

Preker, Jurek; Karos, Dominik

Working Paper

Strategic information selection

Center for Mathematical Economics Working Papers, No. 689

Provided in Cooperation with:

Center for Mathematical Economics (IMW), Bielefeld University

Suggested Citation: Preker, Jurek; Karos, Dominik (2024) : Strategic information selection, Center for Mathematical Economics Working Papers, No. 689, Bielefeld University, Center for Mathematical Economics (IMW), Bielefeld,
<https://nbn-resolving.de/urn:nbn:de:0070-pub-29883821>

This Version is available at:

<https://hdl.handle.net/10419/289846>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.




<https://creativecommons.org/licenses/by/4.0/>

March 2024

Strategic Information Selection

Jurek Preker and Dominik Karos



Strategic Information Selection*

Jurek Preker[†]

Dominik Karos[‡]

March 28, 2024

Abstract

Before choosing her action to match the state of the world, an agent observes a stream of messages generated by some unknown binary signal. The agent can either learn the underlying signal for free and update her belief accordingly or ignore the observed message and keep her prior belief. After each period the stream stops with positive probability and the final choice is made.

We show that a Markovian agent with Gilboa-Schmeidler preferences learns and updates after confirming messages, but she ignores contradicting messages if her belief is sufficiently strong. Her threshold solely depends on the least precise signal.

The agent has strictly higher anticipatory utility than an agent who uses every message to update. However, the latter has a higher chance to choose the correct outcome in the end. In a population of strategic agents, who only differ in their initial beliefs, polarization is inevitable.

Keywords: Dynamic Decision Problem, Ambiguity, Gilboa-Schmeidler Preferences, Confirmation Bias, Polarization

JEL Classification: D81, D83, D91

*The authors thank Chris Woolnough, Gerrit Bauch, Giorgio Ferrari, Ehud Lehrer, Fynn Louis Närmann, and Frank Riedel as well as many seminar participants at the universities of Bielefeld, Durham, and Tel Aviv for their comments and suggestions.

[†]Center for Mathematical Economics, Bielefeld University, Postfach 100131, 33501 Bielefeld, Germany. jurek.preker@uni-bielefeld.de, corresponding author

[‡]Center for Mathematical Economics, Bielefeld University, Postfach 100131, 33501 Bielefeld, Germany. dominik.karos@uni-bielefeld.de

1 Introduction

In an average minute of December 2023, social media users sent out 360,000 tweets on \mathbb{X} , sent 694,000 reels on Instagram via direct messages, and liked four million Facebook posts.¹ A user of any of these social media platforms cannot read all the articles whose headlines appear in her feed, but has to choose what headlines to follow up on and what to ignore. This makes information acquirement on social media fundamentally different from using print media. While a printed newspaper, even if read from cover to cover, will provide the reader only with a fixed number of articles, incoming news in social media form an endless stream. Thus, the information acquirement process in social media requires a pre-step: before reading and processing an article, the user needs to select what to read from a huge supply.

It has been well established that this causes people to learn differently: for instance, they learn less from online than from print sources (Eveland and Dunwoody, 2002; Yang and Grabe, 2011); but at the same time they are more focused on specific topics selected by themselves rather than the news editors (Kruikemeier, Lecheler, and Boyer, 2017). Empirical studies suggest the main drivers behind information selection be *attitude consistency* and *source credibility*. Sülflow, Schäfer, and Winter (2019) investigated the impact of both and showed that attitude consistency has no impact on the time people spend looking at a headline in their news feed but increases the probability they follow up on it and actually read the underlying article. Moreover, they found that the assumed source credibility has a positive effect on both the time spent with a headline as well as on the follow-up probability.²

As social media have become increasingly important as a source of news,³ understanding the mechanics behind information selection as well as its individual and social consequences, such as polarization⁴ and collective decision making, are crucial. This arti-

¹See Statista (2024).

²See also Chiang and Knight (2011) for the importance of credibility.

³The share of people in Germany who weekly use social media as a source of news has increased from 18 % in 2013 to 29 % in 2023, having overtaken classical print media in 2019. In the US, this number has risen from 27 % to 48 % in the same time. The effect is particularly strong in the group of young adults (Reuters Institute for the Study of Journalism, 2023). According to a study by the Pew Research Center, a majority across 19 countries agrees that social media is an effective way to raise public awareness about political or social issues (77 % agree), to change people's minds (65 %), get elected officials to pay attention to issues (64 %), and to influence policy decisions in their own country (61 %) (Pew Research Center, 2022).

⁴In the above mentioned report by the Pew Research Center, 65 % of the investigated subjects say

cle provides a dynamic model of strategic information selection that rationalizes observed behavior and investigates its implications.

We consider a decision maker who has to match an unknown binary state of the world. Before making her choice, she receives a discrete stream of binary and symmetric messages, to be thought of as articles on social media that endorse one of the two states, that are correct with some unknown probability of at least 50%. She has the option to learn this probability; in the news media example, after obtaining a headline (the message) she may click on the link, read the article, learn how well the article is researched, and infer the probability that the message of the headline was correct. If she chooses to ignore the message, her belief does not change; if she chooses to learn the underlying signal, she has to perform a Bayesian update. The underlying assumption is that people can easily ignore a headline they only see briefly, but once they have engaged with it, it cannot be unseen and will affect their belief. For her belief update she uses her previous belief, the message, and the signal. That is, we assume that the decision maker is Markovian, meaning that she does not use the entire history of previous messages and signals.⁵

In each period the stream of messages stops with some exogenous probability and the decision maker has to make her choice about the state; otherwise a new signal and a new message will be drawn. Maximizing her utility thus becomes a dynamic optimization problem, where with every decision she has to take into account how her belief will develop in the future. The Markovian structure gives rise to a Bellman equation whose solution is the decision maker's *anticipatory utility function*.⁶ We show that optimal strategies exist and describe their general structure. Afterwards, we focus on a decision maker with max-min preferences (cf. Gilboa and Schmeidler, 1989) and show that any optimal strategy requires her to ignore messages that oppose her belief whenever her belief is sufficiently strong. Hence, she exhibits what Stone and Wood (2018) call “non-ego-based

that access to the internet and social media has made people more divided in their political opinions. Allcott, Braghieri, Eichmeyer, and Gentzkow (2020) showed that regular Facebook users who turned off the network for a month were less polarized in their attitudes afterwards. See Prior (2013) and Kubin and von Sikorski (2021) for political reviews on how (social) media drive polarization.

⁵There are two major reasons for this assumption: On the one hand, people see and read so many headlines and articles on social media that eventually, they cannot keep track of the entire history of their messages, signals, and beliefs. On the other hand, it allows us to condition the agent's behavior only on her belief and the signal, independent of time and signal history, so that the analysis remains tractable.

⁶We call this utility *anticipatory* as it takes into account all possible further paths on which beliefs might change. This relates to the concept of *utility from anticipation*, introduced by Loewenstein (1987) and later investigated, for instance, by Caplin and Leahy (2001) and Kőszegi (2017).

cognitive dissonance”⁷: media consumers avoid contradicting information “not because this threatens their egos, but simply because uncertainty is unsettling”.

Ignoring free information is obviously not optimal in the framework of Blackwell (1951, 1953). However, we show that it is optimal in our setting, rationalizing human behavior in social media.⁸ In particular, when comparing our strategic decision maker with a Bayesian decision maker, who updates his belief after every round, we find that she obtains a strictly higher anticipatory utility, while her probability to match the true state of the world is strictly lower.

Surely, this model describes social media only in a very simplified way. For instance, we assume that the observed messages are independent of the user’s belief, and we ignore algorithms, echo chambers, as well as the targeted use of fake news by political agents. Yet, even without all these or any other form of homophily in social media networks⁹ we obtain quite extreme polarization results where society only becomes more divided over time.

The rest of the paper is organized as follows. In Section 2, we present further related literature. Section 3 develops the model, discusses its assumptions, and derives general existence and uniqueness results. In Section 4, we derive optimal strategies in the absence of ambiguity, while Section 5 covers agents with Gilboa-Schmeidler preferences under ambiguity. In Section 6, we compare our strategic decision maker with one who updates after every message he observes. Social consequences in terms of polarization and collective decision making are covered in Section 7. Section 8 provides several ways to extend the model and Section 9 concludes the paper. Proofs are relegated to the appendix.

2 Further Related Literature

A large class of models in which agents do not fully exploit the available information focuses on limited cognitive abilities. For instance, the observed tendency to favor and search for information that reinforces prior beliefs while ignoring opposing information

⁷The psychological concept of *cognitive dissonance* has been popularized by Festinger (1957): Uncertainty about what to believe and facing a trade-off between feeling validated and wanting to know the truth makes people feel uncomfortable.

⁸Social media are not the only place where information might be ignored: even some reader of this paper might reject our model and put it aside without letting it affect their belief about human behavior.

⁹See Aiello et al. (2012) for a statistical analysis of homophily in social media.

is known as *selective exposure*.¹⁰ Lord, Ross, and Lepper (1979) demonstrate that when people with strong opinions on controversial issues—in that case, the death penalty in the USA—are given equally convincing evidence for both positions, they become even stronger in their opinions. So, different pieces of information favoring contrary states of the world do not neutralize, but people pick what they want to hear, interpret it as confirming evidence, and ignore the rest. This *confirmation bias*¹¹ is not what drives our results, as our decision maker correctly processes all information that she uses, but deliberately chooses to ignore some of it. This distinction relates to the psychological distinction between automatic and controlled processes (Schneider and Shiffrin, 1977).

Among the models of strategic information selection, an important strand of the literature, following Sims (2003), focuses on *rational inattention* where agents only obtain partial information due to cost. (See Maćkowiak, Matějka, and Wiederholt (2023) for a review.) In contrast, information in our model comes at no cost and may still be rejected.

In many models agents choose ex ante how to gather information, that is, they only obtain the actual message after their decision. Machina (1989) gives a stylized example of such behavior based on the Allais Paradox, where non-expected utility maximizers choose not to obtain information. Carrillo and Mariotti (2000) present a decision maker with time-inconsistent preferences who decides to acquire only incomplete information about a future externality. In Suen (2004), agents observe coarsened signals and decide themselves on a coarsening rule. Similarly, decision makers in Che and Mierendorff (2019) have access to different biased sources and choose the optimal one.

In contrast, our decision maker first obtains a message and then chooses what to do with it. This brings our model closer to Chen (2022), who considers agents with different models at hand who update their beliefs according to the one that best supports their bias, and to Fryer, Harms, and Jackson (2019) where agents deal with “ambiguous” signals that can be used in different ways.¹² Yet, our model is distinct in several key features. First, agents can decide to entirely discard and ignore a message. This brings us closer to Allahverdyan and Galstyan (2014), where an agent does not react to advice (sent by another agent who wants to persuade the former one) that conflicts with her prior conviction. Second, a decision maker with Gilboa-Schmeidler preferences will not

¹⁰Hart et al. (2009) provide a psychological meta-study on that subject.

¹¹See Nickerson (1998) for a psychological overview as well as Rabin and Schrag (1999).

¹²Note that their use of “ambiguous” differs from our terminology as we refer to ambiguity as missing knowledge about probabilities.

operate under the assumption that the message has been generated by the signal that fits best, but rather by the signal that provides her with the lowest utility. Third, even if there is no ambiguity at all, that is, if the signal is known, some of the messages will be ignored.

The concept of *motivated beliefs* might serve as a bridge between our model and the literature on ex-ante information design. According to Epley and Gilovich (2016), agents do not directly choose their beliefs, but “their motivations guide what information they consider, resulting in favorable conclusions that seem mandated by the available evidence”. Even though our agent acts ex post, she follows precisely this approach: if she expects a more favorable conclusion from learning the signal and interpreting the message, she does so; otherwise she discards the message entirely.

The *good news-bad news effect* or *optimism bias*—experimentally confirmed by Eil and Rao (2011) and axiomatically characterized by Bracha and Brown (2012)—prescribes people to respect the strength of a signal when the news is favorable. In Brunnermeier and Parker (2005), agents choose subjective probabilities such that their current well-being is maximized, at the cost of possibly taking a wrong decision afterwards. This effect differs from our model as our states are neutral. However, agents in our setup are happier, the more convinced they are of any state; being unsure about the true state of the world makes them uncomfortable.

Signal ambiguity as a reason for information selection has been described by Gentzkow and Shapiro (2006). The main argument is that a decision maker who does not like what she reads can conveniently tell herself that the article must be poorly researched and can be safely ignored. Ambiguity as a driver of polarization has been modeled by Glaeser and Sunstein (2013). In their model, people believe different parts of the information they get. In Baliga, Hanany, and Klibanoff (2013), polarization occurs as a result of ambiguity aversion. This is not the case here: even though we consider extremely ambiguity averse decision makers, our results about polarization would remain true even if there were no ambiguity at all.

3 Optimal Information Selection

An agent has to match the correct *state of the world* from the state space $\Omega = \{0, 1\}$. That is, the *choice space* is $C = \Omega$ and her von Neumann-Morgenstern utility function

$v: C \times \Omega \rightarrow \{0, 1\}$ is given by $v(c, \omega) = 1$ if $c = \omega$ and $v(c, \omega) = 0$ otherwise.

To keep notation simple, we identify any belief $\Lambda \in \Delta(\Omega)$ she might entertain with $\lambda = \Lambda(1) \in [0, 1]$. With such a belief her expected utilities from choosing 0 or 1 are

$$\begin{aligned}\mathbb{E}[v(0, \cdot)] &= 1 \cdot (1 - \lambda) + 0 \cdot \lambda = 1 - \lambda, \text{ and} \\ \mathbb{E}[v(1, \cdot)] &= 1 \cdot \lambda + 0 \cdot (1 - \lambda) = \lambda,\end{aligned}$$

respectively. Thus, her indirect expected utility at belief λ is $u(\lambda) := \max\{\lambda, 1 - \lambda\}$.

Prior to her choice the agent obtains information, which is modeled as a discrete stream of *symmetric binary signals*, that is, maps $\zeta: \Omega \rightarrow \Delta(\Omega)$ with $\zeta(1|1) = \zeta(0|0)$, and realized *messages* $m \in \{0, 1\}$. We assume without loss of generality that $\zeta(1|1) \geq \frac{1}{2}$, so that any signal ζ can be identified with $z = \zeta(1|1) \in [\frac{1}{2}, 1]$.¹³

Let $Z \subseteq [\frac{1}{2}, 1]$ be a compact set of signals. At each period $t \in \mathbb{N}$, the signal stream ends with some (exogenous) probability $1 - p$ and the agent has to make her decision. With probability p , the stream continues. In this case a signal $z_t \in Z$ is generated, in a potentially unknown way, and afterwards a message $m_t \in \{0, 1\}$ realizes according to z_t , that is, $\mathbb{P}_{z_t}(m_t = 0|\omega = 0) = \mathbb{P}_{z_t}(m_t = 1|\omega = 1) = z_t$.¹⁴ The decision maker only observes m_t .

For any signal z the expected probabilities that 0 and 1 realize are

$$q_z^0(\lambda) := \lambda(1 - z) + (1 - \lambda)z \quad \text{and} \quad q_z^1(\lambda) := \lambda z + (1 - \lambda)(1 - z),$$

respectively. Define for each $z \in [\frac{1}{2}, 1]$ two functions $g_z, g_z^{-1}: [0, 1] \rightarrow [0, 1]$ by

$$g_z^{-1}(\lambda) := \frac{\lambda(1 - z)}{\lambda(1 - z) + (1 - \lambda)z} \quad \text{and} \quad g_z(\lambda) := \frac{\lambda z}{\lambda z + (1 - \lambda)(1 - z)} \quad (1)$$

for $z \in [\frac{1}{2}, 1)$, as well as $g_1^{-1} \equiv 0$ and $g_1 \equiv 1$.¹⁵ Surely, if the agent knew z_t and m_t , then her posterior beliefs in case of a message of 0 or 1 were $\lambda_t = g_{z_t}^{-1}(\lambda_{t-1})$ or $\lambda_t = g_{z_t}(\lambda_{t-1})$,

¹³The triple $(\Omega, \zeta(\cdot|0), \zeta(\cdot|1))$ is a symmetric binary Blackwell experiment, see Blackwell (1951, 1953).

¹⁴The signals $(z_t)_{t \geq 1}$ can be correlated. Yet, given the signals, the probability measures $(\mathbb{P}_{z_t})_{t \geq 1}$ are independent.

¹⁵The latter definitions coincide with (1) for $\lambda \in [0, 1)$ or $\lambda \in (0, 1]$, respectively. The values of $g_1^{-1}(1)$ and $g_1(0)$ are irrelevant for the agent's optimal decision: if she is certain that the true state is 0 (or 1, respectively) and receives a message stating the opposite, she will consider it impossible that the signal is perfectly accurate.

respectively, by Bayes' rule. For later reference, we collect some basic properties of g_z and g_z^{-1} in the following lemma. The proof is straightforward and therefore omitted.

Lemma 3.1. *For any $z \in [\frac{1}{2}, 1)$, the functions g_z^{-1} and g_z are inverse to each other, and strictly increasing in $\lambda \in [0, 1]$. Moreover, for any $\lambda \in (0, 1)$, the map $g_z^{-1}(\lambda)$ is strictly decreasing, while $g_z(\lambda)$ is strictly increasing in $z \in [\frac{1}{2}, 1]$. For all $z \in (\frac{1}{2}, 1]$ and $\lambda \in (0, 1)$, it holds that $g_z^{-1}(\lambda) < \lambda < g_z(\lambda)$, and $g_z(1 - \lambda) = 1 - g_z^{-1}(\lambda)$.*

3.1 The Stage Decisions

If some message m_t is drawn, the agent has two options: she can either discard m_t entirely and stick to her prior λ_{t-1} , or she can (without cost) learn z_t and update her belief. Naturally, she cannot update her belief without knowing z_t . Additionally, we assume she cannot discard m_t once she has learned z_t , reflecting the idea that information cannot be “unseen” once it has been engaged with. Formally, after having received a message she has to choose one of two *actions*: she can *ignore* (i) the message or she can *learn* (l) signal z_t , and perform a Bayesian belief update as above.¹⁶

We assume that the decision maker is Markovian. That is, her action in any period only depends on message m_t and her belief λ_{t-1} . Thus, her set of (pure) Markovian strategies is $\mathcal{S} := (\{i, l\}^2)^{[0,1]}$, where $\lambda \mapsto \sigma(\lambda) = (\sigma_0(\lambda), \sigma_1(\lambda))$ means that at belief λ , she plays $\sigma_0(\lambda) \in \{i, l\}$ if $m = 0$, and $\sigma_1(\lambda) \in \{i, l\}$ if $m = 1$. For instance, the benchmark case of a *Bayesian decision maker*, who uses every message, is described by the strategy ρ with $\rho(\lambda) = (l, l)$ for all $\lambda \in [0, 1]$.

Any strategy σ together with an initial belief λ and a finite stream of signals $(z_t)_{t=1}^T$ and messages $(m_t)_{t=1}^T$ generates a *history*, i.e., a finite sequence $h = h^T(\sigma) = (z_t, m_t, \lambda_t)_{t=1}^T$, where λ_t is the decision maker's belief at the end of period t , that is, $\lambda_0 = \lambda$ and for $t = 1, \dots, T$

$$\lambda_t = \begin{cases} g_{z_t}^{-1}(\lambda_{t-1}) & \text{if } m_t = 0 \text{ and } \sigma_0(\lambda_{t-1}) = l, \\ g_{z_t}(\lambda_{t-1}) & \text{if } m_t = 1 \text{ and } \sigma_1(\lambda_{t-1}) = l, \\ \lambda_{t-1} & \text{else.} \end{cases}$$

¹⁶Throughout, we distinguish between *actions* in each period and the *choice* that must be made once after the stream of signals has ended.

For any strategy σ , set of signals $Z \subseteq [\frac{1}{2}, 1]$, and belief $\lambda \in [0, 1]$ define

$$B_{Z,i}^0(\lambda) = B_{Z,i}^1(\lambda) := \{\lambda\}, \quad B_{Z,l}^0(\lambda) := \{g_z^{-1}(\lambda) : z \in Z\}, \quad B_{Z,l}^1(\lambda) := \{g_z(\lambda) : z \in Z\}.$$

These are the sets of possible posterior beliefs if the decision maker ignores her message or learns the underlying signal upon observing message 0 or 1, respectively.

We assume that at the end of the stream the decision maker chooses optimally for $\lambda_T \neq \frac{1}{2}$ and chooses either option with equal probability if $\lambda_T = \frac{1}{2}$.

3.2 Aggregators

The decision maker chooses a Markovian strategy such that for each belief $\lambda \in [0, 1]$, her expected utility is maximized. This expected utility in turn depends on the utility at potential future beliefs: when a decision maker holds belief λ and receives a message, she can learn z and perform a belief update to $g_z(\lambda)$ or $g_z^{-1}(\lambda)$. Thus, her utility at λ depends on her utility at $g_z(\lambda)$ and $g_z^{-1}(\lambda)$ for all possible signals $z \in Z$, as well as her attitude towards the ambiguity about the signal.

Definition 3.2. Let $D \subseteq \mathbb{R}^n$ be a set, let $\mathcal{F}(D)$ be the set of functions from D to \mathbb{R} , and let $\mathcal{P}(D)$ be the power set of D . An *aggregator* is a map $H: \mathcal{F}(D) \times \mathcal{P}(D) \rightarrow \mathbb{R}$ that satisfies the following two conditions:

Consistency. For all $x \in D$, $H(U, \{x\}) = U(x)$.

Monotonicity. For all $B \subseteq D$ and $U, V: D \rightarrow \mathbb{R}$ with $U|_B \geq V|_B$ it holds that $H(U, B) \geq H(V, B)$. If $U|_B \geq V|_B + \varepsilon$ for some $\varepsilon > 0$, then $H(U, B) > H(V, B)$. \square

An aggregator describes how an agent with payoff function U evaluates sets of objects $B \subseteq D$ without knowing what element of B is the true one. Consistency requires that if the agent knows that the set of potential objects only contains x , then she only evaluates x . Monotonicity requires that if each option in B is evaluated higher according to U than according to V , then this also holds true under aggregation.

In later sections we will have a closer look on the *Gilboa-Schmeidler* aggregator. For any set D of beliefs, expected utility function $U: D \rightarrow \mathbb{R}$, and $B \subseteq D$ it is defined by

$$H[U, B] = \inf_{\lambda \in B} U(\lambda). \tag{2}$$

One can easily see that H is both consistent and monotonic.

3.3 Anticipatory Utility

As the decision maker faces a repeated decision problem, her utility at any belief λ depends on her utility at potential future beliefs. Since she is Markovian, this gives rise to a functional equation. In order to keep notation simple, we use $\mathcal{F} = \mathcal{F}([0, 1])$ and $\mathcal{P} = \mathcal{P}([0, 1])$.

Definition 3.3. Let $Z \subseteq [\frac{1}{2}, 1]$ be a compact set of signals, and let $H: \mathcal{F} \times \mathcal{P} \rightarrow \mathbb{R}$ be an aggregator. A function $U \in \mathcal{F}$ with

$$U(\lambda) = (1 - p)u(\lambda) + pH [q^0(\lambda) H [U, B_{Z, \sigma_0(\lambda)}^0(\lambda)] + q^1(\lambda) H [U, B_{Z, \sigma_1(\lambda)}^1(\lambda)], Z] \quad (3)$$

for all $\lambda \in [0, 1]$ is called *anticipatory utility function* (given H , Z , σ , and p). \square

The second term on the right-hand side of (3) represents the continuation utility if a message is received. Two sources of uncertainty are aggregated here: first, ex ante, about the message that will be received (the distribution of messages depends on the unknown z), and second, at the interim stage, about z . The latter arises only if she chooses action l after her message and Z is not a singleton; otherwise, the set $B_{Z,a}^m(\lambda)$ collapses to the singleton $\{\lambda\}$.

Example 3.4. Let $Z \subseteq [\frac{1}{2}, 1]$ be compact, let H be the Gilboa-Schmeidler aggregator in (2), and consider the strategy ρ with $\rho(\lambda) = (l, l)$ for all $\lambda \in [0, 1]$. In this case any anticipatory utility function U must satisfy

$$U(\lambda) = (1 - p)u(\lambda) + p \min_{z \in Z} \left\{ q_z^0(\lambda) \min_{z' \in Z} \{U(g_{z'}^{-1}(\lambda))\} + q_z^1(\lambda) \min_{z'' \in Z} \{U(g_{z''}(\lambda))\} \right\}.$$

Altogether the decision maker uses three aggregators, and they are used independently of one another.¹⁷ This reflects the idea that a decision maker might have a change of heart about what constitutes the worst case signal after the message has been observed, and that she is aware of this. We will consider the alternative approach where she remains unaware of it in Subsection 5.2. \square

¹⁷A decision maker might use three different aggregators. We shall not pursue this path here as it will add nothing to the general discussion in the current section, and the subsequent sections will focus on Gilboa-Schmeidler decision makers exclusively.

The anticipatory utility function in (3) depends on the decision maker’s strategy σ . We first show that under some mild conditions on the aggregator H there is a unique anticipatory utility function for each $\sigma \in \mathcal{S}$.

Proposition 3.5. *Let $Z \subseteq [\frac{1}{2}, 1]$ be a compact set of signals and let $H: \mathcal{F} \times \mathcal{P} \rightarrow \mathbb{R}$ be an aggregator. If H satisfies*

$$|H[U, B] - H[V, B]| \leq \|U - V\|_\infty \tag{4}$$

for all $B \subseteq D$, then for each $\sigma \in \mathcal{S}$ there is a unique anticipatory utility function, denoted by U_Z^σ .¹⁸

Throughout the paper we restrict ourselves to aggregators that satisfy (4). Before we derive any further results, some comments are in order: With our symmetry assumption we exclude any idiosyncratic preferences over the states that the agent might have. Information selection is, hence, not driven by ego-based cognitive dissonance. Neither do we assume that information is costly or limited. The agent can learn as much as she wants, but we give her the ability to discard messages if she wishes to.¹⁹

Moreover, the anticipatory utility function describes her evaluation of the problem at every belief. This evaluation takes place before the state of the world is revealed: once the stream of messages stops, utilities realize according to u , that is, at the time of utility realization, the state remains unknown. Thus, an agent who has a very firm but incorrect belief, might still be very happy at the time of her decision and only learn much later about its dire consequences.²⁰

3.4 The Agent’s Problem

As each (Markovian) strategy σ leads to some anticipatory utility function U_Z^σ , the decision maker’s optimal utility at prior λ is

$$U_Z^*(\lambda) = \sup_{\sigma \in \mathcal{S}} U_Z^\sigma(\lambda). \tag{5}$$

¹⁸Surely, U_Z^σ depends on H and p as well. As this will not play any role in the further analysis, we drop H and p from the notation.

¹⁹Any user of social media will agree that they ignore a large extent of “news” there and will not let them affect their worldview in any way.

²⁰In this context the \vee -shape of the utility function u reflects the agent’s cognitive dissonance (Festinger, 1957).

The map U_Z^* is well-defined for any aggregator that satisfies Equation (4). If there is a universal strategy σ^* such that $U_Z^*(\lambda) = U_Z^{\sigma^*}(\lambda)$ for all λ , we call this strategy *optimal* (given H , Z , and p). Our first result is that an optimal strategy σ^* exists.

Proposition 3.6. *Let $Z \subseteq [\frac{1}{2}, 1]$ be a compact set of signals, let $H: \mathcal{F} \times \mathcal{P} \rightarrow \mathbb{R}$ be an aggregator that satisfies (4), and let U_Z^* be defined as in (5). Then $U_Z^* = U_Z^{\sigma^*}$ if and only if σ^* satisfies*

$$\sigma_m^*(\lambda) = \begin{cases} l & \text{if } H[U_Z^*, B_{Z,l}^m(\lambda)] > U_Z^*(\lambda) \\ i & \text{if } H[U_Z^*, B_{Z,l}^m(\lambda)] < U_Z^*(\lambda) \end{cases} \quad \text{for all } \lambda \in [0, 1] \text{ and } m \in \{0, 1\}. \quad (6)$$

Before an agent with belief λ who observes message m learns the true z , she faces a set of possible posteriors, $B_{Z,l}^m(\lambda)$. A strategy satisfying (6) prescribes to follow up on m if these posteriors aggregate to a strictly higher utility than $U_Z^*(\lambda)$, and to ignore the message if they aggregate to a strictly lower utility than $U_Z^*(\lambda)$. Thus, a decision maker who solves (5) might not use all messages she receives. In the following two sections we shall find the set of beliefs where the decision maker ignores either message.

4 Optimal Strategies for ex-ante Known Signals

Following the literature on ambiguous signals, one might be inclined to attribute the decision maker's strategic ignorance to the ambiguity she faces. But this is not the case. In this section we shall analyze the situation of a decision maker who faces a stream of messages that are all generated by the same known signal $z \in [\frac{1}{2}, 1]$. In this case, for any λ , a pair of a message and an action can generate only one single posterior belief $b \in \{g_z^{-1}(\lambda), \lambda, g_z(\lambda)\}$, and any aggregator satisfies $|H[U, \{b\}] - H[V, \{b\}]| = |U(b) - V(b)| \leq \|U - V\|_\infty$. Thus, by the same arguments as in Proposition 3.5, $U_{\{z\}}^*$ is well-defined and, by consistency, independent of H .²¹

For $z = 1$ the problem simplifies and one finds that $U_{\{1\}}^*(\lambda) = (1 - p)u(\lambda) + p$. In particular, any optimal strategy requires $\sigma^*(\lambda) = (l, l)$ for all $\lambda \in (0, 1)$. For the other extreme, $z = \frac{1}{2}$, there is no difference between learning and ignoring as $g_{\frac{1}{2}}(\lambda) = g_{\frac{1}{2}}^{-1}(\lambda) = \lambda$. In this case every strategy is optimal and we find that $U_{\{\frac{1}{2}\}}^*(\lambda) = u(\lambda)$ for all $\lambda \in [0, 1]$.

²¹This observation serves solely as a sanity check: as there is no ambiguity involved, the decision maker's ambiguity attitude should not affect the result.




Figure 1: Beliefs where the agent prefers to learn or ignore or is indifferent

Thus, we shall focus on $z \in (\frac{1}{2}, 1)$. For such z let $\lambda^*(z)$ be the unique positive solution²² of

$$g_z(\lambda) = 1 - \lambda. \quad (7)$$

If the decision maker has prior $\lambda^*(z)$, her updated belief after message 1 will be $1 - \lambda^*(z)$; and if she has prior $1 - \lambda^*(z)$, her updated belief after message 0 will be $\lambda^*(z)$. Our first theorem shows that the decision maker will update towards 1 only if she has a belief of at least $\lambda^*(z)$ and towards 0 only if she has a belief of at most $1 - \lambda^*(z)$. In particular, her optimal strategy is independent of the continuation probability p .

Theorem 4.1. *Let $z \in (\frac{1}{2}, 1)$ and let $H: \mathcal{F} \times \mathcal{P} \rightarrow \mathbb{R}$ be an aggregator. Let $U_{\{z\}}^*$ be defined as in (5), and let $\lambda^*(z)$ as in (7). Then a strategy $\sigma \in \mathcal{S}$ is optimal given H and $\{z\}$ if and only if*

$$\sigma_0(\lambda) = \begin{cases} l & \text{for } 0 < \lambda < 1 - \lambda^*(z), \\ i & \text{for } 1 - \lambda^*(z) < \lambda < 1, \end{cases} \quad \sigma_1(\lambda) = \begin{cases} i & \text{for } 0 < \lambda < \lambda^*(z), \\ l & \text{for } \lambda^*(z) < \lambda < 1. \end{cases} \quad (8)$$

Figure 1 illustrates Theorem 4.1 for a message of 1. If $\lambda < \lambda^*(z)$ (Panel 1a), $U_{\{z\}}^*(g_z(\lambda)) < U_{\{z\}}^*(\lambda)$ so that updating would decrease the agent's utility, hence any optimal strategy requires $\sigma_1(\lambda) = i$. If $\lambda = \lambda^*(z)$ (Panel 1b), she is precisely indifferent between learning and ignoring, and whenever $\lambda > \lambda^*(z)$ (Panel 1c), optimality requires $\sigma_1(\lambda) = l$. For later reference, we collect some properties of $U_{\{z\}}^*$ which follow from the proof of Theorem 4.1 in the following corollary.

²²The properties of $\lambda^*(z)$ are investigated more thoroughly in Corollary 4.3.




Figure 2: The graphs of $z \mapsto \lambda^*(z)$ (solid line) and $z \mapsto 1 - z$ (dashed line)

Corollary 4.2. *The function $U_{\{z\}}^*$ is continuous, satisfies $U_{\{z\}}^*(1 - \lambda) = U_{\{z\}}^*(\lambda)$ for all $\lambda \in [0, 1]$ as well as $U_{\{z\}}^*(0) = U_{\{z\}}^*(1) = 1$, and is strictly decreasing on $[0, \lambda^*(z)]$. Moreover,*

$$\begin{aligned} U_{\{z\}}^*(g_z(\lambda)) &< U_{\{z\}}^*(g_z^{-1}(\lambda)) && \text{for } \lambda \in (0, \tfrac{1}{2}) \text{ and} \\ U_{\{z\}}^*(g_z(\lambda)) &> U_{\{z\}}^*(g_z^{-1}(\lambda)) && \text{for } \lambda \in (\tfrac{1}{2}, 1). \end{aligned}$$

Theorem 4.1 shows that even without ambiguity there are beliefs where it is optimal for a strategic decision maker to ignore contradicting messages. This region, denoted by $I(z) = [0, \lambda^*(z)] \cup [1 - \lambda^*(z), 1]$ increases as z decreases.

Corollary 4.3. *Let $z, z' \in (\frac{1}{2}, 1]$. Then $I(z) \subseteq I(z')$ if and only if $z \geq z'$.*

The relation between z and $\lambda^*(z)$ is depicted in Figure 2. As can be seen there (and analytically verified), $\lambda^*(z) > 1 - z$ for all $z \in (\frac{1}{2}, 1)$.

5 Information Selection under Ambiguity

In this section we shall return to the decision problem when z is not fixed, but stems from some compact set $Z \subseteq (\frac{1}{2}, 1]$.²³ We first derive the optimal strategy of a decision maker

²³We exclude $\frac{1}{2} \in Z$ to avoid some tedious case distinctions.

with Gilboa-Schmeider preferences. Afterwards we consider an alternative approach in which the decision maker is unaware that she might reassess her view on possible signals after having received a message.

5.1 Optimal Strategies

Our main finding is that a Gilboa-Schmeidler decision maker who faces a compact set Z of signals will behave exactly as if she knew that every message were generated by $\underline{z} = \min Z$.

Theorem 5.1. *Let $Z \subseteq (\frac{1}{2}, 1]$ be compact and let $H: \mathcal{F} \times \mathcal{P} \rightarrow \mathbb{R}$ be the Gilboa-Schmeidler aggregator in (2). Let $\underline{z} = \min Z$. Then a strategy is optimal given Z if and only if it is optimal given $\{\underline{z}\}$. In particular, $U_Z^* = U_{\{\underline{z}\}}^*$.*

A Gilboa-Schmeidler decision maker facing signal set Z is behaviorally indistinguishable from a decision maker who knows that \underline{z} is the only possible signal. Moreover, such an agent always follows up on messages that confirm her prior belief but ignores contradictory information if her prior is sufficiently strong. The respective threshold λ^* only depends on the least precise signal that is deemed possible.

One might be inclined to argue that Theorem 5.1 is not surprising: after all a Gilboa-Schmeidler decision maker is pessimistic and should, hence, presume the message stem from the “worst” possible signal. We shall briefly illustrate that this line of thought is incorrect though as there is a difference between the “worst” and the “least precise” signal.

Example 5.2. Let the signal set be $Z = \{\frac{8}{10}, \frac{9}{10}\}$ such that $U_Z^* = U_{\{\frac{8}{10}\}}^*$ and $\lambda^*(\frac{8}{10}) = \frac{1}{3}$. Let the agent’s belief be $\lambda = \frac{1}{20}$ and suppose that message 1 realizes. Updating with respect to signal $z = \frac{8}{10}$ leads to an updated belief of $g_{\frac{8}{10}}(\frac{1}{20}) = \frac{4}{23}$, while updating with respect to $z = \frac{9}{10}$ would lead to $g_{\frac{9}{10}}(\frac{2}{10}) = \frac{9}{28}$. As $0 < \frac{4}{23} < \frac{9}{28} < \frac{1}{3}$ and $U_{\{\frac{8}{10}\}}^*$ is strictly decreasing in this area by Corollary 4.2, the more precise signal would lead to a lower utility. \square

The reason why the constellation in the previous example does not affect our finding in Theorem 5.1 is that all beliefs where the higher signal would lead to a lower utility lie in the area $I(\underline{z})$ where the strategic decision maker does not update at all.

5.2 Naive Decision Makers

Recall that at the beginning of any period the decision maker faces two sources of uncertainty: uncertainty about what message she will observe and uncertainty about the signal that will have generated that message. The definition of the anticipatory utility function in (3) made the implicit assumption that she evaluates these two sources independently of one another and that she is aware of this independence at the beginning of the period. Alternatively, she might be unaware that her preferences could require her to reassess what constitutes a worst case scenario. The following example illustrates the issue.

Example 5.3. Recall Example 5.2 with $Z = \{\frac{8}{10}, \frac{9}{10}\}$. At the prior of $\lambda = \frac{1}{20} < \frac{1}{3} = \lambda^*(\frac{8}{10})$, the worst that could happen to the decision maker is observing message 1. Thus, ex ante, she will evaluate her situation under the presumption that $z = \frac{8}{10}$ as for this z the probability of observing message 1 is maximal. Upon observing message 1, however, the worst case for her is that the message was generated by signal $z = \frac{9}{10}$ as seen above. The strategic decision maker in (3) takes this into account and uses two different values of z at the two different stages. \square

The decision maker we considered so far was aware that she might have a change of heart along the way, reflected by the use of $B_{Z, \sigma_m(\lambda)}^m(\lambda)$ as the set of possible beliefs for the inner aggregator in (3). A decision maker who is not aware of this will have an anticipatory utility function that satisfies

$$\hat{U}(\lambda) = (1 - p) u(\lambda) + p H \left[q^0(\lambda) H \left[\hat{U}, B_{\{\cdot\}, \sigma_0(\lambda)}^0(\lambda) \right] + q^1(\lambda) H \left[\hat{U}, B_{\{\cdot\}, \sigma_1(\lambda)}^1(\lambda) \right], Z \right]. \quad (9)$$

We call a decision maker with such an anticipatory utility function *naive*, as opposed to *sophisticated*, which would refer to the decision maker in (3), following a similar distinction made by O'Donoghue and Rabin (1999) in the context of time-inconsistent preferences.

Showing that Proposition 3.5 carries over to the naive decision maker, i.e., that there exists a unique naive anticipatory utility function for each $\sigma \in \mathcal{S}$, henceforth denoted by \hat{U}_Z^σ , can be done analogously to the previous analysis and is omitted.

Example 5.4. Recall Example 3.4 with strategy ρ defined by $\rho(\lambda) = (l, l)$ for all $\lambda \in [0, 1]$. The anticipatory utility function of a naive Gilboa-Schmeidler decision is the unique

solution of

$$\hat{U}_Z^\rho(\lambda) = (1-p)u(\lambda) + p \min_{z \in Z} \left\{ q_z^0(\lambda) \hat{U}_Z^\rho(g_z^{-1}(\lambda)) + q_z^1(\lambda) \hat{U}_Z^\rho(g_z(\lambda)) \right\}$$

for all $\lambda \in [0, 1]$. □

Since $B_{\{z\}, \sigma_m(\lambda)}^m(\lambda)$ is a singleton set for all $z \in Z$ and all $\sigma \in \mathcal{S}$, the naive Gilboa-Schmeidler decision maker in (9) presumes only one expression take its minimum, while their sophisticated counterpart in (3) expects the minimum over all possible signals twice—she is being pessimistic “once more”. Hence, it is not surprising that the anticipatory utility of a naive decision maker is higher.²⁴

Lemma 5.5. *Let $H: \mathcal{F} \times \mathcal{P} \rightarrow \mathbb{R}$ be the Gilboa-Schmeidler aggregator in (2). For any compact set $Z \subseteq [\frac{1}{2}, 1]$ and any $\sigma \in \mathcal{S}$, it holds that $U_Z^\sigma(\lambda) \leq \hat{U}_Z^\sigma(\lambda)$ for all $\lambda \in [0, 1]$.*

For the naive decision maker, the optimal utility at prior λ is

$$\hat{U}_Z^*(\lambda) = \sup_{\sigma \in \mathcal{S}} \hat{U}_Z^\sigma(\lambda). \tag{10}$$

We show that in terms of optimal strategies and optimal anticipatory utilities there is no difference between a naive and a sophisticated decision maker.

Theorem 5.6. *Let $Z \subseteq (\frac{1}{2}, 1]$ be compact and let $H: \mathcal{F} \times \mathcal{P} \rightarrow \mathbb{R}$ be the Gilboa-Schmeidler aggregator in (2). Let $\underline{z} = \min Z$. Then $\hat{U}_Z^* = U_Z^*$, and $\hat{U}_Z^* = \hat{U}_Z^\sigma$ if and only if σ satisfies (8) for $z = \underline{z}$.*

While Example 5.3 demonstrates that naive and sophisticated Gilboa-Schmeidler decision makers might evaluate some situations differently, Theorem 5.6 shows that these differences have no effect on the agents’ behavior or her anticipatory utility function. In particular, $\hat{U}_Z^* = U_Z^* = U_{\{\underline{z}\}}^* = \hat{U}_{\{\underline{z}\}}^*$, so that no matter whether she is sophisticated or naive, a Gilboa-Schmeidler decision maker will always act and feel as if she knew that \underline{z} were the only possible signal.

²⁴We are not the first to highlight this conflict between being happy and overthinking uncertainty. It has been beautifully illustrated, for instance, in Woody Allen’s “Annie Hall”. But there is also empirical evidence supporting the other direction: that happy people appear less clever (cf. Barasch, Levine, and Schweitzer, 2016).

6 Benchmark Comparison

We have derived the optimal behavior of a decision maker who can strategically ignore information. We shall call such a decision maker *strategic* to distinguish her from a Bayesian decision maker, who mechanically updates his belief after every signal. In this section, we shall contrast the two with respect to their anticipatory utility as well as their success probabilities, i.e., the probabilities with which they make the correct choice after the stream of messages stops.

6.1 Ignorance is a Bliss

Recall that a Bayesian decision maker uses strategy ρ with $\rho(\lambda) = (l, l)$ for all $\lambda \in [0, 1]$. He will learn z after every message and update his belief accordingly. Both for the sophisticated and the naive decision maker there are, by Proposition 3.5, unique anticipatory utility functions U_Z^ρ and \hat{U}_Z^ρ , respectively. Moreover, $U_Z^*(\lambda) \geq U_Z^\rho(\lambda)$ and $\hat{U}_Z^*(\lambda) \geq \hat{U}_Z^\rho(\lambda)$ for all $\lambda \in [0, 1]$ by definition. We show that these inequalities are strict for all $\lambda \in (0, 1)$.

Theorem 6.1. *Let $Z \subseteq (\frac{1}{2}, 1)$ be compact. Then $U_Z^*(\lambda) > U_Z^\rho(\lambda)$ and $\hat{U}_Z^*(\lambda) > \hat{U}_Z^\rho(\lambda)$ for all $\lambda \in (0, 1)$.*

For any non-degenerate belief, the anticipatory utility of a strategic decision maker is higher than that of a Bayesian decision maker. That is, strategic ignorance makes people happier, as long as the stream of signals continues.

6.2 Final Outcome

We shall turn to the probabilities with which the strategic and the Bayesian decision maker choose correctly at the end of the message stream. Throughout the subsection, we assume that the agents' prior belief λ is correct and that they are facing a known signal $z \in (\frac{1}{2}, 1)$. The ex-ante probability that a decision maker with initial belief λ and strategy σ makes the correct choice is denoted by $r_z^\sigma(\lambda)$.²⁵

²⁵At this point, we are a bit inaccurate as we do not formally define the concerning probability space. This space would have to contain the correct state of the world ω , the number of signals T , and for each T , the sequence of messages $(m_t)_{t=1}^T$. Defining such a space is not very difficult; however, it requires some cumbersome notation and does not provide additional insight.

As the optimal strategy for the strategic decision maker is not unique, we will focus on σ_z^{**} , which is defined by

$$\sigma_z^{**}(\lambda) = \begin{cases} (l, i) & \text{for } \lambda \in [0, \lambda^*(z)], \\ (l, l) & \text{for } \lambda \in (\lambda^*(z), 1 - \lambda^*(z)), \\ (i, l) & \text{for } \lambda \in [1 - \lambda^*(z), 1]. \end{cases} \quad (11)$$

By Theorem 4.1, this strategy is optimal. Observe that if $T \geq 1$, i.e., if there is at least one message, then $\lambda_1 \in I(z)$, so that for $t \geq 2$ the strategic decision maker will not learn upon a contradicting message. Thus, the agent exhibits an extreme *primacy effect* and her final choice is determined at latest after the first message. This allows us to easily derive the probability that her choice is correct.

Proposition 6.2. *For initial belief $\lambda \in [0, 1]$, the probability for a correct decision under σ_z^{**} is*

$$r_z^{\sigma_z^{**}}(\lambda) = \begin{cases} u(\lambda) & \text{if } \lambda \in I(z), \\ pz + (1 - p)u(\lambda) & \text{otherwise.} \end{cases}$$

An immediate consequence of this proposition is that the ex-ante probability of a correct choice is increasing in z . While it does not change for “extreme” initial beliefs, it strictly increases for “moderate” beliefs.

Corollary 6.3. *Let $1 \geq z > z' \geq \frac{1}{2}$. Then $r_z^{\sigma_z^{**}}(\lambda) = r_{z'}^{\sigma_{z'}^{**}}(\lambda)$ for all $\lambda \in I(z)$ and $r_z^{\sigma_z^{**}}(\lambda) > r_{z'}^{\sigma_{z'}^{**}}(\lambda)$ for all $\lambda \in (\lambda^*(z), 1 - \lambda^*(z))$.*

Finding the ex-ante probability that a Bayesian decision maker chooses correctly is more complex, as it does not solely depend on the first message. Its precise formula for each $\lambda \in (0, 1)$ is derived within the proof of the next theorem.

Theorem 6.4. *For any $\lambda \in (0, 1)$ and $z \in (\frac{1}{2}, 1)$ it holds that $r_z^\rho(\lambda) > r_z^{\sigma_z^{**}}(\lambda)$.*

Theorem 6.4 shows that strategic news selection lowers the decision maker’s long-run welfare as it decreases her probability to choose correctly in the end. This contrasts with Theorem 6.1 which states that as long as the stream of signals continues, the strategic

decision maker is happier than the Bayesian decision maker. Thus, while strategic ignorance increases instantaneous well-being, it is detrimental in the long run if the initial belief is correct.

7 Social Outcome

Consider a population N of decision makers with initial beliefs $(\lambda^i)_{i \in N}$ who face a stream of signals drawn from some compact set $Z \subseteq (\frac{1}{2}, 1)$ and who select their information according to $\sigma_{\underline{z}}^{**}$ as in (11). By Theorems 4.1 and 5.6 it holds that $\lambda_t^i \in I(z)$ for all $t \geq 1$ and all $i \in N$. Thus, already after the first signal, polarization is inevitable and extreme: There will only be “extremists”, who never look at messages contradicting their worldview. The “political center” $(\lambda^*(\underline{z}), 1 - \lambda^*(\underline{z}))$ is empty and will stay so forever.

Polarization even increases over time. After the first message, all beliefs will be in $[0, \lambda^*(\underline{z})]$ (“left-extremists”) or in $[1 - \lambda^*(\underline{z}), 1]$ (“right-extremists”). Subsequently, the beliefs of all left-extremists will (for $z \neq 1$) monotonically decrease and approach 0, while the beliefs of all right-wing extremists will monotonically increase towards 1. In particular, for large T , all beliefs will be close to 0 or 1.

The set of decision makers who choose the correct outcome depends on the initial distribution of priors and Z . Corollaries 4.3 and 6.3 imply that this set increases as \underline{z} increases. Hence, more people will make a correct choice if the least precise information source becomes more precise.

8 Extensions

Two extensions of our model seem straightforward: using aggregators other than Gilboa-Schmeider, and using a non-binary state space. Neither of them is trivial and general results about the structure of optimal strategies remain unclear.

8.1 Aggregation of Second Order Beliefs

If the decision maker knows the distribution of signals, she can exploit this additional information. Modifying Definition 3.2, one can then define an aggregator as a map $H: \mathcal{F}(D) \times \Delta(D) \rightarrow \mathbb{R}$ that satisfies:

Consistency. For all $x \in D$ and the point measure δ_x on x it holds that $H(U, \delta_x) = U(x)$.

Monotonicity. For all $\phi \in \Delta(D)$ and $U, V: D \rightarrow \mathbb{R}$ with $U|_{\text{supp}(\phi)} \geq V|_{\text{supp}(\phi)}$ it holds that $H(U, \phi) \geq H(V, \phi)$. If $U|_{\text{supp}(\phi)} \geq V|_{\text{supp}(\phi)} + \varepsilon$ for some $\varepsilon > 0$ it holds that $H(U, \phi) > H(V, \phi)$.

For instance, an ambiguity neutral decision maker with smooth preferences in the sense of Klibanoff, Marinacci, and Mukerji (2005) with expected utility function U and second order belief $F \in \Delta(D)$ uses the aggregator

$$H[U, F] = \mathbb{E}_F[U]. \quad (12)$$

Let $F \in \Delta\left(\left[\frac{1}{2}, 1\right]\right)$ be a probability measure over the set of signals and let F_λ^m denote the probability measure on Z conditional on the prior belief being $\lambda \in (0, 1)$ and the message being m . That is, for any measurable subset $Z' \subseteq Z$,

$$F_\lambda^m(Z') = \frac{\int_{Z'} q_z^m(\lambda) F(dz)}{\int_Z q_z^m(\lambda) F(dz)}.$$

Moreover, define four maps

$$\begin{aligned} \psi_{\lambda,l}: Z &\rightarrow [0, 1], z \mapsto g_z^{-1}(\lambda), & \psi_{\lambda,i}: Z &\rightarrow [0, 1], z \mapsto \lambda, \\ \varphi_{\lambda,l}: Z &\rightarrow [0, 1], z \mapsto g_z(\lambda), & \varphi_{\lambda,i}: Z &\rightarrow [0, 1], z \mapsto \lambda \end{aligned}$$

and push-forward measures $G_{F,a}^0(\lambda) := F_\lambda^0 \circ \psi_{\lambda,a}^{-1}$ and $G_{F,a}^1(\lambda) := F_\lambda^1 \circ \varphi_{\lambda,a}^{-1}$. Then $G_{F,l}^0(\lambda)$ is the probability measure over posterior beliefs given prior λ and observed message 0 before the generated signal is learned. The measure $G_{F,l}^1(\lambda)$ is defined accordingly for observed message 1. The anticipatory utility function of a sophisticated decision maker with aggregator H is then defined by

$$U_F^\sigma(\lambda) = (1-p)u(\lambda) + pH \left[q^0(\lambda) H \left[U_F^\sigma, G_{F,\sigma_0(\lambda)}^0(\lambda) \right] + q^1(\lambda) H \left[U_F^\sigma, G_{F,\sigma_1(\lambda)}^1(\lambda) \right], F \right]$$

and the anticipatory utility function of a naive decision maker is defined by

$$\hat{U}_F^\sigma(\lambda) = (1-p)u(\lambda) + pH \left[q^0(\lambda) H \left[\hat{U}_F^\sigma, G_{\delta^{\cdot}, \sigma_0(\lambda)}^0(\lambda) \right] + q^1(\lambda) H \left[\hat{U}_F^\sigma, G_{\delta^{\cdot}, \sigma_1(\lambda)}^1(\lambda) \right], F \right].$$

It is easy to see that Proposition 3.5 carries over and that the aggregator in (12) satisfies

its premise. Hence, both the sophisticated and the naive ambiguity neutral decision maker with smooth preferences have a unique anticipatory utility function which are fixed points of the functionals $T_F^\sigma, \hat{T}_F^\sigma: \mathcal{F}([0, 1]) \rightarrow \mathcal{F}([0, 1])$, defined by

$$\begin{aligned} T_F^\sigma U(\lambda) &= (1 - p)u(\lambda) + pH [q^0(\lambda)H [U, G_{F, \sigma_0(\lambda)}^0(\lambda)] + q^1(\lambda)H [U, G_{F, \sigma_1(\lambda)}^1(\lambda)], F], \\ \hat{T}_F^\sigma U(\lambda) &= (1 - p)u(\lambda) + pH [q^0(\lambda)H [U, G_{\delta, \sigma_0(\lambda)}^0(\lambda)] + q^1(\lambda)H [U, G_{\delta, \sigma_1(\lambda)}^1(\lambda)], F]. \end{aligned}$$

We show that these functionals coincide. This contrasts with Example 5.3 which shows that such a statement fails to hold in the case of Gilboa-Schmeidler preferences.

Lemma 8.1. *Let H be defined as in (12). Then $T_F^\sigma U(\lambda) = \hat{T}_F^\sigma U(\lambda)$ for any $F \in \Delta([\frac{1}{2}, 1])$, $\sigma \in \mathcal{S}$, $U \in \mathcal{F}$, and $\lambda \in [0, 1]$.*

The previous lemma might not be entirely surprising as the sophisticated decision maker essentially takes an expected value of a conditional expected value, while the naive decision maker take the expected value only once. In contrast, it might not be true for a decision maker who is not ambiguity neutral. An immediate consequence is that the anticipatory utility functions coincide for the sophisticated and the naive decision maker.

Corollary 8.2. *Let H be defined as in (12). For all $F \in \Delta[\frac{1}{2}, 1]$ and all $\sigma \in \mathcal{S}$, it holds that $U_F^\sigma = \hat{U}_F^\sigma$.*

Hence, in case of ambiguity neutral preferences, we do not have to distinguish between sophisticated and naive decision makers. The existence of optimal strategies can be derived as in Proposition 3.6, and they are identical for both types of decision makers, so that $U_F^*(\lambda) = \sup_{\sigma \in \mathcal{S}} U_F^\sigma(\lambda) = \sup_{\sigma \in \mathcal{S}} \hat{U}_F^\sigma(\lambda) = \hat{U}_F^*(\lambda)$ for all λ . What remains unclear, however, is whether optimal strategies have a similar threshold structure as those of the Gilboa-Schmeidler decision maker. In the following example, a numerical analysis suggests that they might.

Example 8.3. Recall Example 5.2 where $Z = \{\frac{8}{10}, \frac{9}{10}\}$ and let F be the uniform distribution on Z . Suppose the decision maker has smooth preferences and is ambiguity neutral. The anticipatory utility function U_F^* is depicted by the solid line in Figure 3. The function $\lambda \mapsto \mathbb{E}_{F_\lambda^1} [U_Z^*(g(\lambda))]$ is depicted by the dotted line. The two functions have a unique intersection, denoted $\lambda^*(F)$. The optimal strategies require that the agent ignores any 1-message for $\lambda < \lambda^*(F)$ and updates after any 1-message for $\lambda > \lambda^*(F)$. Numerically, we obtain that $\lambda^*(F) \approx 0.2823 \in (\frac{1}{4}, \frac{1}{3}) = (\lambda^*(\frac{9}{10}), \lambda^*(\frac{8}{10}))$. \square




Figure 3: Anticipatory utility function for ambiguity neutral smooth preferences.

8.2 More than Two States

The model can be generalized to more than two symmetric states, a message space that coincides with the state space, and a set of symmetric signals that generate with probability $z \geq \frac{1}{n}$ a correct message and with probability $\frac{1-z}{n-1}$ each other message. While the existence of optimal strategies can be proven as before, it remains unclear whether they will have the same threshold structure, as the proof of Theorem 4.1, in particular part 2b, does not generalize to more than two states.

9 Conclusion

This paper investigates the behavior of a decision maker who receives a stream of messages of known or unknown quality about the true state of the world. At each time period, she faces the decision whether to learn how the message has been generated and update her belief accordingly, or to discard the message and stick to her previous conviction. This “free disposal of information” is inspired by the way users selectively read news articles posted on social media platforms. If there is no ambiguity and under the assumption that the decision maker is Markovian, she will discard contradictory messages if her belief is sufficiently strong. This remains true under ambiguity if she has Gilboa-Schmeidler preferences, where the respective belief threshold depends solely on the worst possible signal.

When comparing her behavior to a decision maker who updates after each message,

we find that she is happier as long as new signals keep coming. However, the possibility to ignore signals leads to a higher probability of the wrong decision in the end. Moreover, on the macroscopic level, such behavior will lead to polarization.

The model neglects some crucial features of modern social media by assuming that messages are generated independently of the agent’s prior belief. Yet, extending the models by algorithms, echo chambers, or and targeted news that support the decision maker’s belief more often than not, will even intensify its already quite extreme results.

Our analysis has normative consequences on the individual as well as on the social level. For the single agent who wants to make a well-founded decision, it pays off to leave the own comfort zone and also have a look at information that challenges one’s prior conviction, even if this decreases the own welfare in the short run. A social planner that wants to have well-informed citizens who decide wisely and also listen to people with different opinions should ensure that the minimal possible signal precision—i.e., the quality of the “worst media outlet”—is as high as possible.

References

- Aiello, L., Barrat, A., Schifanella, R., Cattuto, C., Markines, B., and Menczer, F. (2012): “Friendship Prediction and Homophily in Social Media”. In: *ACM Transactions on the Web* 6(2), 9:1–9:33.
- Allahverdyan, A. E. and Galstyan, A. (2014): “Opinion Dynamics with Confirmation Bias”. In: *PLoS ONE* 9(7), e99557.
- Allcott, H., Braghieri, L., Eichmeyer, S., and Gentzkow, M. (2020): “The Welfare Effects of Social Media”. In: *American Economic Review* 110(3), pp. 629–676.
- Baliga, S., Hanany, E., and Klibanoff, P. (2013): “Polarization and Ambiguity”. In: *American Economic Review* 103(7), pp. 3071–3083.
- Barasch, A., Levine, E. E., and Schweitzer, M. E. (2016): “Bliss is ignorance: How the magnitude of expressed happiness influences perceived naiveté and interpersonal exploitation”. In: *Organizational Behavior and Human Decision Processes* 137, pp. 184–206.
- Blackwell, D. (1951): “Comparison of experiments”. In: *Proceedings of the Second Berkeley Symposium on Mathematical Statistics and Probability*, pp. 93–102.
- (1953): “Equivalent comparisons of experiments”. In: *The Annals of Mathematical Statistics* 24(2), pp. 265–272.

- Bracha, A. and Brown, D. J. (2012): “Affective Decision Making: A Theory of Optimism Bias”. In: *Games and Economic Behavior* 75, pp. 67–80.
- Brunnermeier, M. K. and Parker, J. A. (2005): “Optimal Expectations”. In: *American Economic Review* 95(4), pp. 1092–1118.
- Caplin, A. and Leahy, J. (2001): “Psychological Expected Utility Theory and Anticipatory Feelings”. In: *The Quarterly Journal of Economics* 116(1), pp. 55–79.
- Carrillo, J. D. and Mariotti, T. (2000): “Strategic Ignorance as a Self-Disciplining Device”. In: *Review of Economic Studies* 67, pp. 529–544.
- Che, Y.-K. and Mierendorff, K. (2019): “Optimal Dynamic Allocation of Attention”. In: *American Economic Review* 109(8), pp. 2993–3029.
- Chen, J. Y. (2022): “Biased Learning Under Ambiguous Information”. In: *Journal of Economic Theory* 203, p. 105492.
- Chiang, C.-F. and Knight, B. (2011): “Media Bias and Influence: Evidence from Newspaper Endorsements”. In: *Review of Economic Studies* 78, pp. 795–820.
- Eil, D. and Rao, J. M. (2011): “The Good News-Bad News Effect: Asymmetric Processing of Objective Information about Yourself”. In: *Review of Economic Studies* 74, pp. 1275–1303.
- Epley, N. and Gilovich, T. (2016): “The Mechanics of Motivated Reasoning”. In: *Journal of Economic Perspectives* 30(3), pp. 133–140.
- Eveland Jr., W. P. and Dunwoody, S. (2002): “An Investigation of Elaboration and Selective Scanning as Mediators of Learning From the Web Versus Print”. In: *Journal of Broadcasting & Electronic Media* 46(1), pp. 34–53.
- Festinger, L. (1957): *A Theory of Cognitive Dissonance*. Stanford University Press.
- Fryer Jr., R. G., Harms, P., and Jackson, M. O. (2019): “Updating Beliefs when Evidence is Open to Interpretation: Implications for Bias and Polarization”. In: *Journal of the European Economic Association* 17(5), pp. 1470–1501.
- Gentzkow, M. and Shapiro, J. M. (2006): “Media Bias and Reputation”. In: *Journal of Political Economy* 114(2), pp. 280–316.
- Gilboa, I. and Schmeidler, D. (1989): “Maxmin Expected Utility with Non-Unique Prior”. In: *Journal of Mathematical Economics* 18, pp. 141–153.
- Glaeser, E. L. and Sunstein, C. R. (2013): “Why Does Balanced News Produce Unbalanced Views?” In: *NBER Working Paper Series* Working Paper 18975.

- Hart, W., Albarracin, D., Eagly, A. H., Brechan, I., Lindberg, M. J., and Merrill, L. (2009): “Feeling Validated Versus Being Correct: A Meta-Analysis of Selective Exposure to Information”. In: *Psychological Bulletin* 135(4), pp. 555–588.
- Klibanoff, P., Marinacci, M., and Mukerji, S. (2005): “A Smooth Model of Decision Making under Ambiguity”. In: *Econometrica* 76(6), pp. 1849–1892.
- Kószegi, B. (2017): “Utility from anticipation and personal equilibrium”. In: *Economic Theory* 44(3), pp. 415–444.
- Kruikemeier, S., Lecheler, S., and Boyer, M. M. (2017): “Learning From News on Different Media Platforms: An Eye-Tracking Experiment”. In: *Political Communication* 35(1), pp. 75–96.
- Kubin, E. and von Sikorski, C. (2021): “The Role of (Social) Media in Political Polarization: a Systematic Review”. In: *Annals of the International Communication Association* 45(3), pp. 188–206.
- Loewenstein, G. (1987): “Anticipation and the Valuation of Delayed Consumption”. In: *The Economic Journal* 97(387), pp. 666–684.
- Lord, C. G., Ross, L., and Lepper, M. R. (1979): “Biased Assimilation and Attitude Polarization: The Effects of Prior Theories on Subsequently Considered Evidence”. In: *Journal of Personality and Social Psychology* 37(11), pp. 2098–2109.
- Machina, M. J. (1989): “Dynamic Consistency and Non-Expected Utility Models of Choice Under Uncertainty”. In: *Journal of Economic Literature* 27(4), pp. 1622–1668.
- Maćkowiak, B., Matějka, F., and Wiederholt, M. (2023): “Rational Inattention: A Review”. In: *Journal of Economic Literature* 61(1), pp. 226–273.
- Nickerson, R. S. (1998): “Confirmation Bias: A Ubiquitous Phenomenon in Many Guises”. In: *Review of General Psychology* 2(2), pp. 175–220.
- O’Donoghue, R. and Rabin, M. (1999): “Doing it now or Later”. In: *American Economic Review* 89(1), pp. 103–124.
- Pew Research Center (2022): *Social Media Seen as Mostly Good for Democracy Across Many Nations, But U.S. is a Major Outlier*. URL: https://www.pewresearch.org/global/2022/12/06/social-media-seen-as-mostly-good-for-democracy-across-many-nations-but-u-s-is-a-major-outlier/?utm_content=buffer865be&utm_medium=social&utm_source=twitter.com&utm_campaign=buffer (visited on 08/22/2023).
- Prior, M. (2013): “Media and Political Polarization”. In: *Annual Review of Political Science* 16, pp. 101–127.

- Rabin, M. and Schrag, J. L. (1999): “First Impressions Matter: A Model of Confirmatory Bias”. In: *The Quarterly Journal of Economics* 114(1), pp. 37–82.
- Reuters Institute for the Study of Journalism (2023): *Reuters Institute Digital News Report 2023*. URL: <https://reutersinstitute.politics.ox.ac.uk/digital-news-report/2023> (visited on 08/22/2023).
- Schneider, W. and Shiffrin, R. M. (1977): “Controlled and Automatic Human Information Processing: I. Detection, Search, and Attention”. In: *Psychological Review* 84(1), pp. 1–66.
- Sims, C. A. (2003): “Implications of rational inattention”. In: *Journal of Monetary Economics* 50, pp. 665–690.
- Statista (2024): *Media usage in an internet minute*. URL: <https://www.statista.com/statistics/195140/new-user-generated-content-uploaded-by-users-per-minute/> (visited on 03/25/2024).
- Stone, D. F. and Wood, D. H. (2018): “Cognitive Dissonance, Motivated Reasoning, and Confirmation Bias: Applications in Industrial Organization”. In: *Handbook of Behavioral Industrial Organization*. Ed. by V. J. Tremblay, E. Schroeder, and C. H. Tremblay.
- Suen, W. (2004): “The Self-Perpetuation of Biased Beliefs”. In: *The Economic Journal* 114, pp. 377–396.
- Sülflow, M., Schäfer, S., and Winter, S. (2019): “Selective Attention in the News Feed: An Eye-Tracking Study on the Perception and Selection of Political News Posts on Facebook”. In: *new media & society* 21(1), pp. 168–190.
- Yang, J. and Grabe, M. E. (2011): “Knowledge Acquisition Gaps: A Comparison of Print Versus Online News Sources”. In: *new media & society* 13(8), pp. 1211–1227.

A Proofs

Proof of Proposition 3.5. For any strategy σ define the operator $T_Z^\sigma: \mathcal{F}([0, 1]) \rightarrow \mathcal{F}([0, 1])$ by

$$T_Z^\sigma U(\lambda) = (1 - p) u(\lambda) + p H [q^0(\lambda) H [U, B_{Z, \sigma_0}^0(\lambda)] + q^1(\lambda) H [U, B_{Z, \sigma_1}^1(\lambda)], Z] \quad (13)$$

for all $\lambda \in [0, 1]$. Then T_Z^σ is a contraction since

$$\begin{aligned}
\|T_Z^\sigma U - T_Z^\sigma V\|_\infty &= p \sup_{\lambda \in [0,1]} |H [q^0(\lambda) H [U, B_{Z,\sigma_0(\lambda)}^0(\lambda)] + q^1(\lambda) H [U, B_{Z,\sigma_1(\lambda)}^1(\lambda)], Z] \\
&\quad - H [q^0(\lambda) H [V, B_{Z,\sigma_0(\lambda)}^0(\lambda)] + q^1(\lambda) H [V, B_{Z,\sigma_1(\lambda)}^1(\lambda)], Z]| \\
&\leq p \sup_{\lambda \in [0,1]} \left\| q^0(\lambda) H [U, B_{Z,\sigma_0(\lambda)}^0(\lambda)] + (1 - q^0(\lambda)) H [U, B_{Z,\sigma_1(\lambda)}^1(\lambda)] \right. \\
&\quad \left. - q^0(\lambda) H [V, B_{Z,\sigma_0(\lambda)}^0(\lambda)] - (1 - q^0(\lambda)) H [V, B_{Z,\sigma_1(\lambda)}^1(\lambda)] \right\|_{\infty, Z} \\
&= p \sup_{\lambda \in [0,1]} \sup_{z \in Z} |q_z^0(\lambda) (H [U, B_{Z,\sigma_0(\lambda)}^0(\lambda)] - H [V, B_{Z,\sigma_0(\lambda)}^0(\lambda)]) \\
&\quad + (1 - q_z^0(\lambda)) (H [U, B_{Z,\sigma_1(\lambda)}^1(\lambda)] - H [V, B_{Z,\sigma_1(\lambda)}^1(\lambda)])| \\
&\leq p \sup_{\lambda \in [0,1]} \sup_{z \in Z} (q_z^0(\lambda) \|U - V\|_\infty + (1 - q_z^0(\lambda)) \|U - V\|_\infty) \\
&= p \|U - V\|_\infty.
\end{aligned}$$

Thus, as $\mathcal{F}([0, 1])$ is a complete metric space, T_Z^σ has a unique fixed point by Banach's fixed point theorem, and this fixed point satisfies (3). \blacksquare

Proof of Proposition 3.6. Let σ^* satisfy (6). By definition it holds that $U_Z^* \geq U_Z^{\sigma^*}$, so we show that $U_Z^{\sigma^*} \geq U_Z^*$. By the monotonicity of H it holds that $T_Z^{\sigma^*} U \geq T_Z^{\sigma^*} V$ for all $U, V \in \mathcal{F}([0, 1])$ with $U \geq V$. Thus, since σ^* satisfies (6), and since $U_Z^* \geq U_Z^\sigma$ for every $\sigma \in \mathcal{S}$,

$$T_Z^{\sigma^*} U_Z^* \geq T_Z^\sigma U_Z^* \geq T_Z^\sigma U_Z^\sigma = U_Z^\sigma.$$

Hence, $T_Z^{\sigma^*} U_Z^*(\lambda) \geq \sup_{\sigma \in \mathcal{S}} U_Z^\sigma(\lambda) = U_Z^*(\lambda)$ for all $\lambda \in [0, 1]$. Iterated use of the monotonicity of $T_Z^{\sigma^*}$ reveals that $(T_Z^{\sigma^*})^n U_Z^* \geq U_Z^*$ for all $n \in \mathbb{N}$. Thus, $U_Z^{\sigma^*} = \lim_{n \rightarrow \infty} (T_Z^{\sigma^*})^n U_Z^* \geq U_Z^*$, as claimed.

For the other direction let $\sigma, \sigma^* \in \mathcal{S}$ be such that σ^* satisfies Equation (6) for all $\lambda \in [0, 1]$ and σ does not. By means of contradiction assume that $U_Z^\sigma = U_Z^*$. Then U_Z^* is a fixed point of T_Z^σ so that $T_Z^\sigma U_Z^*(\lambda) = U_Z^*(\lambda)$. Moreover, by the first part of the proof, $T_Z^{\sigma^*} U_Z^*(\lambda) = U_Z^*(\lambda)$. Thus,

$$T_Z^{\sigma^*} U_Z^*(\lambda) = T_Z^\sigma U_Z^*(\lambda) \quad \text{for all } \lambda \in [0, 1]. \quad (14)$$

However, there are $\delta > 0$ and, without loss of generality, $\lambda \geq \frac{1}{2}$ such that $H [U_Z^*, B_{Z,l}^1(\lambda)] - U_Z^*(\lambda) = \delta$ and $\sigma_1(\lambda) = i$. For all $z \in Z$ it holds that

$$\begin{aligned} & \left(q_z^0(\lambda) H [U_Z^*, B_{Z,\sigma_0^*}^0(\lambda)] + q_z^1(\lambda) H [U_Z^*, B_{Z,\sigma_1^*}^1(\lambda)] \right) \\ & \quad - \left(q_z^0(\lambda) H [U_Z^*, B_{Z,\sigma_0}^0(\lambda)] + q_z^1(\lambda) H [U_Z^*, B_{Z,\sigma_1}^1(\lambda)] \right) \\ & \geq q_{\min Z}^1(\lambda) \delta > 0, \end{aligned}$$

so that $T_Z^{\sigma^*} U_Z^*(\lambda) > T_Z^\sigma U_Z^*(\lambda)$ in contradiction to (14). ■

Proof of Theorem 4.1. In order to keep notation simple, let $\lambda^* = \lambda^*(z)$. One can easily verify that for all $\lambda > \lambda^*$ it holds that $u(g_z(\lambda)) - u(\lambda) > 0$. We need the following two lemmas for the proof. The first is straightforward and the proof is omitted.

Lemma A.1. *Let x, y, x' , and y' be positive real numbers with $x' \geq x$ and $y' \geq y$, as well as $x \geq y$ or $x' \geq y'$. Moreover, let $\alpha, \beta \in (0, 1)$ with $\alpha \geq \beta$. Then*

$$\alpha x' + (1 - \alpha) y' \geq \beta x + (1 - \beta) y. \quad (15)$$

If at least one of the inequalities is strict, then (15) is strict as well.

Lemma A.2. *Let \mathcal{C}_z^* be the set of all continuous functions on $[0, 1]$ that satisfy*

$$U(\lambda) = U(1 - \lambda) \quad \text{for all } \lambda \in [0, 1], \quad (16)$$

$$U(0) = U(1) = 1, \quad (17)$$

$$U(\lambda') - U(\lambda) \geq (1 - p)(\lambda' - \lambda) \quad \text{for all } 1 - \lambda^* < \lambda < \lambda' \leq 1, \quad (18)$$

$$U(g_z(\lambda)) - U(\lambda) \geq (1 - p) [u(g_z(\lambda)) - u(\lambda)] \quad \text{for all } \lambda \in [\lambda^*, 1]. \quad (19)$$

Then \mathcal{C}_z^ is a nonempty and closed (with respect to the supremum norm) subset of $\mathcal{F}([0, 1])$, and $T_{\{z\}}^\sigma U \in \mathcal{C}_z^*$ for all $U \in \mathcal{C}_z^*$ and all strategies σ that satisfy (8).*

Proof. It is clear that $u \in \mathcal{C}_z^*$ and that all the properties are preserved under uniform convergence, such that \mathcal{C}_z^* is nonempty and closed. Let $U \in \mathcal{C}_z^*$. To keep notation simple let $T = T_{\{z\}}^\sigma$. Verifying that TU satisfies (16) and (17) is straightforward. Continuity of TU at $\lambda \neq \lambda^*, 1 - \lambda^*$ is clear. Continuity at λ^* and $1 - \lambda^*$ follows from the properties of g_z and g_z^{-1} in Lemma 3.1 in combination with (7) and (16).

Let $U \in \mathcal{C}_z^*$. We show that TU satisfies (18), so let $1 - \lambda^* \leq \lambda < \lambda' \leq 1$. Then, by (18), it holds that $U(\lambda') > U(\lambda)$ and, by the strict monotonicity of g_z , we have $U(g_z(\lambda')) > U(g_z(\lambda))$ as well as $U(g_z(\lambda')) > U(\lambda')$. Hence, we can apply Lemma A.1 with $\alpha = q_z^1(\lambda')$, $\beta = q_z^1(\lambda)$, $x' = U(g_z(\lambda'))$, $y' = U(\lambda')$, $x = U(g_z(\lambda))$, and $y = U(\lambda)$ and get

$$q_z^1(\lambda')U(g_z(\lambda')) + q_z^0(\lambda')U(\lambda') > q_z^1(\lambda)U(g_z(\lambda)) + q_z^0(\lambda)U(\lambda)$$

where we also used $1 - q_z^0(\cdot) = q_z^1(\cdot)$. Together with the definition of σ we have

$$\begin{aligned} (TU)(\lambda') - (TU)(\lambda) &= (1 - p)[u(\lambda') - u(\lambda)] \\ &\quad + p[q_z^1(\lambda')U(g_z(\lambda')) + q_z^0(\lambda')U(\lambda') - q_z^1(\lambda)U(g_z(\lambda)) - q_z^0(\lambda)U(\lambda)] \\ &> (1 - p)[u(\lambda') - u(\lambda)] \\ &= (1 - p)(\lambda' - \lambda) \end{aligned}$$

as claimed.

We show that T preserves (19). For all $\lambda \in [0, 1]$ define

$$(\tilde{T}U)(\lambda) := \begin{cases} q_z^0(\lambda)U(g_z^{-1}(\lambda)) + q_z^1(\lambda)U(\lambda), & \lambda \in [0, \lambda^*], \\ q^0(\lambda)U(g_z^{-1}(\lambda)) + q_z^1(\lambda)U(g_z(\lambda)), & \lambda \in (\lambda^*, 1 - \lambda^*), \\ q_z^0(\lambda)U(\lambda) + q_z^1(\lambda)U(g_z(\lambda)), & \lambda \in [1 - \lambda^*, 1], \end{cases}$$

and observe that $(TU)(\lambda) = (1 - p)u(\lambda) + p(\tilde{T}U)(\lambda)$. Thus,

$$(TU)(g_z(\lambda)) - (TU)(\lambda) = (1 - p)[u(g_z(\lambda)) - u(\lambda)] + p[(\tilde{T}U)(g_z(\lambda)) - (\tilde{T}U)(\lambda)]$$

for all λ . Hence, in order to prove (19), it is sufficient to show that $(\tilde{T}U)(g_z(\lambda)) \geq (\tilde{T}U)(\lambda)$ for all $\lambda \in [\lambda^*, 1]$.

1. First, let $\lambda \in [1 - \lambda^*, 1]$. Then, since $z > \frac{1}{2}$, it holds that $g_z(\lambda) \in (1 - \lambda^*, 1]$, so that

$$\begin{aligned} (\tilde{T}U)(g_z(\lambda)) &= q_z^1(g(\lambda))U(g_z(g_z(\lambda))) + q_z^0(g_z(\lambda))U(g_z(\lambda)) \\ &\geq q_z^1(\lambda)U(g_z(\lambda)) + q_z^0(\lambda)U(\lambda) \\ &= (\tilde{T}U)(\lambda) \end{aligned}$$

where the inequality follows from (15), using that $U(g_z(g_z(\lambda))) \geq U(g_z(\lambda)) \geq U(\lambda)$ by the monotonicity of U on $[1 - \lambda^*, 1]$, and $q_z^1(g_z(\lambda)) \geq q_z^1(\lambda)$ by the monotonicity of q_z^1 .

2. Second, let $\lambda \in (\lambda^*, 1 - \lambda^*)$. Then, $g_z^{-1}(\lambda) \in [0, \lambda^*)$ and $g_z(\lambda) \in (1 - \lambda^*, 1]$. As U is strictly monotone in these areas, $U(g_z(\lambda))$ is increasing, and $U(g_z^{-1}(\lambda))$ is decreasing in λ .

By Lemma 3.1, it holds that $g_z(\frac{1}{2}) = 1 - g_z^{-1}(\frac{1}{2})$, so that by the symmetry of U , $U(g_z(\frac{1}{2})) = U(1 - g_z^{-1}(\frac{1}{2})) = U(g_z^{-1}(\frac{1}{2}))$. Consequently,

$$\begin{aligned} U(g_z(\lambda)) &> U(g_z^{-1}(\lambda)) \text{ for } \lambda \in (\frac{1}{2}, 1 - \lambda^*) \text{ and} \\ U(g_z(\lambda)) &< U(g_z^{-1}(\lambda)) \text{ for } \lambda \in (\lambda^*, \frac{1}{2}). \end{aligned} \tag{20}$$

- (a) Suppose first that $\lambda \in [\frac{1}{2}, 1 - \lambda^*)$. Then $U(g_z(g_z(\lambda))) \geq U(g_z(\lambda)) \geq U(g_z^{-1}(\lambda))$ and therefore

$$\begin{aligned} (\tilde{T}U)(g_z(\lambda)) &= q_z^1(g_z(\lambda))U(g_z(g_z(\lambda))) + q_z^0(g_z(\lambda))U(g_z(\lambda)) \\ &\geq U(g_z(\lambda)) \\ &\geq q_z^1(\lambda)U(g_z(\lambda)) + q_z^0(\lambda)U(g_z^{-1}(\lambda)) \\ &= (\tilde{T}U)(\lambda). \end{aligned}$$

- (b) Finally, let $\lambda \in (\lambda^*, \frac{1}{2})$. Since

$$U(g_z(g_z(\lambda^*))) = U(g_z(1 - \lambda^*)) = U(1 - g_z^{-1}(\lambda^*)) = U(g_z^{-1}(\lambda^*)),$$

the monotonicity of U on $[0, \lambda^*)$ and $(1 - \lambda^*, 1]$ together with (20) imply that

$$U(g_z(g_z(\lambda))) \geq U(g_z^{-1}(\lambda)) \geq U(g_z(\lambda)).$$

Thus, using that $q^1(g_z(\lambda)) > 0$ and $q_z^1(\lambda) = q_z^0(1 - \lambda) > q_z^0(g_z(\lambda))$ for $\lambda > \lambda^*$, we obtain

$$\begin{aligned} (\tilde{T}U)(g_z(\lambda)) &= q_z^1(g_z(\lambda))U(g_z(g_z(\lambda))) + q_z^0(g_z(\lambda))U(g_z(\lambda)) \\ &= q_z^1(g_z(\lambda))U(g_z(g_z(\lambda))) + (q_z^1(\lambda) + q_z^0(g_z(\lambda)) - q_z^1(\lambda))U(g_z(\lambda)) \end{aligned}$$

$$\begin{aligned}
&\geq q_z^1(g_z(\lambda)) U(g_z^{-1}(\lambda)) + q_z^1(\lambda) U(g_z(\lambda)) \\
&\quad + (q_z^0(g_z(\lambda)) - q_z^1(\lambda)) U(g_z^{-1}(\lambda)) \\
&= q_z^1(\lambda) U(g_z(\lambda)) + (1 - q_z^1(\lambda)) U(g_z^{-1}(\lambda)) \\
&= (\tilde{T}U)(\lambda)
\end{aligned}$$

as required. ■

Let $U \in \mathcal{C}_z^*$. By (19) it holds that

$$U(g_z(\lambda)) > U(\lambda) \quad \text{for all } \lambda \in (\lambda^*, 1). \quad (21)$$

Let $\lambda \in (1 - \lambda^*, 1)$. Then $g_z^{-1}(\lambda) \in (\lambda^*, 1)$, so that we find with (21)

$$U(\lambda) = U(g_z(g_z^{-1}(\lambda))) > U(g_z^{-1}(\lambda)) \quad \text{for all } \lambda \in (1 - \lambda^*, 1). \quad (22)$$

Moreover, using the symmetry of U in (16), we obtain by (19), or similar arguments as above

$$U(g_z^{-1}(\lambda)) > U(\lambda) \quad \text{for all } \lambda \in (0, 1 - \lambda^*), \quad (23)$$

$$U(\lambda) > U(g_z(\lambda)) \quad \text{for all } \lambda \in (0, \lambda^*). \quad (24)$$

Consider now the functional operator $S_z: \mathcal{F}([0, 1]) \rightarrow \mathcal{F}([0, 1])$ that is defined by

$$(S_z U)(\lambda) = (1 - p)u(\lambda) + p(q_z^0(\lambda) \max(U(\lambda), U(g_z^{-1}(\lambda))) + q_z^1(\lambda) \max(U(\lambda), U(g_z(\lambda))))$$

for all $U \in \mathcal{F}([0, 1])$ and all $\lambda \in [0, 1]$. One can easily check that S_z is a contraction on $\mathcal{F}([0, 1])$, so S_z has a unique fixed point, denoted by \tilde{U} . Let $\tilde{\sigma}$ be a fixed but arbitrary strategy that satisfies (8). By Equations (21)–(24) we have for all $U \in \mathcal{C}_z^*$ that $S_z U = T_{\{\tilde{\sigma}\}}^{\tilde{\sigma}} U \in \mathcal{C}_z^*$, which implies $\tilde{U} \in \mathcal{C}_z^*$. Moreover $U_{\{z\}}^* = T_{\{z\}}^{\sigma^*} U_{\{z\}}^* = S_z U_{\{z\}}^*$ for any σ^* that satisfies (6), so that $U_{\{z\}}^* = \tilde{U} \in \mathcal{C}_z^*$. This implies that a strategy σ satisfies (8) if and only if it satisfies (6), that is, if and only if it is optimal. ■

Proof of Corollary 4.3. It is sufficient to show that $\lambda^*(z) \leq \lambda^*(z')$ if and only if $z \geq z'$. But this follows from the observation that $\lambda^*(z) = \frac{z-1+\sqrt{(1-z)z}}{2z-1}$, which is strictly decreasing

on $(\frac{1}{2}, 1]$. ■

Proof of Theorem 5.1. If $Z = \{1\}$, there is nothing to show. Hence, let $Z \neq \{1\}$, so that $\underline{z} < 1$. Let σ^* be some optimal strategy given Z , and let σ be some optimal strategy given $\{\underline{z}\}$. It is sufficient to show the following four (in-)equalities, as they imply the claim:

$$U_{\{\underline{z}\}}^\sigma(\lambda) = U_{\{\underline{z}\}}^{\sigma^*}(\lambda) \geq U_{\{\underline{z}\}}^{\sigma^*}(\lambda) \geq U_Z^*(\lambda) \geq U_Z^\sigma(\lambda) = U_{\{\underline{z}\}}^\sigma(\lambda). \quad (25)$$

The first equality follows from Theorem 4.1 and the first inequality by the definition of $U_{\{\underline{z}\}}^*$. For the second inequality note that $H[U, Z] \leq H[U, \{\underline{z}\}]$ for every $U \in \mathcal{F}([0, 1])$ by (2). Hence, $T_Z^{\sigma^*}U(\lambda) \leq T_{\{\underline{z}\}}^{\sigma^*}U(\lambda)$ for all $U \in \mathcal{F}([0, 1])$ by the monotonicity of H , so that $(T_Z^{\sigma^*})^n U(\lambda) \leq (T_{\{\underline{z}\}}^{\sigma^*})^n U(\lambda)$ for all $n \in \mathbb{N}$. Thus, $U_Z^*(\lambda) = \lim_n (T_Z^{\sigma^*})^n U(\lambda) \leq \lim_n (T_{\{\underline{z}\}}^{\sigma^*})^n U(\lambda) = U_{\{\underline{z}\}}^{\sigma^*}(\lambda)$. The third inequality is satisfied by the definition of U_Z^* in (5).

For the last equality it is sufficient to show that $U_{\{\underline{z}\}}^\sigma$ is a fixed point of T_Z^σ . So, without loss of generality let $\lambda \leq \frac{1}{2}$, and recall that $U_{\{\underline{z}\}}^\sigma \in \mathcal{C}_{\underline{z}}^*$. By Theorem 4.1, σ is optimal for $U_{\{\underline{z}\}}^\sigma$ in the sense of Equation (6). Thus,

$$\begin{aligned} H[U_{\{\underline{z}\}}^\sigma, B_{Z, \sigma_0(\lambda)}^0(\lambda)] &= U_{\{\underline{z}\}}^\sigma(g_{\underline{z}}^{-1}(\lambda)) \quad \text{and} \\ H[U_{\{\underline{z}\}}^\sigma, B_{Z, \sigma_1(\lambda)}^1(\lambda)] &= \max\{U_{\{\underline{z}\}}^\sigma(g_{\underline{z}}(\lambda)), U_{\{\underline{z}\}}^\sigma(\lambda)\} \end{aligned}$$

by the monotonicity properties of $U_{\{\underline{z}\}}^\sigma$ in Corollary 4.2. Thus, again by Corollary 4.2,

$$H[U_{\{\underline{z}\}}^\sigma, B_{Z, \sigma_0(\lambda)}^0] = U_{\{\underline{z}\}}^\sigma(g_{\underline{z}}^{-1}(\lambda)) \geq \max\{U_{\{\underline{z}\}}^\sigma(g_{\underline{z}}(\lambda)), U_{\{\underline{z}\}}^\sigma(\lambda)\} = H[U_{\{\underline{z}\}}^\sigma, B_{Z, \sigma_1(\lambda)}^1].$$

As $q^0(\lambda)$ is weakly increasing in z (recall that $\lambda \leq \frac{1}{2}$) this means that

$$\begin{aligned} &H[q^0(\lambda) H[U_{\{\underline{z}\}}^\sigma, B_{Z, \sigma_0(\lambda)}^0(\lambda)] + q^1(\lambda) H[U_{\{\underline{z}\}}^\sigma, B_{Z, \sigma_1(\lambda)}^1(\lambda)], Z] \\ &= q_{\underline{z}}^0(\lambda) H[U_{\{\underline{z}\}}^\sigma, B_{Z, \sigma_0(\lambda)}^0(\lambda)] + q_{\underline{z}}^1(\lambda) H[U_{\{\underline{z}\}}^\sigma, B_{Z, \sigma_1(\lambda)}^1(\lambda)] \\ &= H[q^0(\lambda) H[U_{\{\underline{z}\}}^\sigma, B_{Z, \sigma_0(\lambda)}^0(\lambda)] + q^1(\lambda) H[U_{\{\underline{z}\}}^\sigma, B_{Z, \sigma_1(\lambda)}^1(\lambda)], \{\underline{z}\}]. \end{aligned}$$

Hence, it holds that

$$T_Z^\sigma U_{\{\underline{z}\}}^\sigma(\lambda) = T_{\{\underline{z}\}}^\sigma U_{\{\underline{z}\}}^\sigma(\lambda) = U_{\{\underline{z}\}}^\sigma(\lambda).$$

Thus, $U_{\{z\}}^\sigma$ is a fixed point of T_Z^σ as claimed. \blacksquare

Proof of Lemma 5.5. Let T_Z^σ be as defined in (13). Analogously, let $\hat{T}_Z^\sigma: \mathcal{F}([0, 1]) \rightarrow \mathcal{F}([0, 1])$ be defined by

$$\hat{T}_Z^\sigma U(\lambda) = (1 - p)u(\lambda) + pH [q^0(\lambda) H [U, B_{\{\cdot\}, \sigma_0(\lambda)}^0(\lambda)] + q^1(\lambda) H [U, B_{\{\cdot\}, \sigma_1(\lambda)}^1(\lambda)], Z]$$

for all $\lambda \in [0, 1]$. For all $z \in Z$ and any function $U \in \mathcal{F}([0, 1])$,

$$H [U, B_{Z, \sigma_0(\lambda)}^0(\lambda)] \leq H [U, B_{\{z\}, \sigma_0(\lambda)}^0(\lambda)] \quad \text{and} \quad H [U, B_{Z, \sigma_1(\lambda)}^1(\lambda)] \leq H [U, B_{\{z\}, \sigma_1(\lambda)}^1(\lambda)].$$

By the monotonicity of H , $T_Z^\sigma U(\lambda) \leq \hat{T}_Z^\sigma U(\lambda)$ for any function U and all $\lambda \in [0, 1]$. Thus, as in the proof of Theorem 5.1, $U_Z^\sigma(\lambda) = \lim_n (T_Z^\sigma)^n U(\lambda) \leq \lim_n (\hat{T}_Z^\sigma)^n U(\lambda) = \hat{U}_Z^\sigma(\lambda)$. \blacksquare

Proof of Theorem 5.6. If $Z = \{1\}$, there is nothing left to show, so let $Z \neq \{1\}$, so that $z < 1$. Suppose first that σ satisfies (8). In order to show that $U_Z^* = \hat{U}_Z^* = \hat{U}_Z^\sigma$ it is sufficient to show that

$$\hat{U}_Z^*(\lambda) \leq \hat{U}_{\{z\}}^\sigma(\lambda) = U_{\{z\}}^\sigma(\lambda) = U_Z^\sigma(\lambda) \leq \hat{U}_Z^\sigma(\lambda) \leq \hat{U}_Z^*(\lambda) \quad (26)$$

for all $\lambda \in [0, 1]$. As $z \in Z$, H is the Gilboa-Schmeidler aggregator, and σ is optimal given $\{z\}$, one can use the same arguments as in the proofs of Theorem 5.1 and Lemma 5.5 to show that $\hat{U}_Z^{\sigma'}(\lambda) \leq \hat{U}_{\{z\}}^{\sigma'}(\lambda) \leq \hat{U}_{\{z\}}^\sigma(\lambda)$ for all $\sigma' \in \mathcal{S}$. Taking the supremum over all $\sigma' \in \mathcal{S}$ yields the first inequality.

For the singleton set $\{z\}$, there is no difference between sophisticated and naive decision makers, as (3) and (9) coincide in this case. Hence, we arrive at the first equality.

The second equality holds by Theorem 5.1, and the second inequality by Lemma 5.5. The last inequality directly follows from the definition of \hat{U}_Z^* in (10).

It is left to show that $\hat{U}_Z^* > \hat{U}_Z^{\sigma'}$ for any $\sigma' \in \mathcal{S}$ that does not satisfy (8). Assume that there is $\sigma' \in \mathcal{S}$ with $\hat{U}_Z^* = \hat{U}_Z^{\sigma'}$ that does not satisfy (8). Then $\hat{U}_Z^{\sigma'}(\lambda) = \hat{U}_Z^*(\lambda) = U_{\{z\}}^*(\lambda) > U_{\{z\}}^{\sigma'}(\lambda) = \hat{U}_{\{z\}}^{\sigma'}(\lambda) \geq \hat{U}_Z^{\sigma'}(\lambda)$, where the first equality follows from the assumption, the second one from (26), the strict inequality from Theorem 4.1, the third equality holds as for the singleton $\{z\}$ Equations (3) and (9) coincide, and the weak inequality follows from the same arguments as above. But the overall inequality is impossible. \blacksquare

Proof of Theorem 6.1. As $U_Z^\rho(\lambda) \leq \hat{U}_Z^\rho(\lambda)$ by Lemma 5.5 and $U_Z^*(\lambda) = \hat{U}_Z^*(\lambda)$ by Theorem 5.6, it suffices to show that $\hat{U}_Z^\rho(\lambda) < \hat{U}_Z^*(\lambda)$. So, let σ^* be an optimal strategy given Z .

We first prove the claim for $\lambda \in (0, \lambda^*(z)) \cup (1 - \lambda^*(z), 1)$. Without loss of generality let $\lambda \in (0, \lambda^*(z))$. Then

$$\begin{aligned}
& H \left[q^0(\lambda) H \left[\hat{U}_Z^*, B_{\{\cdot\}, \rho_0(\lambda)}^0(\lambda) \right] + q^1(\lambda) H \left[\hat{U}_Z^*, B_{\{\cdot\}, \rho_1(\lambda)}^1(\lambda) \right], Z \right] \\
&= \min_{z \in Z} \left(q_z^0(\lambda) \hat{U}_Z^*(g_z^{-1}(\lambda)) + q_z^1(\lambda) \hat{U}_Z^*(g_z(\lambda)) \right) \\
&\leq q_z^0(\lambda) \hat{U}_Z^*(g_z^{-1}(\lambda)) + q_z^1(\lambda) \hat{U}_Z^*(g_z(\lambda)) \\
&< q_z^0(\lambda) \hat{U}_Z^*(g_z^{-1}(\lambda)) + q_z^1(\lambda) \hat{U}_Z^*(\lambda) \\
&= H \left[q^0(\lambda) H \left[\hat{U}_Z^*, B_{\{\cdot\}, \sigma_0^*(\lambda)}^0(\lambda) \right] + q^1(\lambda) H \left[\hat{U}_Z^*, B_{\{\cdot\}, \sigma_1^*(\lambda)}^1(\lambda) \right], Z \right].
\end{aligned}$$

Thus,

$$\hat{U}_Z^\rho(\lambda) = \hat{T}_Z^\rho \hat{U}_Z^\rho(\lambda) \leq \hat{T}_Z^\rho \hat{U}_Z^*(\lambda) < \hat{T}_Z^{\sigma^*} \hat{U}_Z^*(\lambda) = \hat{U}_Z^*(\lambda),$$

where we used the calculation above in the first and the monotonicity of \hat{T}_Z^ρ in the second inequality.

Suppose next that $\lambda \in [\lambda^*(z), 1 - \lambda^*(z)]$. Then $g_z^{-1}(\lambda) \in (0, \lambda^*(z)]$ and $g_z(\lambda) \in [1 - \lambda^*(z), 1)$ and at least one of them lies in the respective open interval. By the definition of \hat{U}_Z^* it holds that $\hat{U}_Z^*(g_z^{-1}(\lambda)) \geq \hat{U}_Z^\rho(g_z^{-1}(\lambda))$ and $\hat{U}_Z^*(g_z(\lambda)) \geq \hat{U}_Z^\rho(g_z(\lambda))$, and, by the first part of the proof, at least one of these inequalities is strict. Hence,

$$\begin{aligned}
& H \left[q^0(\lambda) H \left[\hat{U}_Z^\rho, B_{\{\cdot\}, \rho_0(\lambda)}^0(\lambda) \right] + q^1(\lambda) H \left[\hat{U}_Z^\rho, B_{\{\cdot\}, \rho_1(\lambda)}^1(\lambda) \right], Z \right] \\
&= \min_{z \in Z} \left(q_z^0(\lambda) \hat{U}_Z^\rho(g_z^{-1}(\lambda)) + q_z^1(\lambda) \hat{U}_Z^\rho(g_z(\lambda)) \right) \\
&\leq q_z^0(\lambda) \hat{U}_Z^\rho(g_z^{-1}(\lambda)) + q_z^1(\lambda) \hat{U}_Z^\rho(g_z(\lambda)) \\
&< q_z^0(\lambda) \hat{U}_Z^*(g_z^{-1}(\lambda)) + q_z^1(\lambda) \hat{U}_Z^*(g_z(\lambda)) \\
&= H \left[q^0(\lambda) H \left[\hat{U}_Z^*, B_{\{\cdot\}, \sigma_0^*(\lambda)}^0(\lambda) \right] + q^1(\lambda) H \left[\hat{U}_Z^*, B_{\{\cdot\}, \sigma_1^*(\lambda)}^1(\lambda) \right], Z \right].
\end{aligned}$$

Thus,

$$\hat{U}_Z^\rho(\lambda) = \hat{T}_{Z, \rho} \hat{U}_Z^\rho(\lambda) < \hat{T}_{Z, \rho}^{\sigma^*} \hat{U}_Z^*(\lambda) = \hat{U}_Z^*(\lambda)$$

as required. ■

Proof of Proposition 6.2. The first part is clear as a strategic decision maker with $\lambda \in I(z)$ will not update her belief at all and choose the option that is optimal given her (correct) prior.

If $\lambda \notin I(z)$, with probability $1-p$ she will obtain no message at all and follow her prior, and with probability p she will obtain a message and update her belief. Since necessarily $\lambda_1 \in I(z)$, she will make a correct choice if and only if the message was correct, which happens with probability z and is independent of p . ■

Proof of Corollary 6.3. For $\lambda \in I(z)$ and $\lambda \in (\lambda^*(z'), 1 - \lambda^*(z'))$ the claim follows immediately from Proposition 6.2. Without loss of generality let $\lambda \in (\lambda^*(z), \lambda^*(z'))$. Since $\lambda > \lambda^*(z) \geq 1 - z$, we have

$$r_z^{\sigma_z^{**}}(\lambda) = pz + (1-p)(1-\lambda) > 1-\lambda = r_{z'}^{\sigma_{z'}^{**}}(\lambda)$$

as required. ■

Proof of Theorem 6.4. We assume without loss of generality that $\lambda \leq \frac{1}{2}$. For any $\ell \in \mathbb{N}$ denote by g_z^ℓ the ℓ -fold application of g_z , and by $g_z^{-\ell}$ the ℓ -fold application of g_z^{-1} . For any history $h^T(\rho) = (m_t, z, \lambda_t)_{t=0}^T$ let $M^1 = |\{t: m_t = 1\}|$ denote the number of generated 1-messages, and $M^0 = T - M^1$ the number of generated 0-messages. Since g_z and g_z^{-1} commute, we have that $\lambda_T = g_z^{M^1 - M^0}(\lambda)$.

Let $G^\ell = \left(g_z^{-\ell} \left(\frac{1}{2} \right), g_z^{-(\ell-1)} \left(\frac{1}{2} \right) \right]$ and observe that a Bayesian decision maker with prior $\lambda \in G^\ell$ who faces history $h^T(\rho)$ will have final belief $\lambda^T > \frac{1}{2}$ if and only if $M^1 - M^0 \geq \ell$, and he will have final belief $\lambda^T = \frac{1}{2}$ if and only if $\lambda = g_z^{-(\ell-1)} \left(\frac{1}{2} \right)$ and $M^1 - M^0 = \ell - 1$.

For any $T \in \mathbb{N}_0$ and $\kappa \in \mathbb{Z}$ denote by $Z[T, \kappa]$ the probability that at least κ correct messages have been observed, conditional on T messages having been generated altogether, that is,

$$W[T, \kappa] = \begin{cases} 0 & \text{if } \kappa > T, \\ \sum_{k=\kappa}^T \binom{T}{k} z^k (1-z)^{T-k} & \text{if } 0 \leq \kappa \leq T, \\ 1 & \text{if } \kappa < 0. \end{cases}$$

The probability that the number of correct messages exceeds the number of incorrect messages by at least ℓ , conditional on T messages being generated, is $W \left[T, \left\lfloor \frac{T+\ell+1}{2} \right\rfloor \right]$. Similarly, the probability that the number of incorrect messages does not exceed the number of correct messages by (strictly) more than $\ell - 1$, conditional on T messages being generated, is $W \left[T, \left\lfloor \frac{T-\ell}{2} \right\rfloor + 1 \right]$. Hence,

$$R_z^\ell(\lambda) = \lambda \sum_{T=0}^{\infty} (1-p)p^T W \left[T, \left\lfloor \frac{T+\ell+1}{2} \right\rfloor \right] + (1-\lambda) \sum_{T=0}^{\infty} (1-p)p^T W \left[T, \left\lfloor \frac{T-\ell}{2} \right\rfloor + 1 \right]$$

denotes the ex-ante probability that either the state is 1 and the number of 1-messages exceeds the number of 0-messages by at least ℓ or the state is 0 and the number of 1-messages does not exceed the number of 0-messages by more than $\ell - 1$. If $\lambda \in \text{int}(G^\ell)$, this is exactly the ex-ante probability that the Bayesian decision maker will make a correct choice. We therefore define $\ell(\lambda)$ to be the unique integer with $\lambda \in G^\ell$ and define $R_z(\lambda) = R_z^{\ell(\lambda)}(\lambda)$. If $\lambda \in \text{int}(G^\ell)$ for some $\ell \in \mathbb{N}$, then $r_z^\rho(\lambda) = R_z(\lambda)$.

The map $R_z: [0, \frac{1}{2}] \rightarrow [0, 1]$ is continuous in the interior of G^ℓ for all $\ell \in \mathbb{N}$. We prove that $R_z^\ell(g_z^{-\ell}(\frac{1}{2})) = R_z^{\ell+1}(g_z^{-\ell}(\frac{1}{2}))$ for all $\ell \in \mathbb{N}$ as this implies that R_z is continuous everywhere. To this end, we first show that

$$g_z^{-\ell}(\frac{1}{2}) z^\ell = (1 - g_z^{-\ell}(\frac{1}{2})) (1 - z)^\ell \quad \text{for all } \ell \in \mathbb{N}_0. \quad (27)$$

Indeed, this is surely true for $\ell = 0$. Suppose (27) is true for some $\ell \in \mathbb{N}_0$. Then

$$\begin{aligned} g_z^{-(\ell+1)}(\frac{1}{2}) z^{\ell+1} &= \frac{g_z^{-\ell}(\frac{1}{2}) z^\ell (1-z)z}{g_z^{-\ell}(\frac{1}{2}) (1-z) + (1 - g_z^{-\ell}(\frac{1}{2})) z} \\ &= \frac{(1 - g_z^{-\ell}(\frac{1}{2})) (1-z)^\ell (1-z)z}{g_z^{-\ell}(\frac{1}{2}) (1-z) + (1 - g_z^{-\ell}(\frac{1}{2})) z} \\ &= \left(1 - \frac{g_z^{-\ell}(\frac{1}{2}) (1-z)}{g_z^{-\ell}(\frac{1}{2}) (1-z) + (1 - g_z^{-\ell}(\frac{1}{2})) z} \right) (1-z)^{\ell+1} \\ &= (1 - g_z^{-(\ell+1)}(\frac{1}{2})) (1-z)^{\ell+1}, \end{aligned}$$

as required. Next observe that

$$W \left[T, \left\lfloor \frac{T+\ell+2}{2} \right\rfloor \right] - W \left[T, \left\lfloor \frac{T+\ell+1}{2} \right\rfloor \right] = \begin{cases} 0 & \text{if } T - \ell \text{ is odd,} \\ -\binom{T}{\frac{T-\ell}{2}} z^{\frac{T-\ell}{2}} (1-z)^{\frac{T+\ell}{2}} & \text{if } T - \ell \text{ is even,} \end{cases}$$

and

$$W \left[T, \left\lfloor \frac{T-\ell-1}{2} \right\rfloor + 1 \right] - W \left[T, \left\lfloor \frac{T-\ell}{2} \right\rfloor + 1 \right] = \begin{cases} 0 & \text{if } T - \ell \text{ is odd,} \\ \binom{T}{\frac{T+\ell}{2}} z^{\frac{T+\ell}{2}} (1-z)^{\frac{T-\ell}{2}} & \text{if } T - \ell \text{ is even.} \end{cases}$$

Since $\binom{T}{\frac{T+\ell}{2}} = \binom{T}{\frac{T-\ell}{2}}$, this implies

$$\begin{aligned} & R_z^{\ell+1} \left(g_z^{-\ell} \left(\frac{1}{2} \right) \right) - R_z^\ell \left(g_z^{-\ell} \left(\frac{1}{2} \right) \right) \\ &= \sum_{T=0}^{\infty} (1-p) p^T \mathbb{1}_{\{T+\ell \text{ even}\}} \binom{T}{\frac{T+\ell}{2}} z^{\frac{T-\ell}{2}} (1-z)^{\frac{T-\ell}{2}} \cdot \left(-g_z^{-\ell} \left(\frac{1}{2} \right) z^\ell + (1 - g_z^{-\ell} \left(\frac{1}{2} \right)) (1-z)^\ell \right) \\ &= 0 \end{aligned}$$

by (27). Hence, R_z is continuous.

Recall that $r_z^\rho(\lambda) = R_z(\lambda)$ for all $\lambda \in (0, \frac{1}{2}]$ with $\lambda \neq g_z^{-\ell}(\frac{1}{2})$ for all $\ell \in \mathbb{N}_0$. We next show that this is also true if $\lambda = g_z^{-\ell}(\frac{1}{2})$ for some $\ell \in \mathbb{N}_0$. To this end note that if the state is 1 and the decision maker observes exactly ℓ more correct than incorrect messages, he will choose either action with equal probability, and the same is true if the state is 0 and he observes exactly ℓ more incorrect than correct messages. Thus, the ex-ante probability of a correct choice is

$$\begin{aligned} & g_z^{-\ell} \left(\frac{1}{2} \right) \left[\sum_{T=0}^{\infty} (1-p) p^T \left(W \left[T, \left\lfloor \frac{T+\ell}{2} \right\rfloor + 1 \right] + \frac{1}{2} \mathbb{1}_{\{T+\ell \text{ even}\}} \binom{T}{\frac{T+\ell}{2}} z^{\frac{T-\ell}{2}} (1-z)^{\frac{T-\ell}{2}} \right) \right] \\ & \quad + (1 - g_z^{-\ell} \left(\frac{1}{2} \right)) \left[\sum_{T=0}^{\infty} (1-p) p^T \left(W \left[T, \left\lfloor \frac{T-\ell}{2} \right\rfloor + 1 \right] + \frac{1}{2} \mathbb{1}_{\{T-\ell \text{ even}\}} \binom{T}{\frac{T-\ell}{2}} z^{\frac{T-\ell}{2}} (1-z)^{\frac{T+\ell}{2}} \right) \right] \\ &= g_z^{-\ell} \left(\frac{1}{2} \right) \left[\sum_{T=0}^{\infty} (1-p) p^T \left(W \left[T, \left\lfloor \frac{T+\ell+1}{2} \right\rfloor \right] - \frac{1}{2} \mathbb{1}_{\{T+\ell \text{ even}\}} \binom{T}{\frac{T+\ell}{2}} z^{\frac{T+\ell}{2}} (1-z)^{\frac{T-\ell}{2}} \right) \right] \\ & \quad + (1 - g_z^{-\ell} \left(\frac{1}{2} \right)) \left[\sum_{T=0}^{\infty} (1-p) p^T \left(W \left[T, \left\lfloor \frac{T-\ell}{2} \right\rfloor + 1 \right] + \frac{1}{2} \mathbb{1}_{\{T-\ell \text{ even}\}} \binom{T}{\frac{T-\ell}{2}} z^{\frac{T-\ell}{2}} (1-z)^{\frac{T+\ell}{2}} \right) \right] \\ &= R_z^\ell \left(g_z^{-\ell} \left(\frac{1}{2} \right) \right) + \frac{1}{2} \sum_{T=0}^{\infty} (1-p) p^T \mathbb{1}_{\{T+\ell \text{ even}\}} \binom{T}{\frac{T+\ell}{2}} z^{\frac{T-\ell}{2}} (1-z)^{\frac{T-\ell}{2}} \left(-g_z^{-\ell} \left(\frac{1}{2} \right) z^\ell + (1 - g_z^{-\ell} \left(\frac{1}{2} \right)) (1-z)^\ell \right) \end{aligned}$$

$$= R_z^\ell (g_z^{-\ell} (\frac{1}{2}))$$

by (27). Hence, $r_z^\rho(\lambda) = R_z(\lambda)$ for all $\lambda \in (0, \frac{1}{2}]$.

We next show that $r_z^\rho(\lambda) > r_z^{\sigma_z^{**}}(\lambda)$ for $\lambda > \lambda^*(z)$. Since $\lambda > \lambda^*(z) > 1 - z = g_z^{-1}(\frac{1}{2})$ it holds that $\ell(\lambda) = 1$. Define for all $T \in \mathbb{N}_0$

$$f_z^T(\lambda) = \lambda W [T, \lfloor \frac{T+2}{2} \rfloor] + (1 - \lambda) W [T, \lfloor \frac{T+1}{2} \rfloor] \quad (28)$$

and note that $r_z^\rho(\lambda) = R_z(\lambda) = R_z^1(\lambda) = \sum_{T=0}^{\infty} (1-p)p^T f_z^T(\lambda)$. It is straightforward that $f_z^0(\lambda) = 1 - \lambda$, $f_z^1(\lambda) = z$, and $f_z^2(\lambda) = 2(1 - \lambda)(1 - z)z + z^2 \geq z$. We show that $f_z^{T+1}(\lambda) \geq f_z^T(\lambda) > z$ for all $T \geq 3$. To this end, observe first that

$$f_z^T(\lambda) = \begin{cases} W [T, \frac{T+1}{2}] & \text{if } T \text{ is odd,} \\ W [T, \frac{T}{2}] - \lambda \binom{T}{\frac{T}{2}} z^{\frac{T}{2}} (1 - z)^{\frac{T}{2}} & \text{if } T \text{ is even.} \end{cases}$$

Next, one finds that for all $1 \leq \kappa \leq T$

$$\begin{aligned} (1 - z)W [T, \kappa] + zW [T, \kappa - 1] &= \sum_{k=\kappa}^T \binom{T}{k} z^k (1 - z)^{T-k+1} + \sum_{k=\kappa-1}^T \binom{T}{k} z^{k+1} (1 - z)^{T-k} \\ &= \sum_{k=\kappa}^T \binom{T}{k} z^k (1 - z)^{T-k+1} + \sum_{k=\kappa}^{T+1} \binom{T}{k-1} z^k (1 - z)^{T-k+1} \\ &= \sum_{k=\kappa}^T \binom{T+1}{k} z^k (1 - z)^{T-k+1} + \binom{T}{T} z^{T+1} \\ &= W [T+1, \kappa]. \end{aligned}$$

Suppose first that T is even. Then

$$\begin{aligned} f_z^{T+1}(\lambda) &= W [T+1, \frac{T+2}{2}] \\ &= (1 - z)W [T, \frac{T+2}{2}] + zW [T, \frac{T}{2}] \\ &= (1 - z) \left(W [T, \frac{T}{2}] - \binom{T}{\frac{T}{2}} z^{\frac{T}{2}} (1 - z)^{\frac{T}{2}} \right) + zW [T, \frac{T}{2}] \\ &= W [T, \frac{T}{2}] - (1 - z) \binom{T}{\frac{T}{2}} z^{\frac{T}{2}} (1 - z)^{\frac{T}{2}} \end{aligned}$$

$$\begin{aligned}
&> W\left[T, \frac{T}{2}\right] - \lambda \binom{T}{\frac{T}{2}} z^{\frac{T}{2}} (1-z)^{\frac{T}{2}} \\
&= f_z^T(\lambda)
\end{aligned}$$

where the inequality follows from $\lambda > 1 - z$. Thus, $f_z^3(\lambda) > f_z^2(\lambda) = z$.

Suppose next that T is odd. In this case

$$\begin{aligned}
f_z^{T+1}(\lambda) &= W\left[T+1, \frac{T+1}{2}\right] - \lambda \binom{T+1}{\frac{T+1}{2}} z^{\frac{T+1}{2}} (1-z)^{\frac{T+1}{2}} \\
&= (1-z)W\left[T, \frac{T+1}{2}\right] + zW\left[T, \frac{T-1}{2}\right] - \lambda \binom{T+1}{\frac{T+1}{2}} z^{\frac{T+1}{2}} (1-z)^{\frac{T+1}{2}} \\
&= (1-z)W\left[T, \frac{T+1}{2}\right] + z\left(W\left[T, \frac{T+1}{2}\right] + \binom{T}{\frac{T-1}{2}} z^{\frac{T-1}{2}} (1-z)^{\frac{T+1}{2}}\right) \\
&\quad - \lambda \binom{T+1}{\frac{T+1}{2}} z^{\frac{T+1}{2}} (1-z)^{\frac{T+1}{2}} \\
&= W\left[T, \frac{T+1}{2}\right] + \binom{T}{\frac{T+1}{2}} z^{\frac{T+1}{2}} (1-z)^{\frac{T+1}{2}} - 2\lambda \binom{T}{\frac{T+1}{2}} z^{\frac{T+1}{2}} (1-z)^{\frac{T+1}{2}} \\
&= f_z^T(\lambda) + (1-2\lambda) \binom{T}{\frac{T+1}{2}} z^{\frac{T+1}{2}} (1-z)^{\frac{T+1}{2}} \\
&\geq f_z^T(\lambda),
\end{aligned}$$

since $\lambda \leq \frac{1}{2}$. Thus,

$$\begin{aligned}
r_z^\rho(\lambda) &= \sum_{T=0}^{\infty} (1-p)p^T f_z^T(\lambda) = (1-p)(1-\lambda) + \sum_{T=1}^{\infty} (1-p)p^T f_z^T(\lambda) \\
&> (1-p)(1-\lambda) + pz = r_z^{\sigma_z^{**}}(\lambda),
\end{aligned}$$

which completes the proof for $\lambda \in (\lambda^*(z), \frac{1}{2}]$.

Let $\lambda \leq \lambda^*(z)$. We can rewrite $R_z^\ell(\lambda)$ as

$$R_z^\ell(\lambda) = \sum_{T=0}^{\infty} (1-p)p^T W\left[T, \lfloor \frac{T-\ell}{2} \rfloor + 1\right] - \lambda \sum_{T=0}^{\infty} (1-p)p^T (W\left[T, \lfloor \frac{T-\ell}{2} \rfloor + 1\right] - W\left[T, \lceil \frac{T+\ell}{2} \rceil\right]).$$

Since $W[T, \kappa]$ is decreasing in κ and uniformly bounded for all $T \in \mathbb{N}$ and $\kappa \in \mathbb{Z}$, this shows that each R_z^ℓ is an affine linear and decreasing function in λ . Moreover, the absolute value of the slope grows in ℓ , that is, the larger ℓ , the steeper R_z^ℓ . As $R_z^\ell(g_z^{-\ell}(\frac{1}{2})) =$

$R_z^{\ell+1}(g_z^{-\ell}(\frac{1}{2}))$, it follows that $R_z^\ell(\lambda) \geq R_z^{\ell+1}(\lambda)$ if and only if $\lambda \geq g_z^{-\ell}(\frac{1}{2})$. Thus, $R_z(\lambda) = R_z^{\ell(\lambda)}(\lambda) = \max_{\ell \geq 1} R_z^\ell(\lambda)$. In particular $R_z^{\ell(\lambda)}(\lambda) \geq R_z^{\ell(\lambda)+1}(\lambda) > R_z^{\ell(\lambda)+2}(\lambda) \geq R_z^{\ell(\lambda)+n}(\lambda)$ for all $n \geq 2$. As $W[T, \kappa]$ is uniformly bounded for all $T \in \mathbb{N}$ and $\kappa \in \mathbb{Z}$

$$\begin{aligned} \lim_{\ell \rightarrow \infty} R_z^\ell(\lambda) &= \sum_{T=0}^{\infty} (1-p)p^T \lim_{\ell \rightarrow \infty} W\left[T, \left\lfloor \frac{T-\ell}{2} \right\rfloor + 1\right] \\ &\quad - \lambda \sum_{T=0}^{\infty} (1-p)p^T \left(\lim_{\ell \rightarrow \infty} W\left[T, \left\lfloor \frac{T-\ell}{2} \right\rfloor + 1\right] - \lim_{\ell \rightarrow \infty} W\left[T, \left\lceil \frac{T+\ell}{2} \right\rceil \right] \right) \\ &= 1 - \lambda. \end{aligned}$$

Hence, $R_z^{\ell(\lambda)+n}(\lambda)$ converges monotonically to $1 - \lambda$ for $n \rightarrow \infty$, so that $r_z^\rho(\lambda) = R_z^{\ell(\lambda)}(\lambda) > 1 - \lambda = r_z^{\sigma_z^{**}}(\lambda)$ as claimed. \blacksquare

Proof of Lemma 8.1. Using the definitions, we observe that

$$\begin{aligned} &H[q_z^0(\lambda)H[U, G_{F, \sigma_0(\lambda)}^0(\lambda)] + q_z^1(\lambda)H[U, G_{F, \sigma_1(\lambda)}^1(\lambda)], F] \\ &= \int_Z q_z^0(\lambda) \left(\int_{[0,1]} U(\lambda') G_{F, \sigma_0(\lambda)}^0(\lambda)(d\lambda') \right) + q_z^1(\lambda) \left(\int_{[0,1]} U(\lambda') G_{F, \sigma_1(\lambda)}^1(\lambda)(d\lambda') \right) F(dz) \\ &= \int_Z q_z^0(\lambda) \left(\int_Z U(\psi_{\lambda, \sigma_0(\lambda)}(z')) F_\lambda^0(dz') \right) + q_z^1(\lambda) \left(\int_Z U(\varphi_{\lambda, \sigma_1(\lambda)}(z')) F_\lambda^1(dz') \right) F(dz) \\ &= \int_Z q_z^0(\lambda) \left(\int_Z U(\psi_{\lambda, \sigma_0(\lambda)}(z')) \frac{q_{z'}^0(\lambda)}{\int_Z q_{z''}^0(\lambda) F(dz'')} F(dz') \right) \\ &\quad + q_z^1(\lambda) \left(\int_Z U(\varphi_{\lambda, \sigma_1(\lambda)}(z')) \frac{q_{z'}^1(\lambda)}{\int_Z q_{z''}^1(\lambda) F(dz'')} F(dz') \right) F(dz) \\ &= \int_Z q_z^0(\lambda) \left(\frac{\int_Z U(\psi_{\lambda, \sigma_0(\lambda)}(z')) q_{z'}^0(\lambda) F(dz')}{\int_Z q_{z''}^0(\lambda) F(dz'')} \right) F(dz) \\ &\quad + \int_Z q_z^1(\lambda) \left(\frac{\int_Z U(\varphi_{\lambda, \sigma_1(\lambda)}(z')) q_{z'}^1(\lambda) F(dz')}{\int_Z q_{z''}^1(\lambda) F(dz'')} \right) F(dz) \\ &= \left(\int_Z q_z^0(\lambda) F(dz) \right) \left(\frac{\int_Z U(\psi_{\lambda, \sigma_0(\lambda)}(z')) q_{z'}^0(\lambda) F(dz')}{\int_Z q_{z''}^0(\lambda) F(dz'')} \right) \\ &\quad + \left(\int_Z q_z^1(\lambda) F(dz) \right) \left(\frac{\int_Z U(\varphi_{\lambda, \sigma_1(\lambda)}(z')) q_{z'}^1(\lambda) F(dz')}{\int_Z q_{z''}^1(\lambda) F(dz'')} \right) \\ &= \int_Z U(\psi_{\lambda, \sigma_0(\lambda)}(z')) q_{z'}^0(\lambda) F(dz') + \int_Z U(\varphi_{\lambda, \sigma_1(\lambda)}(z')) q_{z'}^1(\lambda) F(dz') \end{aligned}$$

$$\begin{aligned}
&= \int_Z q_z^0(\lambda) U(\psi_{\lambda, \sigma_0(\lambda)}(z)) + q_z^1(\lambda) U(\varphi_{\lambda, \sigma_1(\lambda)}(z)) F(dz) \\
&= \int_Z q_z^0(\lambda) \left(\int_{[0,1]} U(\lambda') G_{\delta_z, \sigma_0(\lambda)}^0(\lambda) (d\lambda') \right) + q_z^1(\lambda) \left(\int_{[0,1]} U(\lambda') G_{\delta_z, \sigma_1(\lambda)}^1(\lambda) (d\lambda') \right) F(dz) \\
&= H[q^0(\lambda) H[U, G_{\delta, \sigma_0(\lambda)}^0(\lambda)] + q^1(\lambda) H[U, G_{\delta, \sigma_1(\lambda)}^1(\lambda)], F].
\end{aligned}$$

This is clearly equivalent to the claim. ■