias = $(Bias, 6 - Bias)$	\rightarrow	Bias \cup State	\nearrow	Sample = $(S, 3 - S)$
State = $(4, 0)$ or $(0, 4)$			\searrow	y

Stage 2. In stage 2 subjects (Receivers) have *no* sample but observe the stage 1 estimate (*SenderEst*) of another subject (the Sender) who has faced the same state but possibly had a different bias. Refer to the example in Figure 2.1. As in the introduction, we refer to the *Bias* of the Sender as *SenderBias* and that of the Receiver as *ReceiverBias*.

Stage 3. In stage 3 the only difference to stage 2 is that the subject can choose between two Senders, SenderA and SenderB, with different *SenderBias*. When choosing he knows only *SenderABias* and *SenderBBias* but not the estimates of the two Senders.

Table 2.2 recaps the stages. Screen prints are included in the appendix, Section 2.7.³

Table 2.2: Stages

	Stage 1	Stage 2	Stage 3
Sample	Yes	No	No
Observe Estimate from Sender	No	Yes	Yes
Choose Sender	No	No	Yes
Rounds	20	30	10

2.2.2 Scoring Rule

Subjects were paid for 5 randomly chosen rounds (out of the total of 60) in order to limit wealth effects. Payoffs were calculated according to a quadratic scoring rule. In particular, if the subject indicated a probability of q that y is black, he received the following remuneration (expressed in experimental currency, ECU, which was later exchanged at ECU 5,000 = 1 EUR).

$$w(q) = \begin{cases} 100 * (1 - (1 - q)^2) & \text{if } y \text{ is black} \\ 100 * (1 - q^2) & \text{if } y \text{ is white.} \end{cases}$$

Under risk neutrality the quadratic scoring rule is incentive compatible, which is why it was chosen. However, it is often seen that subjects are not risk neutral even over small stakes and therefore we estimate risk attitudes as part of our analysis. Before finalizing an estimate subjects could easily check their potential scores by one click of a button.

³Notice that in the experiment, the colors were red and black and subjects were asked to estimate the probability of drawing a red ball.

2.2.3 Sessions

To collect data 4 sessions of the experiment were carried out with a total of 66 subjects (excluding a number of unpaid pilots with graduate students and a small paid pilot with undergraduate students). All sessions took place in the experimental lab of Universidad Carlos III de Madrid. The subjects were drawn at random from a pool of roughly 1,000 students who had signed up to be contacted. The usual precautions were taken so as to not inform subjects about the contents of the experiment on forehand. A participation fee of EUR 4 was paid on top of any winnings from the experiment. Subjects were mainly undergraduate students from the economics, business and engineering faculties. The gender distribution was 59% males and 41% females.

Table 2.3: Sessions

Session	Date	Input Data ^a	Subjects
S0	31/11/2010	Pilots ^b	18
S1	01/12/2010 (a.m.)	S0	18
S2	01/12/2010 (p.m.)	S0	11
S3	02/12/2010	S0	19

^a The input data refers to the *SenderEst* given to subjects in stage 2 of the experiment. These are estimates from other subjects who have done the experiment at another time.

^b The data for this session was taken mainly from a paid pilot experiment and supplemented with data from an unpaid pilot with graduate students.

Subjects were given a short description of the experiment upon arrival and then lead into to the lab and allowed to start. The experiment was completely computerized and programmed using Urs Fischbacher's ZTree. When finished they were asked to wait while payments were prepared and then paid individually before leaving. Students generally used between 45 minutes and 1 1/2 hours to complete the task with a few being faster or slower but were not allowed to leave before at least 45 minutes had passed. Average

winnings were EUR 11.96 with a standard deviation of 0.8.⁴

2.2.4 Feedback and Matching

There is no feedback on the true state of the world at any stage of the experiment. In stages 2 and 3 Receivers are informed that they are matched with a new Sender each round to rule out learning effects.

2.2.5 Benchmarks

To fix a target against which to evaluate subjects we use three measures. Recall that with probability π then **State** is 4 black balls and otherwise 4 white balls. First, it seems probable that some Receivers will ignore the Sender's advice and focus on the prior probability of drawing a black ball. This is given by

$$Prior = \frac{Bias}{10} + \pi \cdot \frac{4}{10}. \quad (2.1)$$

On the other hand, if the Receiver believes the Sender to have made a correct estimate the maths of the second stage are simple: if the Sender's estimated probability of drawing a black ball is given by *SenderEst*, the Receiver's corresponding estimate should be

$$AdjEst = SenderEst + \frac{ReceiverBias - SenderBias}{10}. \quad (2.2)$$

It is also possible that the Receiver makes a mistake by failing to realize that he must correct for the bias difference, and just uses *SenderEst* as his estimate. This gives us three targets to evaluate subjects against:

- i. The trusting type who uses *AdjEst*.
- ii. The trusting type who fails to adjust correctly and uses *SenderEst*.
- iii. The untrusting type, who uses *Prior* as his estimate.

⁴The standard deviation is quite low reflecting the low power of the quadratic scoring rule. However, most subjects did not seem to notice this given that they spent a great deal of time in completing the task.

Furthermore, we will be interested in testing whether the weight given by Receivers to Senders depends on the distance in their bias. Recall that

$$BiasDist = |ReceiverBias - SenderBias|.$$

We also want to test whether the weight changes when the Sender's estimate is clearly not Bayesian. E.g., in the example in Figure 2.1 the Sender has at least 5 and at most 9 black balls in his cage. Estimates below 0.5 or above 0.9 cannot be said to follow any type of generalized Bayesian updating. We refer to this as an extreme estimate.

$$ExtremeEst = 1 - \mathbb{I}(SenderBias/10 < SenderEst < (SenderBias + 4)/10).$$

2.2.6 Hypotheses

The following main hypotheses are tested in the experiment. Since the Sender's informativeness does not change with her bias, theory would predict that *BiasDist* should not affect the weight the Receiver gives to the Sender. Neither should choosing the Sender (Stage 3). On the other hand, we hypothesize that observing an extreme estimate leads the Receiver to give lower weight to the Sender.

- H1. The weight given to Sender estimates does not depend on *BiasDist*. See Section 2.5.1 for evidence on this.
- H2. Extreme estimates (measured by $ExtremeEst = 1$) are given lower weight. See Section 2.5.2 for evidence on this.
- H3. Choosing the Sender in stage 3 does not change the weight placed on his estimate. See Section 2.5.3 for evidence on this.

2.3 Preliminary Analysis: Aggregated Data

Before moving on to a more detailed statistical analysis we first look at a simple linear regression using the targets described in Section 2.2.5. While not taking account for risk

Table 2.4: Regression Results

	Dep.Var.: <i>ReceiverEst</i> – <i>Prior</i>	
	Stage 2	Stage 3
<i>AdjEst</i> – <i>Prior</i>	0.385 (5.30) ^{***}	0.531 (5.34) ^{***}
$(AdjEst - Prior) \times BiasDist$	-0.271 (-3.41) ^{**}	-0.648 (-3.06) ^{**}
$(AdjEst - Prior) \times ExtremeEst$	-0.170 (-4.06) ^{***}	-0.178 (-2.10) [*]
<i>SenderEst</i> – <i>AdjEst</i>	0.119 (3.34) ^{**}	0.108 (2.87) ^{**}
<i>Constant</i>	0.012 (2.26) [*]	0.006 (0.79)
R^2	0.167	0.192
Observations	1980	480

t-statistics in parentheses. Standard errors clustered around subjects.

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

attitudes and the fact that variables are limited to the unit interval this analysis gives a good idea of what to expect. The results are presented in Table 2.4.

In each column of Table 2.4 the dependent variable is the deviation of the Receiver's estimate from the prior probability ($OwnEst - Prior$). The coefficient on $AdjEst - Prior$ is a measure of how much the Receiver trusts the Sender and we allow this to interact with $BiasDist$ and $ExtremeEst$. The term $SenderEst - AdjEst$ captures whether Receivers fail to adjust for the difference in bias. This coefficient should evaluate to 0 if Receivers do the maths perfectly and 1 if they completely fail to take the relative bias into account.

How do the hypotheses fare? Table 2.4 provides us with the first evidence. *H1: Weight placed on Sender estimate is decreasing in BiasDist.* The results for stage 2 and 3 indicate that Receivers place less weight on the Sender's estimate the greater the distance in bias (the coefficient on the term $(AdjEst - Prior) \times BiasDist$). *H2: Extreme estimates receive lower weight.* In both stages less weight is placed on estimates that are extreme (the coefficient on the term $(AdjEst - Prior) \times ExtremeEst$). *H3: Difference between stages not significant.* Receivers seem to place higher weight on Sender's estimates in stage 3 than in stage 2 (compare the coefficients on $AdjEst - Prior$). But when estimating the model for stages 2 and 3 jointly and including a set of interaction terms for stage 2 we cannot reject the null-hypothesis that the coefficient on $AdjEst - Prior$ is the same in stages 2 and 3 (p-value= 0.087). Hence, the evidence from the regressions rejects H1 but not H2 and H3. We return to the hypotheses in Section 2.5.

Do Receivers realize that the Sender's estimate is valuable? If subjects did not understand the purpose of the exercise the coefficient on $AdjEst - Prior$ should not be significant. It turns out to be significant and have the expected sign. However, $SenderEst - AdjEst$ is also significant (although with a small coefficient) and therefore there are indications that Receivers do not always do the adjustment well.

Learning. Two dummy variables, for periods 11-20 and 21-30 respectively, were interacted with $AdjEst - Prior$ in stage 2. None of these were significant suggesting that there were little or no learning effects.

Do Receivers give lower weight to relatively biased Senders out of computational concerns? One could argue that perhaps Receivers place high weight on Senders with no relative bias out of computational concerns: the formula for bias correction is always the same but of course no computation is needed when there is no relative bias. To test whether subjects respond to the actual size of the relative bias or merely the presence of relative bias, we augment the stage 2 regression with an interaction term between the weight placed on the Sender’s estimate and the dummy $IsBiased = \mathbb{I}(|SenderBias - ReceiverBias| \neq 0)$. The results are presented in Table 2.8. The new interaction term does not enter significantly whereas the interaction between the weight placed on the Sender and $BiasDist$ is still significant. This indicates that what matters is the size of the relative bias and hence Receivers do not seem to be motivated by computational concerns.

The above analysis falls short on two points. First, it uses standard regression techniques without imposing any restrictions on parameters or adjusting for the fact that the dependent variable is limited to $[0,1]$. Second, it does not account for heterogeneity between subjects. The methodology proposed in the next section attempts to deal with these shortcomings.

2.4 Methodology for Disaggregated Analysis

The family of updating rules that we allow for in this section is quite large and includes as special cases pure Bayesianism as well as conservatism (overweighting prior) and baserate neglect (overweighting data) but is not based on a generalized Bayesian model of the type considered by for instance Grether and El-Gamal. While such a model could have been used it does not have the same normative properties as the pure Bayesian model.⁵ Therefore we have preferred an alternative which seems more intuitive although we see no reason it should be better or worse than the generalized Bayesian approach.

⁵The classic objection against the generalized Bayesian approach is: if subjects are not Bayesian why should they err in a Bayesian manner?

2.4.1 Risk Attitudes and Probability Estimates

We concurrently estimate the risk attitude and the probability updating model of Receivers.

Risk Attitude. Even though stakes are fairly small in each period we want to allow for the possibility that subjects are not risk neutral (and as it turns out, they are not). To estimate risk attitudes we use the following specification.⁶

$$RiskAdjEst = \frac{(RawEst)^\alpha}{(RawEst)^\alpha + (1 - RawEst)^\alpha}, \quad (2.3)$$

where *RawEst* is the uncorrected probability estimate of the subject and *RiskAdjEst* is the risk adjusted probability estimate given α . The subject is risk-averse for $\alpha > 1$, risk-neutral for $\alpha = 1$ and risk-seeking for $\alpha < 1$.

Probability updating model. For the Receiver's estimate we consider a family of updating rules, which are convex combinations of the adjusted Sender estimate (*AdjEst*), the prior probability (*Prior*) and the unadjusted Sender estimate (*SenderEst*). The Receiver's predicted estimate is then

$$\begin{aligned} PredEst = & \delta \cdot Prior + (1 - \delta) \cdot AdjSenderEst \\ & + (1 - \delta) \cdot \tau \cdot (SenderEst - AdjSenderEst). \end{aligned} \quad (2.4)$$

The weight given by the Receiver to the estimate of the Sender is parametrized by $1 - \delta$. This measures the trust the Receiver places in the estimate of the Sender. If the Receiver trusts the Sender to have made a correct estimate he should set $\delta = 0$. On the other hand, if he does not trust the Sender at all he should stick to his prior and set $\delta = 1$. If the Receiver correctly adjusts for the bias difference between the two then $\tau = 0$. We

⁶Recall that $w(\cdot)$ is the scoring rule. Our risk-adjustment can be seen as following from the solution to the subject's maximization problem with respect to *RawEst* when he has a power utility function of the type

$$u(w(RawEst)) = -\frac{1}{2(\alpha + 1)} [-(w(RawEst) - 1)]^{\alpha+1}.$$

The specification was chosen because it works well with risk-seeking behavior.

allow him to make mistakes such that $\tau > 0$ if he underadjusts and uses the Sender's "biased" estimate rather than the adjusted estimate.

Error structure. We suppose that subjects make random errors when calculating these estimates and that the errors follow a truncated normal distribution such that the likelihood of the Receiver's risk adjusted estimate being *RiskAdjEst* when the probability updating model predicts *PredEst* is given by a truncated normal distribution.

$$l(PredEst, RiskAdjEst) = \frac{\frac{1}{\sigma} \phi\left(\frac{PredEst - RiskAdjEst}{\sigma}\right)}{\Phi\left(\frac{1 - RiskAdjEst}{\sigma}\right) - \Phi\left(\frac{-RiskAdjEst}{\sigma}\right)}, \quad (2.5)$$

where ϕ is the standard normal density function with distribution function Φ and standard deviation σ . We then substitute (2.4) into (2.5) and estimate the model by maximum likelihood.

Testing hypotheses. To test our hypotheses we make δ flexible and let it depend on two variables which were introduced in Section 2.2.5: the relative bias between Receiver and Sender, *BiasDist*, and the dummy variable indicating extreme estimates, *ExtremeEst*. We make the restriction $\delta \in [0, 1]$ by using a link function

$$\delta = W(\delta_0 + \delta_1 BiasDist + \delta_2 ExtremeEst), \quad (2.6)$$

where $W(\cdot) = \exp(\cdot)/(1 + \exp(\cdot))$. To make the results easier to analyze we use the following (marginal) measures

$$\begin{aligned} \delta_{cons} &= w(\delta_0) & : & \text{baseline trust} \\ \delta_{dist} &= w(\delta_0 + \delta_1 \cdot 0.1) - w(\delta_0) & : & \text{marginal } BiasDist \text{ effect} \\ \delta_{ob} &= w(\delta_0 + \delta_2) - w(\delta_0) & : & \text{marginal } ExtremeEst \text{ effect.} \end{aligned}$$


Thus, $\delta = \delta_{cons}$ in (2.4) if Receiver and Sender have the same bias and the Sender's estimate is not extreme. δ_{dist} is the marginal effect of increasing *BiasDist* from 0 to 1 when *ExtremeEst* = 0. And δ_{ob} is the marginal effect of the Sender's estimate being extreme when *BiasDist* = 0.

2.5 Results

2.5.1 H1: Does Relative Bias Matter?

Recall that δ_{dist} measures the marginal effect of *BiasDist* on trust. The null-hypothesis is that $\delta_{dist} = 0$, i.e. that *BiasDist* should not matter to how much the Receiver trusts the Sender and the most intuitive alternative hypothesis is that $\delta_{dist} > 0$, implying that Receivers trust Sender estimates less the greater the relative bias. Figure 2.2 represents the individual estimates of δ_{dist} from stage 2 and 3. The figure indicates a great deal of difference between the two stages: whereas stage 2 estimates are closely distributed around 0, stage 3 estimates are almost all positive.

Figure 2.2: Effect of Bias on Prior Weight in Stages 2 and 3



In stage 2, a t-test of the hypothesis that the mean of the distribution of δ_{dist} is zero is not rejected ($t = -0.61$) but using a more conservative Wilcoxon sign-rank test we can reject the null-hypothesis in favor of the mean being negative ($z = -2.68$). This goes against the intuition we had initially. However, of the 66 subjects as many as 20 have $\delta_{cons} > 0.95$. This implies that they can hardly increase the weight they give to their prior. If we restrict the sample to subjects with $\delta_{cons} < 0.95$ we get a z -value of

-1.95 and for $\delta_{cons} < 0.8$ we have z -value -0.05 .

In stage 3 subjects were faced with the same problem, except that they had chosen the Sender themselves. Yet in stage 3 the sign-rank test indicates that the mean of the individual δ_{dist} is positive ($z = 5.00$) and this effect remains significant when we restrict the test to the subsamples mentioned above.

Result 1. *There is only weak evidence of Receivers responding to relative bias in stage 2 (evidence only found in aggregate analysis, not in individual estimates) whereas in stage 3 there is strong evidence that Receivers trust Senders less the greater the difference in bias (both at aggregate and individual level).*


There are at least a couple of possible explanations for why results are stronger in stage 3: (a) *Learning effect.* In Section 2.3 we tested for within-stage learning effects and found nothing. It also seems unlikely that after 30 periods in stage 2 subjects should suddenly learn in the 10 periods of stage 3. We therefore discard this option. (b) *Salience.* In stage 3 Receivers are asked to choose between two Senders with different bias. This draws attention to the bias and therefore it is possible that relative bias has a larger effect simply because of increased salience. If this is the case, it seems reasonable to use stage 3 results since they tell us that whenever Receivers pay attention to relative bias they are likely to be affected by it.

2.5.2 H2: Extreme Estimates

We hypothesize that Receivers will trust Senders less when these are outside of the rational risk-neutral interval ($ExtremeEst = 1$) and thus represent, if taken at face value, a logical impossibility (for instance estimating a 90% chance of drawing a black ball when there can be at most 7 black balls in the cage). This translates into $\delta_{ob} > 0$ with the alternative hypothesis that $\delta_{ob} \leq 0$. We want to consider the possibility that the inequality is reversed for the following reason: it might well be that Receivers interpret an extreme Sender estimate as indicating that the Sender has observed an extreme sample, i.e. either 0 or 3 black balls. If the Receiver believes this to be true he is less uncertain about what information to extract and in this case it makes sense to trust the Sender more.

Figure 2.3 represents the distribution of the individual estimates of δ_{ob} in stage 2. The distribution seems to be right-skewed and has positive mean. A t-test on the mean of the individual estimates rejects the null-hypothesis of $\delta_{ob} = 0$ ($t = 2.70$ and 67 degrees of freedom) but a Wilcoxon sign-rank test just fails to reject the null at the 5% significance level ($z = 1.94$). In stage 3 the mean is positive but neither of the two tests rejects the null.

Figure 2.3: Extreme Estimates and Prior Weight in Stage 2



Would Receivers be right to trust extreme Senders less? Table 2.5 presents the distribution of sample information conditional on estimates being above or below the rational interval. If extreme estimates provided strong sample information then we should see the distribution being concentrated around the upper-left and the lower right cell, respectively. In the first row there is a slight tendency toward this - an estimate below the lower bound is more likely when the sample has no red balls but there is a lot of noise still. In the second row there is no such tendency. To support this evidence we check how many times Senders at least go in the right direction when making extreme estimates, i.e. when he is above the rational interval then $Post - Prior > 0$ and vice versa. This happens in 69% of the cases when the estimates are above the rational interval and 62% of the cases when it is below.

We conclude that extreme estimates are quite noisy although more than half the time they indicate the correct direction of the updating. Taking this evidence together it seems reasonable that Receivers trust such estimates less without discounting them completely.

Table 2.5: Sample Distribution in Stage 1 Conditional on “Estimation Error”

# red balls in sample	0	1	2	3
$SenderEst < LowerBound^a$.33	.24	.24	.20
$SenderEst > UpperBound^b$.29	.20	.22	.29

^a $LowerBound = B_S/10$.

^b $UpperBound = (B_S + 4)/10$.

Result 2. *There are strong indications at the aggregate level but less strong indications at the subject level that Receivers trust extreme Sender estimates less. They would be partially right in doing so, in that extreme estimates contain a great deal of noise rather than being signals of “extreme” samples.*

2.5.3 H3: Choosing the Sender in Stage 3


Allowing Receivers to choose the Sender has an effect

We saw earlier that Receivers tend to behave differently in stage 3, although the nature of the task is identical to stage 2 except that Receivers can choose their Sender in stage 3. There does not seem to be an a priori reason to act differently in the two stages, yet this is exactly what we observe. Figure 2.4 documents this. Generally, Receivers trust Senders more in stage 3 (δ_{cons} is lower). A sign-rank test confirms this (z-value 5.07).

Result 3. *Receivers place more weight on Sender estimates when they can choose the Sender (stage 3) than when they cannot (stage 2).*

There seems to be a psychological effect of having chosen the Sender. Even though this choice could in no way affect the Sender’s estimate and Receivers would know

Figure 2.4: Prior Weight in Stages 2 and 3



this (since they know that Senders were chosen from another session) they trust them more. It could be a question of framing: drawing attention to the Sender makes her information more salient. But the choice could also give Receivers an illusion of control, making them feel erroneously that their choice affects the informativeness of the Sender. The additional evidence below shows that the effect is almost exclusively limited to those Receivers that choose the Sender with the smallest relative bias. This favors the hypothesis of illusion of control over salience, since salience is the same for all subjects whereas illusion of control can reasonably be thought to be correlated with choosing the Sender with the smallest relative bias.

Additional evidence: Is stage 3 behavior related to whether the Receiver chooses the most similar Sender?

In stage 3 Receivers choose the Sender with the smallest *BiasDist* 49% of the time. This almost seem random so we want to investigate if Receivers act the same or differently in the two cases. It seems logical that the Receivers who for one reason or another care about relative bias are the ones that will choose the Sender with the smallest relative bias in stage 3. We test this by creating two dummy variables: $S3 = \mathbb{I}(\text{Stage} = 3)$

and $MinBiasS3 = MinBias \times S3$ where $MinBias$ is equal to 1 if in stage 3 if the Receiver chooses the Sender with the smallest $BiasDist$ and 0 otherwise. We then take the model in Table 2.4 and run stage 2 and 3 jointly with all 4 explanatory variables interacted with the two dummies.

The results show the following. (i) The interaction between $MinBiasS3$ and $(AdjEst - Prior)$ is positive and significant ($p = 0.007$). This indicates that Receivers who chose the most similar Sender trusted more in Stage 3 than in Stage 2. (ii) The interaction between $(AdjEst - Prior) \times BiasDist$ and $MinBiasS3$ is significantly negative ($p < 0.001$). Thus, Receivers who chose the most similar Sender punished $BiasDist$ more in Stage 3 than in Stage 2. (iii) The interaction between $MinBiasS3$ and $(AdjEst - Prior) \times ExtremeEst$ is negative and significant ($p = 0.020$). Hence, Receivers who chose the most similar Sender punished $ExtremeEst$ more in Stage 3 than in Stage 2. This translates into the following results.

Result 4. *Receivers who choose the most similar Sender are on average*

- i. more trusting,*
- ii. more sensitive to $BiasDist$ (i.e. relative bias is punished more), and*
- iii. more sensitive to $ExtremeEst$ (i.e. extreme estimates are punished more).*

2.5.4 Comments

Analysis of Stage 1 Behavior

First, we define some benchmarks for evaluating stage 1. Suppose 3 balls are drawn with replacement from a population of 10 balls of which b are black. Denote by $B(x, b)$ the probability that out of these 3 balls x are black. Recall that **State** can be either all black balls or all white balls. Given the subject's $Bias$ and a sample with S black balls the probability that the state is all black is

$$P_{Black} = \frac{\pi \cdot B(S, Bias + 4)}{\pi \cdot B(S, Bias + 4) + (1 - \pi) \cdot B(S, Bias)}.$$

Hence, the Bayes posterior probability of drawing a black ball is

$$Post = P_{Black} \times \frac{Bias + 4}{10} + (1 - P_{Black}) \times \frac{Bias}{10}.$$

We want to allow for baserate neglect as well, which corresponds to using the likelihood $L_{Black} = B(S, Bias + 4)/(B(S, Bias + 4) + B(S, Bias))$ instead of P_{Black} , such that

$$Like = L_{Black} \times \frac{Bias + 4}{10} + (1 - L_{Black}) \times \frac{Bias}{10}.$$

Table 2.6 shows the result of a regression of *OwnEst*–*Prior* on these benchmarks, where *Prior* is as defined previously. If subjects use the Bayesian posterior the coefficient on *Post* – *Prior* should be 1. If they use the Bayesian prior the coefficient should be 0. It turns out to be 0.574 and significantly different to zero. Hence, subjects do use the information contained in the sample but are, on aggregate, not fully Bayesian updaters. We have included the term *Like* – *Post* to allow for baserate neglect. Results indicate that baserate neglect is not present in the aggregate data.

Table 2.6: Stage 1 Regression Results

	Dep.Var.: <i>OwnEst</i> – <i>Prior</i>
<i>Post</i> – <i>Prior</i>	0.574 (3.31)**
<i>Like</i> – <i>Post</i>	-0.134 (-0.90)
<i>Constant</i>	-0.007 (-0.97)
R^2	0.091
Observations	1320

t-statistics in parentheses. Standard errors clustered around subjects.

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Remark 1. *On aggregate subjects use some of the sample information but do not do full Bayesian updating. Furthermore, there is no evidence of baserate neglect.*

Linking Stage 1 and 2 Behavior

It seems plausible that there is a connection between stage 1 and stage 2 behavior but the nature of this link is less obvious. Subjects that are more Bayesian in stage 1 might expect others to behave similarly. Or they may believe that they are smarter than others and not pay attention to them.

To assess this we first need to estimate an individual level model for stage 1. Recall the benchmarks *Prior*, *Post* and *Like* from Section 2.5.4. As in the analysis of stages 2 and 3 we consider a family of updating rules which are essentially a *convex combination* of the three benchmarks. The subject's predicted estimate is then

$$\begin{aligned} PredEst = & \gamma_P \cdot Prior + (1 - \gamma_P) \cdot Post \\ & + (1 - \gamma_P) \cdot \gamma_L \cdot (Like - Post). \end{aligned} \tag{2.7}$$

We make the restriction $\gamma_P, \gamma_L \in [0, 1]$. Hence, γ_P is the weight placed on the prior versus the posterior probability and γ_L is included to allow for baserate neglect. Thus, if the subject is completely “Bayesian” then $\gamma_P = \gamma_L = 0$. If he uses only the prior and ignores the sample then $\gamma_P = 1$ and $\gamma_L = 0$. If he is Bayesian but suffers from baserate neglect then $\gamma_P = 0$ and $\gamma_L > 0$. We substitute (2.7) into (2.5) and estimate the model for each subject by maximum likelihood. We then regress individual estimates of γ_P from stage 1 on δ_{cons} , δ_{dist} and δ_{ob} from stage 2. The results are presented in Table 2.7.

There is a positive and significant correlation between γ_P and δ_{cons} . I.e., the more weight a subject places on his prior in stage 1 the more he does so in stage 2. Turning this around, it implies that the more “Bayesian” a subject is in the first stage, the more he trusts the Sender in the second stage. Looking closer at the data this effect seems to stem from a group of very conservative subjects who always follow their prior and disregard any other information. There are also positive correlations between γ_P and respectively δ_{dist} and δ_{ob} , although these are less significant. This seems to imply that subjects who are more conservative in stage 1 punish more severely Senders who have different biases or have made extreme estimates in stage 2 and give their estimates lower weight.

Remark 2. *There is a positive correlation between the weight that subjects place on*

Table 2.7: Regression of Stage 1 Prior Weight on Stage 2 Parameters


OLS Reg. Results	γ_P
δ_{cons}	0.307 (3.31)**
δ_{dist}	0.690 (2.23)*
δ_{ob}	0.256 (2.20)*
Constant	0.598 (7.68)***
R^2	0.193
Observations	66
<i>t</i> statistics in parentheses	
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$	

their prior in stages 1 and 2. Expressed differently, the more “Bayesian” subjects are in stage 1, the more they trust Senders in stage 2.

Risk-attitude

Turning shortly to risk attitudes, Figure 2.5 plots the distribution of the individual estimates of the risk parameter α in stage 1 and 2. The distribution is quite diffuse, with a peak in both stages just below 0.5. The vast majority of subjects have $\alpha < 1$ which implies risk-seeking behavior. It seems reasonable that this is due to the small stakes employed in our experiment. A sign-rank test does not reject the null-hypothesis of the difference between stage 1 and stage 2 values of α being zero (although only barely with z-value 1.87), so there does not seem to be any material differences in risk attitudes in the two stages.

Figure 2.5: α -estimates in Stages 1 and 2



Remark 3. *Subjects are generally risk-seeking and seem to have the same risk attitude in stage 1 and 2.*

2.6 Discussion

2.6.1 Main Findings Corroborated by Literature

Homophily and trust. Whereas Nyarko, Schotter and Sopher [35] find a tendency for homophily - being more trusting of advisors that have similar personal characteristics - we want to construct a setting where there is no motive for trusting one Sender more than another. Although the evidence is somewhat mixed there are indications that trust is not independent of bias, even in our experiment where bias is exogenous and easily correctable. This is reminiscent of other laboratory regularities such as money illusion and framing effects. However, we also find evidence that many subjects generally trust advice very little or ignore it altogether. This corroborates Nyarko, Schotter and Sopher's finding that subjects generally prefer data over advice and Progrebna [37] who shows that game show contestants often do not follow the audience's advice even though this is in general quite informative. It is also in line with Huck and Weizsacker [28] where

predictions on other subjects' behavior are found to be distorted toward a conservative uniform prior.

Updating with affect. It is striking that when we introduce a small element of choice in the experiment in stage 3 we observe that: (i) trust is generally higher, (ii) relative bias is punished more, and (iii) extreme estimates are trusted less. We also show that this effect seems be driven by the cases where the Receiver actually chooses the most similar Sender. Subjects seem to suffer from “illusion of control”: although their choice is not related to the informational content of the advice they still become more trusting and at the same time more sensitive to relative bias and extreme estimates. This corresponds well to Charness, Karni and Levin [10, 9] who find that subjects make more updating mistakes in the presence of affect (which they test by including/excluding references to outcomes as successes).

Baserate neglect. A stylized fact of probability updating experiments is that a large group of subjects suffer from baserate neglect (too much weight on data) and a smaller group from conservatism (too much weight on prior). We find conservatism to be more predominant in our experiment, which illustrates that the results are very much subject to the particular setup.

2.6.2 How Do the Results Relate to Empirics?

Media bias. Our finding that similar Senders are trusted more can be related to the effect of news on political opinions. Take the example of the US stimulus package of February 2009, which included \$288.000 millions in tax cuts (37% of the total package). This is a hard fact that leaves little, if any, room for interpretation in the media. Yet according to data from the Pew Research Center [8] more frequent viewing of Fox News is correlated with *less* knowledge about this fact whereas more frequent viewing/reading of the other media is correlated with better knowledge. A study by DellaVigna and Kaplan [16] goes further than this and establishes evidence of causality. They use the gradual introduction of Fox News across US states for identification and show that the introduction of Fox News was associated with a right-swing in voting. Since there is no

scarcity of news media in the US, this effect seems to be tied to the slant on news rather than the provision of new information. If news consumers give higher weight to media with the same bias as themselves, merely introducing a conservative news source could have an effect on opinions, as seems to have been the case with Fox News.

Stock markets. Numerous studies deal with bias in stock market recommendations: Dische [17] documents that stock market analysts update in the right direction but with too small magnitude (conservatism); Capstaff, Paudyal and Rees [6] find indications of a deliberate optimistic bias in stock market forecasts; De Bondt and Forbes [15] show evidence consistent with herding, overoptimism and overreaction in forecasts. And Balboa, Gómez-Sala and López-Espinosa [3] find that there are significant differences in recommendation bias across countries and that *using de-biased recommendations is more profitable*, providing an empirical parallel to our results.

Another well-known phenomenon in stock markets is the tendency to invest in companies that have a presence in the investor’s home market (the so-called “home-bias”). In the most literal sense, our results suggest that investors may trust forecasters in their home market more because they tend to be subject to the same bias. Thus they may prefer to invest in home market stocks thereby creating the home-bias.

2.6.3 Policy Implications

We have argued that the evidence from our experiment shows that trust in advice depends on how that piece of advice was acquired. This is relevant, not just to economic modeling, but also to regulators seeking for instance to limit effects of advertisement or trying to “nudge” people into making sounder decisions by providing them with (truthful) information. Given that regulators and the people they regulate will often have different biases that shape their opinions and that this bias is often transparent it seems integral to analyze how people update based on biased advice. Taking our results at face value, such regulators should indeed be worried about the bias that they are perceived to have and can potentially increase the effect attained by their messages by allowing people to choose between different information sources.

2.7 Tables and Figures

2.7.1 Tables

Table 2.8: Regression Results

	Dep.Var.: $OwnEst - Prior$
	Stage 2
$AdjEst - Prior$	0.484 (5.65)***
$(AdjEst - Prior) \times IsBiased$	-0.111 (-1.76)
$(AdjEst - Prior) \times BiasDist$	-0.246 (-2.94)**
$(AdjEst - Prior) \times ExtremeEst$	-0.171 (-4.08)***
$SenderEst - AdjEst$	0.119 (3.34)**
Constant	0.0120 (2.23)*
R^2	0.167
Observations	1980

t statistics in parentheses. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Chapter 3

Equivalence of Non-Common Priors and Non-Common Preferences in Information Transmission

3.1 Introduction

3.1.1 Motivation

The literature on strategic information transmission largely features players with either non-common preferences and common prior beliefs (CB) or non-common priors and common preferences (CP). In some applications one framework may be preferred over another. A CB model captures situations such as the game played between a newspaper journalist and an editor. The journalist may have an incentive to exaggerate his information to gain a “scoop” whereas the editor is more concerned with the reputation of the newspaper. Both players know this and the bias of the journalist has nothing to do with his beliefs. He simply has an incentive to misrepresent his information, even if this information perfectly reveals the true state of affairs. On the other hand, CP models describe situations where players have the same preferences conditional on the state of the world but different a priori beliefs. For instance, investment in renewable energy. If we knew the true state of the climate we would supposedly all agree on the

correct action but the state is uncertain and since our prior beliefs differ we interpret information differently, which leads us to prefer different actions. In other applications, however, we may not have a preference for one framework or the other. As an example, differences in investment behavior can be modeled as different beliefs as well as different risk attitudes.

On a practical level both types of models often result in similar expected utility functions and therefore a reasonable question is if we can always find models of both types that yield the same results, in some sense. If we can, then the difference between the two frameworks reduces to a question of interpretation and we can transfer the results we know in one type of models to the other. Since CB models are dominant in the literature (see Section 3.1.2) and more results are available for this type of models we investigate the question: *for which classes of CB models can we always find an “equivalent” CP model?* We consider two measures of equivalence between CB and CP models. Models can be equivalent in the sense that they induce the same individual choice functions. We call this *Choice-Equivalence*. They can also be equivalent in the sense that in equilibrium players have the same strategies. We call this *Strategy-Equivalence*. Perhaps surprisingly, neither is sufficient nor necessary for the other. We say that two models are *equivalent* if they satisfy both types of equivalence.

Our results show that in a discrete world a CB model always has an equivalent CP model when the state space is sufficiently large compared to the action space. In effect, the size of the state space determines how much latitude we have for specifying individual behavior in a CP model, since utilities must be the same. If the state space is too small, certain patterns of behavior that are possible in a CB model are not replicable in a CP model. Moving on to continuous action and state spaces we show a set of sufficiency conditions for equivalence and identify a class of CB models that satisfy these. Last, we show that equivalence in our sense does not imply that other properties of the models are necessarily the same. We construct an example in which the Sender is characterized by his “competence”, i.e. the precision of the signal he observes, and show that even when a CB and a CP model are equivalent the Receiver has a stronger preference for more competent Senders in the CP model than in the CB model.

3.1.2 Modeling Player Heterogeneity in the Information Transmission Literature

The information transmission literature traditionally features players with common priors but different preferences (CB models). The strand of the literature which is occupied with cheap talk games is heavily influenced by the CB-model of Crawford and Sobel [13] and their approach has been used in general. For instance, Gal-Or [20, 21] on information sharing in oligopoly, Farrell and Gibbons [19] on cheap talk with two audiences, Stein [44] on central banks and policy announcements, Watson [47] on information transmission with two-sided information, Morgan and Stocken [33] on stock recommendations when there is uncertainty about the recommender's type, Ottaviani and Sørensen [36] on reputational cheap talk, etc. Models of verifiable information are generally also set in a CB framework, for example Milgrom and Roberts [32], Austen-Smith [2], Seidmann and Winter [43], Wolinsky [48]. One notable exception is Che and Kartik [11] who analyze the interaction between non-common beliefs and non-common preferences and show that the two are different when information acquisition is costly to the Sender.

In the cheap talk literature non-common priors have been modeled in early work by Green and Stokey [24] (circulated in 1980 although only published in 2007) who incorporate both different preferences and different priors and show that improvements in the information structure do not in general improve welfare. But although they allow for non-common priors they do not exploit their structure actively. This, on the other hand, is done by Admati and Pfleiderer [1] who analyze overconfident Senders and show that overconfidence may actually eliminate less informative equilibria and improve information transmission. Kawamura [29] extends this framework to include underconfidence and shows that an overconfident Sender reveals more accurate information when the signal is close to his prior expectation whereas an underconfident Sender reveals more accurate information when the signal is further away from his prior expectation. This is somewhat related to Gentzkow and Shapiro [22] who model a news market and find that news media slant their messages toward the priors of consumers to appear to be of higher quality. In this setting the prior is important in that it governs not just the preferences of consumers but also their expectation of what message they will observe. Rudiger [38] uses a model of non-common priors explain why voters tend to choose news

sources that have more extreme views than the voters themselves.

The paper proceeds as follows. Section 3.2 describes the model and formally introduces the measures of equivalence. Section 3.3 analyzes equivalence when the action and state spaces are discrete and derives a sufficiency condition for equivalence which depends on the relative size of the two spaces. Section 3.4 moves on to continuous action and state spaces and identifies a class of CB models that always have an equivalent CP model. Section 3.5 presents an example in which the Sender has a certain level of competence (probability of observing the true state of the world) and shows that even when a CB and a CP model attain equivalence in the sense described above the Receiver has different preferences over Sender-types in the two models. Section 3.6 concludes.

3.2 Setup

3.2.1 Model

There are two players, the Sender (S) and the Receiver (R). The setup is standard: S observes payoff relevant information in the form of the realization ($x \in \mathbf{S}$) of a random variable. He sends a cheap talk (non-verifiable) message $m \in \mathbf{M}$ to R who then takes an action $y \in \mathbf{Y}$, which affects both players. The players have utility functions $u^i(y, \theta)$, $i = R, S$, which depend on the action y and a state variable, $\theta \in \Theta$. We make the standard assumptions that for each θ there is a unique y that maximizes utility and that this utility maximizing action is increasing in θ (the sorting condition). If \mathbf{Y} and Θ are continuous and $u^i(\cdot)$ is twice-differentiable this corresponds to the familiar derivative assumptions: for each θ then $u_1^i(\cdot) = 0$ for only one y and $u_{12}^i(\cdot) > 0$ for all y .

The players' prior beliefs over the state have probability density $p^i(\theta)$ and their common beliefs over the data-generating process (dgp) have conditional density $\lambda(x|\theta)$. We suppose that the dgp satisfies the increasing likelihood ratio property (ILR). In particular, if $x' > x$ and $\theta' > \theta$ then $\lambda(x'|\theta')/\lambda(x|\theta') \geq \lambda(x'|\theta)/\lambda(x|\theta)$ with strict inequality for some $\theta, \theta' \in \Theta$ and $x, x' \in \mathbf{S}$. We require beliefs about the dgp to be common for two reasons. First, although some work considers different beliefs about the dgp (Admati and Pfleiderer [1]; Kawamura [29]) the by far most common setup

considers players that agree about the interpretation of data but disagree about prior state probabilities. Second, if we consider the extreme case where we can choose the dgp freely for each player the problem we consider is trivial. Hence, what we consider in this paper is in reality the subspace of problems in which the dgp is commonly agreed upon. For instance, imagine that $x \in \{Heads, Tails\}$ is the flip of a coin and the state $\theta \in [0, 1]$ indicates the probability of *Heads*. Both players agree, for instance, that if $\theta = 1/2$ the coin is fair, but they might have different beliefs about the probability of this.

Letting $p^i(\theta|x) \equiv p^i(\theta)\lambda(x|\theta)/\int_{\Theta} p^i(t)\lambda(x|t) dt$, the expected utilities are given by

$$U^i(y, x) \equiv \int_{\Theta} u^i(y, t) p^i(t|x) dt.$$

Our assumptions on $u(\cdot)$ together with ILR imply that there exists a unique choice function $y_*^i(x) \equiv \max_y U^i(y, x)$ which will be increasing in x . In particular, if \mathbf{Y} and \mathbf{S} are continuous and $u^i(\cdot)$ twice-derivable the assumptions imply that for all x then $U_1^i(\cdot) = 0$ for only one y and $U_{12}^i(\cdot) > 0$ for all y . A priori, the players believe the signal x to be distributed according to the density function $l^i(x)$, which is derived from their prior beliefs and the dgp. If the variables are continuous this is given by

$$l^i(x) \equiv \frac{\int_{\Theta} \lambda(x|t) p^i(t) dt}{\iint_{\Theta \times \mathbf{S}} \lambda(s|t') p^i(t') dt' ds}.$$

The solution concept employed is Bayesian Nash Equilibria. In particular, (i) S 's strategy $\sigma(m|x)$ specifies a distribution of messages for each x and R 's strategy $y(m)$ specifies an action for each message m ; and (ii) S chooses $\sigma(m|x)$ to maximize his expected utility given $y(m)$ and R chooses $y(m)$ to maximize his expected utility given $\sigma(m|x)$. Since in equilibrium each message will correspond to an action we can ease notation by interpreting the message m as a recommended action¹. Thus $\mathbf{M} = \mathbf{Y}$ and we can reduce S 's strategy to a deterministic function $m(x)$, which specifies a message $m \in Y$ for each signal $x \in \mathbf{S}$.

¹In particular, any other equilibrium induces the same relationship between x and y and is thus economically equivalent. Crawford and Sobel [13] refer to this as the equilibrium being *essentially unique*.

A *model* is constituted by the 5-tuple $\{u^i(\cdot), p^i(\cdot), \lambda(\cdot)\}_{i=R,S}$. The purpose of the paper is to investigate “common prior and non-common utility” and “non-common prior and common utility” models, which we define precisely as follows.

Definition *Common Prior Beliefs* (CB). A CB model is defined by the two players having a common prior, i.e. $p^R(\theta) = p^S(\theta)$ for all $\theta \in \Theta$.

Definition *Common Preferences* (CP). A CP model is defined by the two players having common preferences, i.e. $u^R(y, \theta) = u^S(y, \theta)$ for all $y \in \mathbf{Y}$ and $\theta \in \Theta$.

The exercise we want to perform consists of starting with a CB model, characterized by $\{u^i(\cdot), p(\cdot), \lambda(\cdot)\}$, and investigate whether we can find a CP model, $\{u(\cdot), p^i(\cdot), \lambda(\cdot)\}$, that is equivalent in a sense we will make precise in the next section.

3.2.2 A Measure of Equivalence

Definition

In this section we introduce two measures of equivalence between models. First, models that are equivalent in the sense that players want to take the same action for all $x \in \mathbf{S}$. We refer to this as Choice-Equivalence. This is an interesting measure in that it tells us when two models deliver the same individual behavior of the players if they act on their own. One of the points we want to make is that this is neither necessary nor sufficient for them to act the same in the two models when we introduce the information transmission game. In particular, even if a CP model and a CB model are Choice-Equivalent the equilibrium strategies that obtain in each model need not be the same. Therefore we define a second measure of equivalence, which is that in both models the two players have the same equilibrium strategies. We refer to this as Strategy-Equivalence. The two measures are defined formally as follows.

Definition *Choice-Equivalence*. Two models are Choice-Equivalent if the choice functions $y_*^i(x)$, $i = R, S$, are the same in both models for all $x \in \mathbf{S}$.

Definition *Strategy-Equivalence*. Two models are Strategy-Equivalent if in equilibrium the strategy functions $m(x)$ and $y(m)$ are the same in both models for all $m \in \mathbf{M}$ and $x \in \mathbf{S}$.

Two models are *equivalent* if they satisfy both Choice- and Strategy-Equivalence. The contrast between the two types of equivalence is that for Choice-Equivalence the players' prior expectations over the distribution of the signal $l^i(\cdot)$ do not matter - once x is realized their preferred action depends only on the conditional distribution of θ given x . In the information transmission game this is not so, since in general S 's equilibrium strategy will be to only reveal the partition of \mathbf{S} to which the signal belongs (as in Crawford and Sobel). Hence, R must infer the distribution of x conditional on this partition and for that purpose he uses $l^i(\cdot)$. This places a strong restriction on the set of priors and utility functions we can consider and hence in some cases it may be that there exist no utility function and prior beliefs that yield Strategy-Equivalence.

Examples

We offer two examples to illustrate the difference between the two kinds of equivalence. Suppose that $\mathbf{Y} = \mathbf{S} = \{L, H\}$ and $\Theta = \{L, M, H\}$ with the dgp given by $\lambda(L|L) = \lambda(H|H) = 1$ and $\lambda(M|H) = \lambda(M|L) = 1/2$. Hence, with probability a half the signal reveals the true state and with probability a half the signal is $x = M$, which is uninformative about the state. Suppose a CB model in which R 's and S 's utilities are given by $u^R(H, H) = 2/3$, $u^S(H, H) = 4/3$, $u^R(L, L) = u^S(L, L) = 1$ and $u^R(s, t) = u^S(s, t) = 0$ for $s, t \in \{L, H\}$ and $s \neq t$. The common prior gives equal weight to both states, $p(H) = 1/2$. This implies that the common beliefs over the distribution of x are $l(L) = l(H) = 1/4$ and $l(M) = 1/2$. Obviously $y_*^i(L) = L$ and $y_*^i(H) = H$ since these signals are completely informative. Moreover, $y_*^R(M) = L$ and $y_*^S(M) = H$ as the players attach equal probability to both states when $x = M$ but have different utilities. Thus there is a conflict of interest.

In an informative equilibrium the only credible strategy for S is $m(L) = L$ and $m(M) = m(H) = H$. Given $x \in \{M, H\}$ the posterior probability of the high state is

$$Pr(\theta = H|x \in \{M, H\}) = [(1/2) * (1/2) + (1/4) * (1)] / (1/2 + 1/4) = 2/3.$$

Thus conditional on $m = H$ then R has expected utility $(2/3) * (2/3) = 4/9$ from choosing $y = H$ and expected utility $(1/3) * (1) = 1/3$ from choosing $y = L$ and he chooses $y = H$. Hence there exists an informative equilibrium with $m(x)$ as above and

$y(m) = m$. Notice that the signals $x = L$ and $x = H$ reveal the true state perfectly and therefore, regardless of beliefs, players will take the same actions conditional on these signals. It follows that for Choice-Equivalence we need only investigate behavior when $x = M$. The first example is a rather trivial demonstration of the fact that we can obtain the same equilibrium strategies with a CP model even if the choice functions are not the same.

Example 2. Suppose a CP model with $u(\cdot) = u^S(\cdot)$ and $p^R(\cdot) = p^S(\cdot) = p(\cdot)$. Clearly, $y_*^S(\cdot)$ is the same as in the CB model above and therefore also $m(\cdot)$. Since $y_*^R(\cdot) = y_*^S(\cdot)$ the equilibrium is obviously informative. Hence, the two models are Strategy-Equivalent. But now $y_*^R(L) = L$ and $y_*^R(M) = y_*^R(H) = H$ so the models are not Choice-Equivalent.

This example is particularly simple since everything is discrete. R 's strategy specifies a discrete action as a function of the discrete message he receives from S . All we do in the example is to maintain S 's strategy and vary R 's choice function. This implies that the two models are not Choice-Equivalent. But since we have only varied incentives slightly the equilibrium of the information transmission game remains the same and therefore there is Strategy-Equivalence. This is perhaps not so surprising. It is more counterintuitive that we can construct two models that are Choice-Equivalent but not Strategy-Equivalent. The next example illustrates this.

Example 3. Suppose a CP model with $p^S(\cdot) = p(\cdot)$ and $u(\cdot) = u^S(\cdot)$ such that $y_*^S(\cdot)$ and $m(\cdot)$ are as in the CB model above. But now suppose $p^R(H) = 1/6$. Clearly $y_*^R(\cdot)$ is as before and therefore the models are Choice-Equivalent. But now $l^R(L) = 5/12$, $l^R(M) = 1/2$ and $l^R(H) = 1/12$. R 's estimated probability of $\theta = H$ conditional on $x = H$ is 1 and conditional on $x = M$ is $1/6$ and hence

$$Pr^R(\theta = H | x \in \{M, H\}) = [(1/2) * (1/6) + (1/12) * (1)] / (1/2 + 1/12) = 2/7.$$

Thus when $m = H$ then R has expected utility $(2/7) * (4/3) = 8/21$ from choosing $y = H$ and expected utility $(3/7) * (1) = 9/21$ from choosing $y = L$. Therefore R always chooses $y = L$, no matter what message S sends and the models are not Strategy-Equivalent.

In both models S has a lower cutoff point for switching from $y = L$ to $y = H$ than does R . In the CB model this is so because S has more utility than R from correctly

predicting $\theta = H$. But in the CP model this occurs because S places a higher prior probability on $\theta = H$. Since R 's utility from correctly predicting $\theta = H$ is higher in the CP model, Choice-Equivalence obtains only because R places a much lower probability on $\theta = H$ in the CP model than in the CB model. But this in turn is exactly the reason that Strategy-Equivalence fails, since this probability affects R 's expectation of the distribution of the signal. Conditional on $m = H$, in the CB model R infers that the probability of $\theta = H$ is twice as high as the probability of $\theta = L$ and therefore he chooses $y = H$. In the CP model this ratio is $2/5$ and therefore he chooses $y = L$. Although R would take the same action conditional on x in the two models, in the CP model he always wants to take the low action regardless of what message S sends him, which implies that the equilibrium is uninformative and Strategy-Equivalence does not obtain.

3.3 A Discrete World

In this section we discretize the action and state spaces to derive some illustrative results. In the most general version we can describe the state space by $\Theta = \{\theta_1, \dots, \theta_T\}$ and the action space by $Y = \{y_1, \dots, y_K\}$. We retain the continuous signal space, and to fix ideas suppose $\mathbf{S} = \mathbb{R}$ although this is not important. Priors are denoted by $p_t \equiv p(\theta_t)$ and posteriors by $p_t(x) \equiv p(\theta_t|x)$ for $t = 1, \dots, T$. Utilities are similarly $u_{k,t} \equiv u(y_k, \theta_t)$ for $t = 1, \dots, T$ and $k = 1, \dots, K$. As before, probabilities and utilities are denoted by a superscript $i = R, S$ when they are non-common. We denote the dgp by $\lambda_t(x) \equiv \lambda(x|\theta_t)$. Expected utilities in the CB model are

$$U_k^i(x) \equiv U^i(y_k, x) = \sum_{t=1}^T p_t(x) u_{k,t}^i,$$

and in the CP model

$$\bar{U}_k^i(x) \equiv \bar{U}^i(y_k, x) = \sum_{t=1}^T p_t^i(x) u_{k,t}.$$

One important restriction in this section is that we focus on cases where the most informative equilibrium exists, i.e. the equilibrium where all actions are taken with positive probability.

3.3.1 The Simplest Case

We start by considering the simplest possible case with two actions and two states of the world, i.e. $Y = \{y_L, y_H\}$ and $\Theta = \{\theta_L, \theta_H\}$. Then $p_L \equiv p(\theta_L)$ and $p_H \equiv 1 - p_L$ denote the priors and $p_L(x) \equiv p(\theta_L|x) = p_L \lambda_L(x) / (p_L \lambda_L(x) + (1 - p_L) \lambda_H(x))$ and $p_H(x) = 1 - p_L(x)$ the posterior probabilities conditional on x . It turns out that in this case, for any CB model we can find an equivalent CP model.

Proposition 7. *Suppose $Y = \{y_L, y_H\}$, $\Theta = \{\theta_L, \theta_H\}$ and $\mathbf{S} = \mathbb{R}$. Any CB model always has an equivalent CP model.*

Proof. For the proof we assume that the inequality of the increasing likelihood ratio property (ILR) of the dgp is strict for all x and θ . The proof is similar in the case where we allow for equality. First we define the “arbitrage functions” for the CB model $V^i(x) \equiv U_L^i(x) - U_H^i(x)$ and for the CP model $\bar{V}^i(x) \equiv \bar{U}_L^i(x) - \bar{U}_H^i(x)$. By ILR and the sorting condition, the optimal action is increasing in x and for each model there will be a unique cutoff point, c^i and \bar{c}^i , defined by $V^i(c^i) = 0$ and $\bar{V}^i(\bar{c}^i) = 0$, such that the players are indifferent between the two actions. Players will prefer the low action when x is below the cutoff point and the high action when x is above.

We start by showing that for any CB model we can always find a Choice-Equivalent CP model characterized by a utility function $u_{k,t}$ and the prior probabilities p_L^i . The sufficient and necessary condition for the CB and the CP models to be Choice-Equivalent is $c^i = \bar{c}^i$, i.e. that $\bar{V}^i(c^i) = 0$. Let $u_{k,k} > u_{k,t}$ for $k, t = L, H$ and $k \neq t$ and write out the arbitrage function in the CP case.

$$\bar{V}^i(x) = p_L^i(x)(u_{L,L} - u_{H,L}) + (1 - p_L^i(x))(u_{L,H} - u_{H,H}).$$

Observe that $\bar{V}^i(\cdot) = u_{L,H} - u_{H,H} < 0$ for $p_L^i = 0$ and $\bar{V}^i(\cdot) = u_{L,L} - u_{H,L} > 0$ for $p_L^i = 1$. Since $\bar{V}^i(\cdot)$ is strictly decreasing in $p_L^i(\cdot)$ and $p_L^i(\cdot)$ is strictly increasing in x then for any c there exists a unique p_L^i such that $\bar{V}^i(c) = 0$. This assures us that for any CB model there exists a Choice-Equivalent CP model.

We then prove Strategy-Equivalence. Since we have assumed that we are in the most informative equilibrium, both actions must be taken. The only credible strategy for S

is to reveal his preferred action. If he did not do so, there would exist x for which he could profitably deviate. S 's equilibrium message strategy in the CB model is hence

$$m(x) = \begin{cases} y_L & \text{if } x \leq c^S \\ y_H & \text{if } x > c^S. \end{cases} \quad (3.1)$$

Similarly, let the message strategy in the CP model be denoted by $\bar{m}(x)$. Clearly, if $c^i = \bar{c}^i$ then $m(x) = \bar{m}(x)$ and hence Choice-Equivalence implies Strategy-Equivalence for S . Furthermore, let $y(m)$ and $\bar{y}(m)$ be R 's equilibrium action strategies in the CB and the CP models, respectively. Since the CB equilibrium is informative then $y(m) = m$. This implies that

$$\begin{aligned} \int_{-\infty}^{c^S} U_L^R(x) l^R(x) dx &\geq \int_{-\infty}^{c^S} U_H^R(x) l^R(x) dx, \text{ and} \\ \int_{c^S}^{\infty} U_H^R(x) l^R(x) dx &\geq \int_{c^S}^{\infty} U_L^R(x) l^R(x) dx. \end{aligned} \quad (3.2)$$

Strategy-equivalence for R obtains if $\bar{y}(m) = m$, which implies

$$\begin{aligned} \int_{-\infty}^{\bar{c}^S} \bar{U}_L^R(x) l^R(x) dx &\geq \int_{-\infty}^{\bar{c}^S} \bar{U}_H^R(x) l^R(x) dx, \text{ and} \\ \int_{\bar{c}^S}^{\infty} \bar{U}_H^R(x) l^R(x) dx &\geq \int_{\bar{c}^S}^{\infty} \bar{U}_L^R(x) l^R(x) dx. \end{aligned} \quad (3.3)$$

The procedure is then the following. We start with a CB model which is characterized by c^R and c^S and which is by assumption informative such that $m(x)$ is given by (3.1) and $y(m) = m$. Assume $\bar{c}^S = c^S$ such that the CB and CP models are Choice- and Strategy-Equivalent for S . Then notice that we can always choose $u_{k,t} = u_{k,t}^R$ and $p_L^R = p_L$ which implies that $\bar{c}^R = c^R$ and hence the CB and CP models are Choice-Equivalent for R . Furthermore, (3.2) and (3.3) are the same and therefore the models are also Strategy-Equivalent for R . Finally, we just need to remark that, by the arguments above, for any c^S and $u_{k,t}$, we can find p_L^S to ensure $\bar{V}^S(c^S) = 0$ such that $\bar{c}^S = c^S$. In conclusion, for any CB model we can find a CP model which is Choice- and Strategy-Equivalent. This yields the proposition. \square

However, this result not generally true and we need look no further than to models with three discrete actions to find a counterexample as the next section shows.

3.3.2 The Size of the State Space Matters

In this section we look at two examples with a size-3 action space. Our strategy is the same as in the proof of the previous section. The sorting condition and ILR property of the dgp assure us the existence of a set of cutoff points, $\{c_k^i\}_{k=1}^{K-1}$, defined by the CB model and $\{\bar{c}_k^i\}_{k=1}^{K-1}$ defined by the CP model, such that the optimal message strategy for S in the CB model is $m(x) = y_k$ for $x \in (c_{k-1}^i, c_k^i]$ and in the CP model $\bar{m}(x) = y_k$ for $x \in (\bar{c}_{k-1}^i, \bar{c}_k^i]$. As before define the functions $V_k^i(x) \equiv U_k^i(x) - U_{k+1}^i(x)$ and $\bar{V}_k^i(x) \equiv \bar{U}_k^i(x) - \bar{U}_{k+1}^i(x)$ for $k = 1, \dots, K-1$.

First, we will show a counter-example to Proposition 7 where Choice-Equivalence fails in all but a small number of cases. Let $K = 3$ and $T = 2$ and recall that we have supposed that the equilibrium is fully informative, i.e., all three actions are taken with positive probability. This requires that $u_{1,1} > u_{k,1}$ for $k = 2, 3$, $u_{3,2} > u_{k,2}$ for $k = 1, 2$ and $u_{2,s} > u_{k,s}$ for some s and $k \in \{1, 3\}$. Otherwise at least one action would never be strictly preferred. Given $\{c_k^i\}_{k=1}^{K-1}$ defined by the CB model, the requirement for Choice-Equivalence is to find beliefs and a common utility function that make $\bar{V}_k^i(c_k^i) = 0$ for $k = 1, \dots, K-1$ and $i = R, S$. Write out the values of the arbitrage function $\bar{V}_k^i(\cdot)$.

$$\begin{aligned}\bar{V}_1^i(x) &= p_1^i(x)u_{1,1} + (1 - p_1^i(x))u_{1,2} - [p_1^i(x)u_{2,1} + (1 - p_1^i(x))u_{2,2}], \text{ and} \\ \bar{V}_2^i(x) &= p_1^i(x)u_{2,1} + (1 - p_1^i(x))u_{2,2} - [p_1^i(x)u_{3,1} + (1 - p_1^i(x))u_{3,2}].\end{aligned}$$

The conditions $\bar{V}_k^i(c_k^i) = 0$ can then be written as

$$u_{2,1} = u_{1,1} + \frac{1 - p_1^i(c_1^i)}{p_1^i(c_1^i)}(u_{1,2} - u_{2,2}), \text{ and} \quad (3.4)$$

$$u_{3,1} = u_{2,1} + \frac{1 - p_1^i(c_2^i)}{p_1^i(c_2^i)}(u_{2,2} - u_{3,2}). \quad (3.5)$$

Since the equilibrium is informative by assumption we must have $p_1^i \in (0, 1)$. Otherwise, some action is never taken. For the two conditions to hold for both R and S simultaneously we must have $p_1^R(c_1^R) = p_1^S(c_1^S)$ and $p_1^R(c_2^R) = p_1^S(c_2^S)$. First consider the case where $c_1^R > c_1^S$. Then we must require $p_1^R > p_1^S$ to solve (3.4). If furthermore $c_2^R < c_2^S$ then we must have $p_1^R < p_1^S$ to satisfy (3.5). This is obviously a contradiction and hence Choice-Equivalence is not always possible. Let $\lambda_{t,k}^i \equiv \lambda_t(c_k^i)$. We can then rewrite the two conditions to arrive at the following proposition.

Proposition 8. *Suppose $T = 2$ and $K = 3$. A CB model has a Choice-Equivalent CP model if and only if*

$$\frac{\lambda_{1,1}^R \lambda_{2,2}^R}{\lambda_{1,2}^R \lambda_{2,1}^R} = \frac{\lambda_{1,1}^S \lambda_{2,2}^S}{\lambda_{1,2}^S \lambda_{2,1}^S}.$$

A particular case where this fails is whenever the preferences exhibit “crossover” (i.e., $c_1^R - c_1^S$ and $c_2^R - c_2^S$ have opposite signs).

Proof. See Appendix. □

Proposition 8 shows that the set of Choice-Equivalent models is very small in this particular case. Choice-Equivalence requires us to solve four equations and although in principle we have many choice parameters we are restricted by the fact that the utilities are pairwise the same in the equations. Therefore a solution exists only if the posterior state probabilities at the cutoff points are equal and this leads to the condition of the proposition.

Since the prior is defined over the states, the size of the state space determines the flexibility we have in specifying the CP model. So we would expect the existence of equivalent models to depend to some extent on T . The next result confirms this. Suppose $T = K = 3$. Our approach is the same as before. It is assumed that all actions are taken in equilibrium and therefore $u_{k,k} > u_{k,t}$ for $k = 1, 3$ and $k \neq t$. Let $p_{t,k}^i \equiv p_t^i(c_k^i)$ and $\hat{p}_{t,k}^i \equiv p_t^i \lambda_{t,k}^i$. Manipulating the conditions $\bar{V}_k^i(c_k^i) = 0$ for $k = 1, 2$ we obtain

$$\hat{p}_{1,1}^i(u_{1,1} - u_{2,1}) + \hat{p}_{2,1}^i(u_{1,2} - u_{2,2}) + \hat{p}_{3,1}^i(u_{1,3} - u_{2,3}) = 0 \quad (3.6)$$

$$\hat{p}_{1,2}^i(u_{2,1} - u_{3,1}) + \hat{p}_{2,2}^i(u_{2,2} - u_{3,2}) + \hat{p}_{3,2}^i(u_{2,3} - u_{3,3}) = 0. \quad (3.7)$$

Cancel out $u_{2,3}$ and write out $\hat{p}_{t,k}^i \equiv p_t^i \lambda_{t,k}^i$ to get

$$\begin{aligned} p_1^i \frac{\lambda_{1,1}^i}{\lambda_{3,1}^i} (u_{1,1} - u_{2,1}) + p_2^i \frac{\lambda_{2,1}^i}{\lambda_{3,1}^i} (u_{1,2} - u_{2,2}) + p_1^i \frac{\lambda_{1,2}^i}{\lambda_{3,2}^i} (u_{2,1} - u_{3,1}) \\ + p_2^i \frac{\lambda_{2,2}^i}{\lambda_{3,2}^i} (u_{2,2} - u_{3,2}) + (1 - p_1^i - p_2^i)(u_{1,3} - u_{3,3}) = 0. \end{aligned} \quad (3.8)$$

First let $p_1^i, p_2^i \rightarrow 0$. Then the left-hand side goes to $(u_{1,3} - u_{3,3}) < 0$. Second let $p_1^i \rightarrow 1$ and $p_2^i \rightarrow 0$. Then the left-hand side goes to $(u_{1,1} - u_{2,1}) > 0$. Hence, there exist p_1^i and

p_2^i such that (3.8) is true. Thus, for any CB model we can find a Choice-Equivalent CP model. Since the exact values of the utilities do not matter for the existence result, we can use the same argument as in the proof of the case with two actions and two states. Hence, we choose $u_{k,t} = u_{k,t}^R$ and $p_t^R = p_t$ for $k, t = 1, 2, 3$ to make the CP model Choice- and Strategy-Equivalent for R and then use the above argument to find p_t^S such that $\bar{c}_k^S = c_k^S$. This yields the following proposition which we state without further proof.

Proposition 9. *Suppose $T = K = 3$. Any CB model always has an equivalent CP model.*

This proposition confirms the previous intuition: with a larger state space we can achieve equivalence. The explanation is simple. To achieve Choice-Equivalence we need to match two cutoff points for each player and for that purpose we need, in general, two free variables *that differ* between the players. With three states we can decide two of the state probabilities freely and by choosing the utilities appropriately we can also make the problem Strategy-Equivalent, which gives us the result.

3.3.3 A Sufficiency Result for Equivalence

In the previous section we found that with 3 actions and 2 states Choice-Equivalence could only be obtained in a small set of cases whereas with 3 states both kinds of equivalence is always attainable. The next step is to use the intuition from these two examples to establish a sufficiency condition for equivalence.

Assume that the model is as before, except that T and K may take any value. We then want to establish a sufficiency condition on T and K such that we can assure the existence of equivalent models. Our approach is to formulate the entire equivalence problem as a system of equations which is linear in $u_{t,k}$. As these are utilities and can be both negative and positive we know that a solution exists conditional on certain properties of the augmented matrix of the problem. We start by “stacking” all the utilities into a KT -dimensional column vector $\mathbf{u} = (u_1, \dots, u_K)'$ where $u_k = (u_{1,k}, \dots, u_{T,k})'$. Using the notation from above, $p_{t,k}^i$ is the probability that i assigns to state t when $x = c_k^S$. We write the vector of such probabilities for each state as $P_k^i = (p_{1,k}^i, \dots, p_{T,k}^i)$.

We then want the solution to satisfy two sets of conditions. First, the arbitrage conditions which for i correspond to $\sum_t p_{t,k}^i (u_{t,k} - u_{t,k+1}) = 0$ for $k = 1, \dots, K-1$. These conditions yield Choice-Equivalence for i and we can write them as $\mathbf{C}^i \mathbf{u} = \mathbf{0}$ where $\mathbf{0}$ denotes a $K-1$ dimensional column vector of zeros and \mathbf{C}^i is a $(K-1) \times KT$ matrix given by

$$\mathbf{C}^i = \begin{pmatrix} P_1^i & -P_1^i & 0 & 0 & \cdots & 0 & 0 \\ 0 & P_2^i & -P_2^i & 0 & \cdots & 0 & 0 \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \\ 0 & 0 & 0 & 0 & \cdots & P_{K-1}^i & -P_{K-1}^i \end{pmatrix}.$$

Second, to assure Strategy-Equivalence for R we want to fix the expected utility in each partition of the message space. Let

$$\bar{p}_{t,k}^R \equiv \frac{\int_{c_k^S}^{c_{k+1}^S} p_t^R(x) l^R(x) dx}{\sum_t \int_{c_k^S}^{c_{k+1}^S} p_t^R(x) l^R(x) dx},$$

be the conditional probability of state t in partition k . Construct the vector \bar{P}_t^R as above. We want to find $u_{t,l}$ such that $\sum_t \bar{p}_{t,k}^R u_{t,l} = v_{k,l}$ for $k, l = 1, \dots, K$, where $v_{k,l}$ is chosen to assure Strategy-Equivalence. Let $\mathbf{v} = (v_1, \dots, v_K)'$ and $v_k = (v_{1,k}, \dots, v_{K,k})'$. It is enough for Strategy-Equivalence that $\mathbf{S}^R \mathbf{u} = \mathbf{v}$, where $\mathbf{S}^R = [S_1^R, \dots, S_K^R]'$ is a $K^2 \times KT$ matrix with the elements $S_k^R = \text{diag}\{\bar{P}_k^R\}$ being $K \times KT$ block-diagonal matrices. A sufficient condition for equivalence is that there exists a vector \mathbf{u} that solves

$$\begin{pmatrix} \mathbf{C}^S \\ \mathbf{C}^R \\ \mathbf{S}^R \end{pmatrix} \mathbf{u} = \begin{pmatrix} \mathbf{0} \\ \mathbf{0} \\ \mathbf{v} \end{pmatrix}.$$

Let $\mathbf{E} = (\mathbf{C}^S \mathbf{C}^R \mathbf{S}^R)'$ and $\mathbf{e} = (\mathbf{0} \mathbf{0} \mathbf{v})'$. The Rouché-Capelli Theorem assures the existence of a solution if the rank of the matrix \mathbf{E} is equal to the rank of the augmented matrix $(\mathbf{E}:\mathbf{e})$. In the proof of the following result we show that this problem boils down to the rank of the matrices $R_k = (-P_{k-1}^S, P_k^S, -P_{k-1}^R, P_k^R, \bar{P}_2^R, \dots, \bar{P}_K^R)'$ for $k = 1, \dots, K-1$ and $R_K = (-P_{K-1}^S, -P_{K-1}^R, \bar{P}_2^R, \dots, \bar{P}_K^R)'$. These must have rank at least $K+1$. Notice that these matrices depend on the function $\lambda(\cdot)$, which is given, and

the prior probabilities p_t^i about which we have made no assumptions so far. To apply the theorem we make the following assumption which for clarity of exposition we make slightly more restrictive (assuming full rank) than necessary.

Assumption 2. *There exist prior probabilities $\{p_t^R\}$ and $\{p_t^S\}$ such that the matrices R_k , $k = 1, \dots, K$, have full rank.*

Since the elements of the matrices R_k depend non-linearly on the prior probabilities, we expect in general to be able to find a vector of probabilities such that the rank condition is satisfied. This leads us to the following result.

Proposition 10. *Suppose Assumption 2. For any CB model, a sufficient condition for the existence of an equivalent CP model is $T \geq K + 1$.*

Proof. See Appendix. □

One immediate corollary to this result is that a CB model with a discrete action space and a continuous state space will always have a Choice- and Strategy-Equivalent CP model, since we can always specify a prior that is degenerate in a finite number of points.

3.4 Equivalence for Continuous Models

When we allow for a continuous actions space, the action taken will depend directly on R 's beliefs and therefore we require even stronger conditions for equivalence. To explore this we assume (denoting partial derivatives in the usual manner) that for each x then $U_{22}^i(\cdot) < 0$ for all y and $y_*^R(\cdot) - y_*^S(\cdot)$ is either strictly positive or strictly negative for all x . These are the assumptions made by Crawford and Sobel [13] and their results show that the (essentially) only equilibrium is a finite partition of the message space such that $m(x) = m_n$ whenever $x_{n-1} < x \leq x_n$, for $n = 1, \dots, N$. The maximum N for which an equilibrium exists depends on the alignment of the players' preferences. Let $y(a, b) \equiv \arg \max_y \int_a^b U^R(y, x) l^R(x) dx$ and $y_n = y(x_{n-1}, x_n)$. The cutoff points $\{x_n\}_{n=0}^N$ must satisfy the condition $U^S(y_n, x_n) = U^S(y_n, x_{n-1})$.

Two models A and B that satisfy the above derivative assumptions are Strategy-Equivalent if and only if the following two conditions are satisfied.

(SE1) Fixing $\{y_n\}_{n=1}^N$, S 's optimal cutoff points $\{x_n\}_{n=0}^N$ are the same in A and B .

(SE2) Fixing $\{x_n\}_{n=0}^N$, R 's optimal actions $\{y_n\}_{n=1}^N$ are the same in A and B .

First, we say that two models are *EU-Equivalent* if they produce the same expected utilities conditional on x . We can quickly derive the following lemma.

Lemma 7. *A sufficient condition for (SE1) is that models A and B are Choice-Equivalent for S . Furthermore, it is sufficient for (SE2) that A and B are EU-Equivalent for R and that for some $k_n > 0$*

$$l_A^R(x) = k_n l_B^R(x) \text{ for all } x \in (x_{n-1}, x_n). \quad (3.9)$$

Proof. See Appendix. □

The condition for (SE2) says that the conditional distribution of x in the interval (x_{n-1}, x_n) remains unchanged. In effect, R can only change his expectation of the probability of the intervals in the partition but not of the relative likelihood of two signals within a given partition. Hence, for a given $\lambda(\cdot)$ we are very restricted in our choice of $p^i(\cdot)$ if we want to achieve Strategy-Equivalence between two models. We start out by digressing a bit further on EU-Equivalence, showing that it may be impossible to obtain for some $\lambda(\cdot)$. EU-Equivalence for a given $u^i(\cdot)$, $\lambda(\cdot)$ and $p(\cdot)$ requires us to find $u(\cdot)$ and $p^i(\cdot)$ such that for any x and y

$$\int_{\Theta} u^i(y, t) p(t|x) dt = \int_{\Theta} u(y, t) p^i(t|x) dt. \quad (3.10)$$

Without requiring $\lambda(\cdot)$ to be common between players and models it is easy to obtain EU-Equivalence. Suppose that the function $u(y, \theta)$ is continuous in both its arguments and that for $i = R, S$ then $\min_{\theta} u(y, \theta) \leq \inf_x U^i(y, x)$ and $\max_{\theta} u(y, \theta) \geq \sup_x U^i(y, x)$ for all y . Hence, for all (y', x') there exists $\theta'_i \in \Theta$ such that $U^i(y', x') = u(y', \theta'_i)$. Then we can choose $\lambda(\cdot)$ such that $p^i(\theta'_i|x') = 1$ and $p^i(\theta|x') = 0$ for $\theta \neq \theta'_i$, and thus (3.10) is true. Supposing that $\lambda(\cdot)$ is fixed the matter is more complicated. The following two examples illustrate the difficulties involved.

Example 4. Suppose $\Theta = \mathbf{S} = [0, 1]$ and $\lambda(x|\theta) = 1$ for all x and θ . I.e., conditional on θ the signal x is uniformly distributed on the unit interval. Then $p^i(\theta|x) = p^i(\theta)$ and

$$\int_{\Theta} u^i(y, t) p(t|x) dt = \int_{\Theta} u(y, t) p^i(t) dt = C(y).$$

Expected utility is constant for all x and hence if $U^i(y, x) \neq U^i(y, x')$ for some triple (y, x, x') EU-Equivalence fails.

In this example the dgp is too vague. No information is revealed by the signal and hence expected utility depends only on prior beliefs over the state. The next example shows the contrarian case.

Example 5. Suppose $\Theta = \mathbf{S} = [0, 1]$ with $\lambda(\theta|\theta) = 1$ and $\lambda(x|\theta) = 0$ for $x \neq \theta$. Then expected utility conditional on x is $u(y, x)$ and hence if $U^R(y, x) \neq U^S(y, x)$ for some pair (y, x) EU-Equivalence fails.

Now the signal is too informative and completely removes the prior beliefs from the expected utility, thereby taking away any flexibility we have for specifying individual behavior in the CP model. These two extreme examples illustrate the problems when the dgp is too peaked or too flat. Assuming that $\lambda(\cdot)$ has been given together with $p(\cdot)$ and $u^i(\cdot)$, we want to investigate which $u(\cdot)$ are admissible to assure (SE2), which will subsequently identify a class of CB models which have equivalent CP models. First, we notice that this class includes any CB model in which $u^S(\cdot)$ satisfies the following condition for some $\hat{p}(\cdot)$.

$$\int_{\Theta} u^R(y, t) \hat{p}(t|x) dt = \int_{\Theta} u^S(y, t) p(t|x) dt.$$

In this case, we can set $u(\cdot) = u^R(\cdot)$ and $p^R(\cdot) = p(\cdot)$ to assure Choice- and Strategy-Equivalence for R and $p^S(\cdot) = \hat{p}(\cdot)$ to assure the same for S . In fact, condition (3.9) restricts our choice of $p^R(\cdot)$ such that in general the only solution will be to set $p^R(\cdot) = p(\cdot)$, leaving as free variables only $u(\cdot)$ and $p^S(\cdot)$. Suppose $u(y, \theta) \neq u^R(y, \theta)$ for some y and θ . Then we can define $a(y, \theta) \equiv u(y, \theta) - u^R(y, \theta)$. The expected utility in the CP model is thus

$$\int_{\Theta} [u^R(y, t) + a(y, t)] \frac{p^R(t) \lambda(x|t)}{\int_{\Theta} p^R(t') \lambda(x|t') dt'} dt.$$

Clearly, this is equal to $U^R(\cdot)$ only if for all y and x

$$\int_{\Theta} a(y, t) \frac{p^R(t) \lambda(x|t)}{\int_{\Theta} p^R(t') \lambda(x|t') dt'} dt = 0.$$

In general this condition will be hard to satisfy unless we impose some more structure on the distributions. We say that $\lambda(\cdot)$ satisfies the Partial Constant Likelihood Property in an interval T if for all pairs (x, x') and $\theta, \theta' \in T$ we have

$$\frac{\lambda(x|\theta')}{\lambda(x|\theta)} = \frac{\lambda(x'|\theta')}{\lambda(x'|\theta)}. \quad (\text{PCLR})$$

The property says that the signal is uninformative about the relative likelihood of states in T . Examples of distributions that satisfy (PCLR) in some interval is (i) any distribution for which $\lambda(x|\theta)$ is equal to a constant times $e^{-(\theta-x)}$ for $\theta \in T$ or (ii) a distribution with three states $\{A, B, C\}$ and $x \in [0, 1]$ where $\lambda(x|A) = x$ and $\lambda(x|B) = \lambda(x|C) = \frac{1-x}{2}$. In the last example, for any x then $\lambda(x|B)/\lambda(x|C) = 1$ such that $T = \{B, C\}$. We use this property to state the following proposition which we prove in the appendix.

Proposition 11. *For some T define $a(y, \theta)$ with $a(y, \theta) = 0$ for $\theta \notin T$ and*

$$\int_T a(y, t) \frac{p(t) \lambda(x|t)}{\int_{\Theta} p(t') \lambda(x|t') dt'} dt = 0 \text{ for some } x.$$

Suppose that in a CB model $\lambda(\cdot)$ satisfies (PCLR) for T and there exists $p^S(\cdot)$ such that

$$U^S(y, x) = \int_{\Theta} [u^R(y, t) + a(y, t)] \frac{p^S(t) \lambda(x|t)}{\int_{\Theta} p^S(t') \lambda(x|t') dt'} dt. \quad (3.11)$$

Then there exists an equivalent CP model.

Proof. See Appendix. □

Lemma 7 identifies a very strict set of sufficiency conditions for equivalence, which the above Proposition relaxes somehow, although we can only apply it to very particular dgp. The difficulty of finding equivalent models in continuous space underpins is important - it appears that often an equivalent model will not exist. The culprit is R 's expectation of the dgp, which is formed using his prior. Since any heterogeneity in individual behavior must be modeled by priors and priors affect strategic behavior as well, it seems intuitive that often it will be impossible to model a particular individual and strategic behavior simultaneously.

3.5 Preferences Over Information: An Example

In this section we show an example which demonstrates that even when a CB model has an equivalent CP model, this CP model may be different in other respects. In particular, we will investigate a model where S is characterized by how precise a signal he observes. Following Kawamura [29] we call this his “competence”. We show that even when a CB and CP model satisfy Choice- and Strategy-Equivalence R has a stronger preference for competence in the CP model than in the CB model. We can relate this to Che and Kartik [11] who analyze an information transmission model with information acquisition in which players are heterogenous either in preferences or beliefs. They show that, due to the information acquisition motive, a change in beliefs is different to a change in preferences, since a change in preferences affects only the preferred action whereas a change in beliefs affects the preferred action as well as the expectation of the information that will be acquired. Our focus here is different. Information acquisition is exogenous and the reason why R has a greater preference for S being more competent in the CP model is that the players’ ex post preferences converge as S becomes more competent and it turns out that they converge more in the CP model than in the CB model. Thus, although our result at first seems reminiscent of Che and Kartik, both the implications and the cause of the result are quite different.

As our example we take a simple situation with two actions and where players’ utilities depend linearly on the distance between the action y , the state θ and a bias parameter b . The state belongs to the unit interval and S receives a signal x which is completely informative with probability π and pure noise with probability $1 - \pi$ (thus π parametrizes S ’s competence). Since payoffs are linear we can characterize payoffs and strategies entirely in terms of the prior expectation of the state, m . I.e. the shape of the prior distribution does not matter. As in previous sections the players’ choice functions have a cutoff point such that they prefer the high action above and the low action below, and this cutoff point depends only on b , m and π . The equilibrium strategy of S will be to reveal his preferred action but not the exact value of x . Our approach will be to derive CP and CB models that are Choice- and Strategy-Equivalent and show that R has a larger utility gain from increasing π in the CP model.

We start out by describing the model in more detail. Initially we denote all variables

by a superscript and then move on to specify the CB and CP models below. The action space is given by $\mathbf{Y} = \{0, 1\}$ whereas the state and signal spaces are the unit interval, $\Theta = \mathbf{S} = [0, 1]$. For $i = R, S$, utility is given by $u^i(0, \theta) = b^i - \theta$ and $u^i(1, \theta) = \theta - b^i$. The prior over the state space is $p^i(\theta)$ with mean m^i . S receives a signal x that is equal to θ with probability π and otherwise distributed uniformly on $[0, 1]$, i.e. not informative about the state. Expected utility is thus

$$U^i(y, x) = \pi u^i(y, x) + (1 - \pi) \int_0^1 u^i(y, t) p^i(t) dt.$$

Let the arbitrage function be defined as in Section 3.3.1, which implies

$$\begin{aligned} V^i(x) &= u(0, x) - u(1, x) + (1 - \pi) \int_0^1 (u^i(0, t) - u^i(1, t)) p^i(t) dt \\ &= 2\pi(m^i - x) + 2(b^i - m^i). \end{aligned}$$

The cutoff point is identified by $V^i(c^i) = 0$ and hence in this case it will only depend on m^i , b^i and π .

$$c^i = \frac{b^i - m^i}{\pi} + m^i.$$

A CB model in this context is defined by $m^R = m^S = m$ and a CP model by $b^R = b^S = b$. If the players have common preferences prior (CP) then $\pi = 1$ implies $c^R - c^S = 0$ whereas if the preferences are non-common (CB) it implies $c^R - c^S = b^R - b^S$. This gives us a hint as to the comparative statics: when preferences are common the ex post bias disappears as S 's competence grows whereas when preferences are non-common there will always be ex post bias, regardless of S 's competence.

Choice- And Strategy-Equivalence. The condition for Choice-Equivalence in this setup is that the CB and CP models have the same cutoff point, c^i , which gives us the following.

$$\frac{b^i - m}{\pi} + m = \frac{b - m^i}{\pi} + m^i.$$

We always start with a CB model and an information structure, which in this case is given by m , b^i and π . The free parameters we have in the CP model are b and m^i . For

each b we can solve the previous condition for m^i to yield

$$m_*^i(b) = \frac{(b - b^i)}{1 - \pi} + m.$$

This is the values of m^i given b that assure Choice-Equivalence. Thus, for the CB model given by m , b^i and π we can identify a Choice-Equivalent CP model if there exists some b such that $m_*^R(b), m_*^S(b) \in [0, 1]$. Furthermore, this will be Strategy-Equivalent for S since his strategy is defined uniquely by c^S . To check that the models are Strategy-Equivalent for R we first remark that if $m_*^S(b^R) \in [0, 1]$ we can set $b = b^R$ and the models will be exactly the same for R and therefore also Strategy-Equivalent. More generally, we need to calculate R 's expected utility from receiving the signals 0 and 1 when S 's cutoff point is c .

$$\begin{aligned}\bar{U}_0^R(y) &\equiv \int_0^c \left[\pi u^R(y, t) + (1 - \pi) \int_0^1 u^R(y, s) p^R(s) ds \right] p^R(t) dt \text{ and} \\ \bar{U}_1^R(y) &\equiv \int_c^1 \left[\pi u^R(y, t) + (1 - \pi) \int_0^1 u^R(y, s) p^R(s) ds \right] p^R(t) dt.\end{aligned}$$

Strategy-Equivalence for R boils down to whether or not the equilibrium is informative, such that $y(m) = m$, or not. In the first case, we must have $\bar{U}_0^R(0) \geq \bar{U}_0^R(1)$ and $\bar{U}_1^R(1) \geq \bar{U}_1^R(0)$. In the second case, one of the inequalities fail. In a CB model we can rewrite the first inequality as

$$b^R \geq \pi t_0 + (1 - \pi)m, \quad (3.12)$$

and in a CP model

$$b \geq \pi t_0^R + (1 - \pi)m^R, \quad (3.13)$$

where $t_0^i \equiv E_{p^i}[\theta | \theta \leq c]$. Substituting m_*^R for m^R we can see that (3.12) and (3.13) are the same if $t_0^R = t_0$. Similarly, the inequality $\bar{U}_1^R(1) \geq \bar{U}_1^R(0)$ reduces, in a CB model, to $b^R \geq \pi t_1 + (1 - \pi)m$, and in a CP model to $b \geq \pi t_1^R + (1 - \pi)m^R$. Hence, the two are equivalent if $t_1^R = t_1$. We can always choose $p^i(\cdot)$ to assure $t_0^R = t_0$ and $t_1^R = t_1$. Let $P^i(c) \equiv \int_0^c p^i(t) dt$. Since $m^R = P^i(c)t_0 + (1 - P^i(c))t_1$ and we can choose $P^i(c)$ to be any value in the unit interval then we can achieve Strategy-Equivalence for R as long as there exists b such that $m_*^S(b) \in [0, 1]$ and $m_*^R(b) \in [t_0, t_1]$. The earlier observation tells us that this condition is also sufficient for Choice-Equivalence.

Preferences Over Informativeness. In this section we use the class of equivalent models identified in the previous section. Suppose a CB model with an informative size-2 equilibrium and that an equivalent CP model exists. R 's ex-ante expected utility conditional on S having cutoff point c and signal precision π is

$$EU^R(c, \pi) \equiv \bar{U}_0^R(0) + \bar{U}_1^R(1).$$

We are interested in how expected utility changes when we increase π and compare the changes between CB and CP models. Since in equilibrium c depends on π we first calculate the partial derivatives of the expected utility and then the total derivative with respect to π . The partial derivative of $EU^R(\cdot)$ with respect to π and holding c fixed is

$$P^R(c)(m^R - t_0^i) + (1 - P^R(c))(t_1^i - m^R).$$

After some rewriting the partial derivative of $EU^R(\cdot)$ with respect to c reduces to

$$2p^R(c)[\pi(m^R - c) + b^R - m^R].$$

Lastly, the derivative of c^S with respect to π is $-(b^S - m^S)/\pi^2$. Putting all this together, we have

$$\begin{aligned} \frac{dEU^R(c^S, \pi)}{d\pi} = & P^R(c^S)(m^R - t_0^R) + (1 - P^R(c^S))(t_1^R - m^R) \\ & - 2p^R(c)[\pi(m^R - c) + b^R - m^R] \frac{(b^S - m^S)}{\pi^2}. \end{aligned}$$

Let us consider the case where $m_*^S(b^R) \in [0, 1]$ such that we can set $b = b^R$, $m^R = m$ and $p^R(\cdot) = p(\cdot)$. In this case, only the last term of the above derivative is different in a CB and a CP model. The difference between $dEU^R(c^S, \pi)/d\pi$ of a CB model and $dEU^R(c^S, \pi)/d\pi$ of its equivalent CP model is then $2p(c^S)$ times

$$[\pi(m^R - c^S) + b - m^R] \frac{b - m^S}{\pi^2} - [\pi(m - c^S) + b^R - m] \frac{b^S - m}{\pi^2}. \quad (3.14)$$

Substituting $m^R = m$, $m^S = m_*^S(b)$, $c^S = (b^S - m)/\pi + m$ and $b = b^R$ in (3.14) we get

$$-\frac{(b^S - b^R)^2}{\pi(1 - \pi)} < 0. \quad (3.15)$$

It follows that for a given CB model and its equivalent CP model, R has a stronger preference for a more precise signal in the CP model. We summarize this in the following proposition.

Proposition 12. *Suppose $m_*^S(b^R) \in [0, 1]$ and let $b = b^R$, $m^R = m$ and $p^R(\cdot) = p(\cdot)$ such that the CB and CP models are equivalent. Then R has a greater utility gain from increasing π in the CP model than in the CB model.*

This proposition demonstrates the point made in the introduction of this section. Although we can, in many situations, write CB and CP models that are both Choice- and Strategy-Equivalent this does not imply that other properties are the same. Suppose that we were to introduce a first stage of the model where R can choose to pay to increase the competence of the Sender (we can think of this as a principal-agent situation where the principal pays to educate the agent). Then, even if we write a CB and a CP model to be exactly identical in the second stage, R 's first stage problem will be different in that he will be more inclined to pay to increase S 's competence in the CP model. Notice that this is different to the result of Che and Kartik [11] who notice that non-common priors and non-common preferences do not affect the Sender's incentive to acquire information in the same way. In our model the Sender's competence is exogenously fixed and we analyze the Receiver's gain from an increase in competence.

3.6 Conclusion

Non-common preferences and non-common priors represent two different ways of modeling player heterogeneity. This paper has analyzed conditions under which a model with common priors and non-common preferences has an equivalent model with non-common priors and common preferences. Our two measures of equivalence have been chosen to illustrate that equivalence at the individual level (Choice-Equivalence) is neither necessary nor sufficient for equivalence in the strategic game between the two players (Strategy-Equivalence). This implies that a CP and a CB model in which players act exactly the same individually may not have the players do so once they engage in information transmission. Whereas this may occur even if both models are CP or CB the reason

why Strategy-Equivalence fails when the models are of different types is particular to the nature of CP models: the players' priors affect not only their own preferred actions conditional on the signal but also their expectation of what the other player has seen. Therefore they may act the same individually but not together.

Often one of the two approaches is chosen over the other for the sake of convenience or because it fits well with a particular application of the model. But it is important to understand the consequences of this choice and we see this paper as a step in that direction. In particular, whenever the nature of a problem does not favor one of the two approaches over the other it might be an interesting exercise to construct two models that are Choice-Equivalent and see if they are also Strategy-Equivalent. As our results show, it is possible that such equivalent models are not available and in this case the choice of approach conditions the possible set of outcomes. And if equivalent models do exist, will they they also be equivalent if other elements are added to the game? Performing such a check will give an idea about to which extent results depend on the framework.

3.7 Proofs

Proof of Proposition 8. Define $Q_i^i(p_1^i) = \lambda_{1,t}^i / (p_1^i \lambda_{1,t}^i + (1-p_1^i) \lambda_{2,t}^i)$. Hence, we can rewrite the two conditions as $p_1^R Q_1^R(p_1^R) = p_1^S Q_1^S(p_1^S)$ and $p_1^R Q_2^R(p_1^R) = p_1^S Q_2^S(p_1^S)$. We can use the first condition to derive equilibrium p_1^R as a function of p_1^S .

$$\bar{p}_1^R(p_1^S) = \frac{\lambda_{2,1}^R p_1^S Q_1^S(p_1^S)}{\lambda_{1,1}^R (1 - p_1^S Q_1^S(p_1^S)) + \lambda_{2,1}^R p_1^S Q_1^S(p_1^S)}.$$

Substitute this into the second condition and cancel out the denominator of $\bar{p}_1^R(p_1^S)$ to get

$$\frac{\lambda_{1,1}^R \lambda_{2,1}^R p_1^S Q_1^S(p_1^S)}{\lambda_{1,1}^R \lambda_{2,1}^R p_1^S Q_1^S(p_1^S) + \lambda_{2,2}^R \lambda_{1,1}^R (1 - p_1^S Q_1^S(p_1^S))} = p_1^S Q_2^S(p_1^S).$$

Writing out the right-hand side and simplifying this we obtain

$$\frac{\lambda_{1,2}^R \lambda_{2,1}^R \lambda_{1,1}^S p_1^S}{\lambda_{1,2}^R \lambda_{2,1}^R \lambda_{1,1}^S p_1^S + \lambda_{1,1}^R \lambda_{2,2}^R \lambda_{2,1}^S (1 - p_1^S)} = \frac{\lambda_{1,2}^S p_1^S}{\lambda_{1,2}^S p_1^S + \lambda_{2,2}^S (1 - p_1^S)}$$

Rearranging we get the condition of the proposition. The second part is proven in the text. \square

Proof of Proposition 10. A sufficient condition for the result is that the rank of the matrix \mathbf{E} equals the row dimension. \mathbf{E} has dimension $(2(K-1) + K^2) \times KT$, and hence we require $T \geq K + 2$. However, we can reduce the problem by setting $u_{t,1} = 0$ for all $t = 1, \dots, T$. This implies that R always has expected utility 0 from choosing action 1, and we can choose the v vector such that if R chooses action k in partition l in the original problem, then $v_{k,l} > 0$ and otherwise $v_{k,l} < 0$. In this case we can delete the first T columns and the K rows corresponding to the conditions for $v_{1,l}$. Therefore, we reduce \mathbf{E} to a matrix $\hat{\mathbf{E}}$ of size $(2(K-1) + (K-1)^2) \times (K-1)T$ and the requirement becomes $T \geq K + 1$.

We can split the matrix $\hat{\mathbf{E}}$ into K blocks such that $\hat{\mathbf{E}} = (\hat{E}_1, \dots, \hat{E}_K)$, i.e. each block constitutes T columns of $\hat{\mathbf{E}}$. For instance,

$$\begin{aligned} \hat{E}_1 &= (-P_2^S, P_2^S, 0, \dots, 0, -P_1^R, P_2^R, 0, \dots, 0, \bar{P}_2^R, 0, \dots, 0, \bar{P}_3^R, \dots)' \text{ and} \\ \hat{E}_2 &= (0, -P_2^S, P_3^S, 0, \dots, 0, -P_2^R, P_3^R, 0, \dots, 0, \bar{P}_2^R, 0, \dots, 0, \bar{P}_3^R, \dots)'. \end{aligned}$$

Any vector in a block k is clearly independent of any combination of vectors that belong to other blocks, since some rows only have strictly positive entries in vectors in block k . Block $K - 1$ will only have non-zero entries in $2 + K - 1 = K + 1$ rows whereas all other blocks have non-zero entries in $4 + K - 1 = K + 3$ rows. What we need to check is that each block has rank at least $K + 1$. This would imply that \mathbf{E} has rank at least $K(K + 1)$, which is sufficient for the result. We can reduce this to the condition that the matrices R_k defined in the main text have rank at least $K + 1$. Assumption 2 assures this. \square

Proof of Lemma 7. If the two models are Choice-Equivalent, then obviously for fixed $\{y_n\}_{n=1}^N$, the conditions $U^S(y_n, x_n) = U^S(y_n, x_{n-1})$ remain unchanged. Similarly, fixing $\{x_n\}_{n=0}^N$ we have

$$\begin{aligned} y_n^B &= \arg \max_y \int_{x_{n-1}}^{x_n} U^R(y, x) l_B^R(x) dx \\ &= \arg \max_y \int_{x_{n-1}}^{x_n} U^R(y, x) k_n l_B^R(x) dx \\ &= \arg \max_y \int_{x_{n-1}}^{x_n} U^R(y, x) l_A^R(x) dx \\ &= y_n^A. \end{aligned}$$

This gives the second part of the lemma. \square

Proof of Proposition 11. Given the assumptions, then for all x

$$\int_{\Theta} a(y, t) \frac{p(t)\lambda(t|x)}{\int_{\Theta} p(t')\lambda(t'|x) dt'} dt = \int_T a(y, t) \frac{p(t)\lambda(t|x)}{\int_{\Theta} p(t')\lambda(t'|x) dt'} dt = 0.$$

The first equality holds since $a(y, \theta) = 0$ for any $\theta \notin T$. The last equality holds because of (PCLR), which implies that if the integral is zero for some x it is zero for all x . Hence, if we set $p^R(\cdot) = p(\cdot)$ then (SE2) is satisfied. Then by Lemma 7, equation (3.11) defines a set of $U^S(\cdot)$ for which we can construct a Strategy-Equivalent CP model. \square

Bibliography

- [1] A.R. Admati and P. Pfleiderer. Broadcasting opinions with an overconfident sender*. *International Economic Review*, 45(2):467–498, 2004.
- [2] D. Austen-Smith. Strategic transmission of costly information. *Econometrica: Journal of the Econometric Society*, pages 955–963, 1994.
- [3] M. Balboa, J.C. Gómez-Sala, and G. López-Espinosa. The value of adjusting the bias in recommendations: International evidence. *European Financial Management*, 15(1):208–230, 2009.
- [4] D.P. Baron. Persistent media bias. *Journal of Public Economics*, 90(1-2):1–36, 2006.
- [5] A. Blume, O. Board, and K. Kawamura. Noisy talk. *Theoretical Economics*, 2(4):395–440, 2007.
- [6] J. Capstaff, K. Paudyal, and W. Rees. A comparative analysis of earnings forecasts in europe. *Journal of Business Finance & Accounting*, 28(5-6):531–562, 2001.
- [7] B. Çelen, S. Kariv, and A. Schotter. Words Speak Louder than Actions and Improve Welfare: An Experimental Test of Advice and Social Learning. *Levine’s Bibliography*, 2005.
- [8] Pew Research Center. 2006 pew research center for the people & the press news consumption and believability study.

- [9] G. Charness, E. Karni, and D. Levin. Individual and group decision making under risk: An experimental study of Bayesian updating and violations of first-order stochastic dominance. *Journal of Risk and uncertainty*, 35(2):129–148, 2007.
- [10] G. Charness and D. Levin. When optimal choices feel wrong: A laboratory study of Bayesian updating, complexity, and affect. *American Economic Review*, 95(4):1300–1309, 2005.
- [11] Y.K. Che and N. Kartik. Opinions as incentives. *Journal of Political Economy*, 117(5):815–860, 2009.
- [12] V. Crawford. A survey of experiments on communication via cheap talk. *Journal of Economic Theory*, 78(2):286–298, 1998.
- [13] V.P. Crawford and J. Sobel. Strategic information transmission. *Econometrica: Journal of the Econometric Society*, 50(6):1431–1451, 1982.
- [14] A. Cukierman and M. Tommasi. When does it take a nixon to go to china? *American Economic Review*, pages 180–197, 1998.
- [15] W.F.M. De Bondt et al. Herding in analyst earnings forecasts: evidence from the united kingdom. *European Financial Management*, 5(2):143–163, 1999.
- [16] S. DellaVigna and E. Kaplan. The fox news effect: Media bias and voting*. *The Quarterly Journal of Economics*, 122(3):1187–1234, 2007.
- [17] A. Dische. Dispersion in analyst forecasts and the profitability of earnings momentum strategies. *European Financial Management*, 8(2):211–228, 2002.
- [18] M.A. El-Gamal and D.M. Grether. Are People Bayesian? Uncovering Behavioral Strategies. *Journal of the American Statistical Association*, 90(432), 1995.
- [19] J. Farrell and R. Gibbons. Cheap talk with two audiences. *The American Economic Review*, 79(5):1214–1223, 1989.
- [20] E. Gal-Or. Information sharing in oligopoly. *Econometrica: Journal of the Econometric Society*, 53(2):329–343, 1985.

- [21] E. Gal-Or. Information transmission—Cournot and Bertrand equilibria. *The Review of Economic Studies*, 53(1):85–92, 1986.
- [22] M. Gentzkow and J.M. Shapiro. Media Bias and Reputation. *Journal of Political Economy*, 114(2), 2006.
- [23] B. Goldberg. *Bias: A CBS Insider Exposes How the Media Distort the News*. Regnery Publishing, 2001.
- [24] J.R. Green and N.L. Stokey. A two-person game of information transmission. *Journal of Economic Theory*, 135(1):90–104, 2007.
- [25] D.M. Grether. Bayes rule as a descriptive model: The representativeness heuristic. *The Quarterly Journal of Economics*, 95(3):537–557, 1980.
- [26] D. Griffin and A. Tversky. The weighing of evidence and the determinants of confidence* 1. *Cognitive Psychology*, 24(3):411–435, 1992.
- [27] T. Groseclose and J. Milyo. A measure of media bias*. *The Quarterly Journal of Economics*, 120(4):1191–1237, 2005.
- [28] Steffen Huck and Georg Weizsäcker. Do players correctly estimate what others do? : Evidence of conservatism in beliefs. *Journal of Economic Behavior & Organization*, 47(1):71 – 85, 2002.
- [29] K. Kawamura. Confidence and competence in communication. 2010.
- [30] A. Kydd. Which side are you on? bias, credibility, and mediation. *American Journal of Political Science*, 47(4):597–611, 2003.
- [31] J.R. Lott Jr and K.A. Hassett. Is newspaper coverage of economic events politically biased? *American Enterprise Institute Mimeo*, 2004.
- [32] P. Milgrom and J. Roberts. Relying on the information of interested parties. *The RAND Journal of Economics*, pages 18–32, 1986.

- [33] J. Morgan and P.C. Stocken. An analysis of stock recommendations. *The RAND Journal of Economics*, 34(1):183–203, 2003.
- [34] S. Mullainathan and A. Shleifer. The market for news. *The American Economic Review*, 95(4):1031–1053, 2005.
- [35] Y. Nyarko, A. Schotter, and B. Sopher. On the informational content of advice: A theoretical and experimental study. *Economic Theory*, 29(2):433–452, 2006.
- [36] M. Ottaviani and P.N. Sørensen. Reputational cheap talk. *The Rand journal of economics*, 37(1):155–175, 2006.
- [37] Ganna Pogrebna. Naïve advice when half a million is at stake. *Economics Letters*, 98(2):148 – 154, 2008.
- [38] J. Rudiger. The market for extreme news. *Mimeo.*, 2012.
- [39] Jesper Rudiger. Information Transmission with Non-Common Priors. *SSRN eLibrary*, 2011.
- [40] A. Schotter and B. Sopher. Trust and trustworthiness in games: An experimental study of intergenerational advice. *Experimental Economics*, 9(2):123–145, 2006.
- [41] A. Schotter and B. Sopher. Advice and behavior in intergenerational ultimatum games: An experimental approach. *Games and Economic Behavior*, 58(2):365–393, 2007.
- [42] A. Schotter, B. Sopher, and CV Starr Center for Applied Economics (New York University). Social learning and coordination conventions in intergenerational games: An experimental study. *Journal of Political Economy*, 111(3):498–529, 2003.
- [43] D. J. Seidmann and E. Winter. Strategic information transmission with verifiable messages. *Econometrica*, 65(1):pp. 163–169, 1997.
- [44] J.C. Stein. Cheap talk and the fed: A theory of imprecise policy announcements. *The American Economic Review*, pages 32–42, 1989.

- [45] W. Suen. The self-perpetuation of biased beliefs*. *The Economic Journal*, 114(495):377–396, 2004.
- [46] A. Tversky and D. Kahneman. Advances in prospect theory: Cumulative representation of uncertainty. *Journal of Risk and uncertainty*, 5(4):297–323, 1992.
- [47] J. Watson. Information transmission when the informed party is confused. *Games and Economic Behavior*, 12:143–161, 1996.
- [48] A. Wolinsky. Information transmission when the sender’s preferences are uncertain. *Games and Economic Behavior*, 42(2):319–326, 2003.