Learning in a Small/Big World*

Benson Tsz Kin Leung†

April 2, 2024

Abstract

Complexity and limited ability affect how we learn and make decisions under uncertainty. Using finite automata to model belief formation, this paper studies the characteristics of optimal learning behavior in small and big worlds, where the complexity of the environment is low and high, respectively, relative to the cognitive ability of the decision maker. Optimal behavior is well approximated by the Bayesian benchmark in very small worlds but is more different as the worlds get bigger. In addition, in big worlds, the optimal learning behavior could exhibit a wide range of well-documented non-Bayesian learning behavior, including heuristics, correlation neglect, persistent over-confidence, inattentive learning, and other behaviors of model simplification or misspecification. These results establish a clear and testable relationship among the prominence of non-Bayesian learning behavior, complexity, and cognitive ability.

Keywords: Learning, Bounded Memory, Bayesian, Complexity, Cognitive Ability
JEL codes: D83, D91

*I am grateful to Tilman Börgers, Jacques Crémer, Matthew Elliott, Renato Gomes, Philippe Jehiel, Hamid Sabourian, Mikhael Safronov, Larry Samuelson, Tak-Yuen Wong, anonymous referees, and the audiences at various seminars and conferences for their insightful discussions and comments.

†Hong Kong Baptist University. Email: btkleung@hkbu.edu.hk, ltkbenson@gmail.com.

1 Introduction

Savage (1972) argues that Bayesian decision theory applies only to small world problems but not big world problems, where the latter refers to scenarios where it is difficult to form prior beliefs on states and signal structures or construct the state space. This paper offers an alternative distinction between "small" and "big" worlds based on the complexity of the inference problem relative to the cognitive ability of individuals. "Small" worlds refers to inference problems that are simple relative to the cognitive ability of individuals. Therefore, by comparing learning behaviors in the "small" and "big" worlds, this paper sheds light on the heterogeneity of learning behaviors across decision problems and individuals. Analyzing a theoretical model with finite automata, I show that a wide range of ignorance behavior, including the use of heuristics (Kahneman et al. (1982)), correlation neglect (Enke and Zimmermann (2019)), persistent over-confidence (Hoffman and Burks (2017)), inattentive learning (Graeber (2019)), arises as optimal learning behaviors in face of complexity in big worlds, but not in small worlds. Thus, these "biased" learning behaviors are more prominent in more complex problems and among individuals with lower cognitive ability.

More specifically, I consider a decision maker (DM) who tries to learn the true state of the world from a finite state space, where the number of possible states N measures the complexity of the inference problem. In each period $t = 1, \cdots, \infty$, the DM guesses the true state of the world and gets a higher utility if he is correct. After making a guess, he receives a signal and updates his belief. To model limited cognitive ability, I assume that the DM's belief is confined to an M sized automaton that captures limited cognitive ability, as in the seminal work of Hellman and Cover (1970) and Wilson (2014).[1] In period 0, he starts in one of the M memory states, makes his guess given his memory state based on a decision rule, receives a signal, transits to another based on a transition rule, and finally enters the next period. In contrast to the Bayesian model, the DM has a coarser idea of the likelihood of different states of the world, and the coarseness decreases in M. Thus M measures the cognitive ability of the DM. I define small worlds as cases where $\frac{N}{M}$ is small, otherwise the decision problem is a big world: Whether a problem is a small or big world depends on the relative complexity of the world with respect to the individual's cognitive ability.[2]

To shed light on how complexity affects learning, I compare the characteristics of the optimal updating mechanisms that maximize the asymptotic utility of the DM in small and big worlds. The results are summarized in Table 1. First, I analyze how the individual's decisions differ from the Bayesian benchmark in small and big worlds. This sheds light on under what circumstances the Bayesian model serves as a good approximation of decision-making under uncertainty. I show that in small worlds, asymptotic behavior is close to

[1]See also Compte and Postlewaite (2012), Monte and Said (2014), Basu and Chatterjee (2015), Chauvin (2019) Chatterjee and Sabourian (2020) in the economic literature that model belief updating and the aversion of complexity with finite automata. See also Oprea (2020) and Banovetz and Ryan (2020) for experimental evidence.

[2]Note that if the individual tracks his belief not with a finite automaton but with a real number statistic, the cardinality of the belief statistics is much larger than N, and the model collapses to a Bayesian model.

	Small Worlds: low complexity relative to cognitive ability $\frac{N}{M}$	Big Worlds: high complexity relative to cognitive ability $\frac{N}{M}$
Is behavior close to Bayesian?	Yes	No
Could ignorance in learning be "optimal"?	No	Yes
Could disagreement be persistent?	No	Yes

Table 1: Differences in learning behaviors in small/big worlds

Bayesian. In particular, the DM almost always makes the same guess as a Bayesian individual as $\frac{N}{M} \to 0$. In contrast, when the world is bigger, i.e., when $\frac{N}{M}$ increases, the DM makes more mistakes and his behavior becomes more different from Bayesian.

The second, and main, result of this paper shows that in big worlds, as the DM faces a trade-off in allocating his scarce cognitive resources, i.e., the M memory states, it could be optimal to ignore some states and focus learning on a subset of states. In contrast, such "ignorant" behavior is never optimal in small worlds. This shows a relationship between complexity and ignorance in learning, which nests different well-documented learning biases, including heuristics, correlation neglect, persistent over-confidence, inattentive learning, and other model simplification and misspecification behaviors.

To see how "ignorance" captures different biases, consider the example of persistent over-confidence (Hoffman and Burks (2017), Heidhues et al. (2018)). Suppose that the state of the world comprises the DM's and his teammate's ability, where both could be high or low, and the DM observes team performances as signals. In the current setting, persistent over-confidence occurs when the DM never guesses the states where his ability is low, behaves as if he always believes he has high ability, and only updates his belief about his teammate's ability, even after observing a sequence of bad team performance. Similarly, for correlation neglect (Enke and Zimmermann (2019)), consider that the state of the world comprises a strong or weak stock market and positive, no, or negative correlation among data. Correlation neglect is captured in the current setting as the individual ignores the possible correlation among data and only updates his belief about the stock market as if he always believes there is no correlation.

This paper shows that such ignorance behavior is optimal only when relative complexity is large (big worlds), especially when the ignored state is a priori unlikely, or when information supporting that state is weak. For example, if the DM is confident about his ability, he focuses on learning his teammate's ability; if it is difficult to distinguish state with and without correlation, the DM ignores the possibility that data are correlated. Importantly, I also show that even if the states and information structures are symmetric, ignorance is optimal in environments where it is difficult to learn, e.g., when signals are noisy or when the state space is large. It echoes that complexity, or difficulty in learning, drives ignorance, even when there are no a priori reasons to ignore any specific states of the world.

Last, I analyze whether disagreements are persistent in small and big worlds. As asymptotic behaviors are close to the Bayesian model in small worlds, intuitively individuals would always eventually agree with each other and make the same guesses. In contrast, in big worlds, because individuals with different prior beliefs and/or cognitive ability adopt different optimal learning mechanisms and could ignore different states, they could disagree with each other with probability 1 even after receiving the same infinite sequence of public information. For example, after observing a large sequence of bad team performance, two individuals with persistent over-confidence would disagree on the assessment of their abilities: they ignore the states where their ability is low and attribute the bad performance to the other person. Moreover, I show a novel driving force of disagreement: even when two individuals have the same prior beliefs and observe the same infinite sequence of public signals with no uncertainties in signal structures, they could eventually disagree with each other when they have different levels of cognitive ability M.

This paper is organized as follows. In the next section, I briefly discuss how this paper relates to the literature. Section 3 presents the model. I analyze the optimal learning behavior in small and big worlds in Sections 4. In Section 5, I conclude by presenting a discussion of the results. The proofs and omitted results are presented in the Appendix, and the extensions are presented in the online Appendix.[3]

2 Literature

In this section, I discuss the existing literature and the contribution of this paper.

First, this paper is obviously related to the literature using finite automata to model learning with aversion to complexity (Hellman and Cover (1970), Börgers and Morales (2004), Compte and Postlewaite (2012), Wilson (2014), Chatterjee and Hu (2021), etc.).[4] While the literature discusses the impact of complexity on learning, they do not analyze how *different levels* of complexity and cognitive ability affect learning and the prominence of learning biases. Moreover, the studies above focus on binary state space because of technical difficulties. While this paper does not "fully" characterize the optimal automaton, the results shed light on its characteristics when $N > 2$. In particular, the results about ignorance suggest that for larger state space, unlike when $N = 2$, ignorance is a crucial feature even when the states are a priori the same.

Sims (2003) and Matějka and McKay (2015) study the implication of rational inattention and show that it explains sticky prices in the market and micro-founds the multinomial logit choice model, Steiner and Stewart (2016) shows that an optimal response to noises in perceiving the details of lotteries leads to probability weighting in prospect theory (Kahneman and Tversky (1979)), and Jehiel and Steiner (2020) and Leung (2020) show that a capacity

[3] The online Appendix could be found on `https://sites.google.com/site/ltkbenson/research`.

[4] Also see Chatterjee and Sabourian (2020) for a review, and Oprea (2020) and Banovetz and Ryan (2020) for empirical evidence.

constraint on the number of signals that individuals could update their belief with drives confirmation bias and other biases in belief formation. This paper also contributes to a growing literature that explains behavioral anomalies as optimal/efficient strategies in light of limited cognitive ability. Steiner and Stewart (2016) shows that an optimal response to noises in perceiving the details of lotteries leads to probability weighting in prospect theory, and Jehiel and Steiner (2020) and Leung (2020) show that a capacity constraint on the number of signals that individuals could update their belief with drives confirmation bias. In contrast, this paper explains a larger class of biases under the same framework and illustrates a relationship between their prominence and the level of complexity.

Similar to Section 4.2 in this paper, Caplin et al. (2019) present conditions where the DM would ignore some actions in a rational inattention setting. Unlike this paper, they only present conditions depending on the prior belief and the utility matrix, but not the level of complexity nor the informativeness of signals. Moreover, in contrast to Caplin et al. (2019), in this paper, I show that it could be optimal to ignore some states even in symmetric environments where prior belief is uniform and the utility matrix is symmetric.

Last, this paper's results on asymptotic disagreement contribute to the large literature that explains the phenomenon. In the existing literature, asymptotic disagreement is driven by differences in signal distributions across states or differences in learning mechanisms (Morris (1994), Mailath and Samuelson (2020), Gilboa, Samuelson and Schmeidler (2020)), the lack of identification or uncertainty in signal distributions (Acemoglu, Chernozhukov and Yildiz (2016)), confirmation bias (Rabin and Schrag (1999)), or model misspecification (Freedman (1963, 1965), Berk (1966)). Differently, this paper looks into the connection between limited ability and disagreement, and shows when asymptotic disagreement could arise and when it will not occur, depending on the relative complexity of the inference problem. Moreover, I show a novel machanism that disagreement could arise solely because of differences in cognitive abilities.

3 Model

I consider a world with N possible true states, i.e., $\omega \in \Omega = \{1, 2, \cdots, N\}$, and a decision-maker (DM), wherein each period $t = 1, \cdots, \infty$, the DM tries to guess what the true state is. Formally, in each period t, the DM takes an action $a_t \in \mathscr{A} = \Omega$ and gets utility $u(a, \omega) \in \mathscr{R}$ where $a = \omega$ is the unique maximizer of $u(a, \omega)$.[5] The DM does not observe his utility after taking action; thus, $u(a, \omega)$ is best interpreted as an intrinsic utility of being correct. Otherwise, the problem becomes trivial as the DM will experiment and learn perfectly the true state after observing the utility. The prior belief of the DM is $(p^\omega)_{\omega=1}^N$ where $\sum_{\omega=1}^N p^\omega = 1$ and $p^\omega > 0$ for all $\omega \in \Omega$.

[5]This assumption rules out "safe" actions that are not maximizers in any state but yield good payoffs in multiple states. In the next section, I discuss how complexity affects the incentive of choosing these "safe" actions.

In each period after taking an action, the DM receives a signal $s_t \in S$ that is independently drawn across different periods from a continuous distribution with p.d.f. f^ω in state ω.[6,7] I assume that no signal perfectly rules out any states of the world: there exists $\varsigma > 0$ such that

$$\inf_{s \in S} \frac{f^\omega(s)}{f^{\omega'}(s)} > \varsigma \text{ for all } \omega,\, \omega' \in \Omega. \tag{1}$$

Without loss of generality, no pairs of signal structures are the same, i.e., there are no ω and $\omega' \neq \omega$ such that $f^\omega(s) = f^{\omega'}(s)$ for (almost) all $s \in S$. This implies that states are identifiable, and a Bayesian learns the true state perfectly as $t \to \infty$. In contrast, I focus on the bounded memory setting I now describe.

The DM is subject to a memory constraint such that he can only update his belief using an M memory states automaton. In each period, his belief is represented by a memory state $m_t \in \{1, 2, \cdots, M\}$. An updating mechanism specifies an initial state $m_1 \in \triangle M$, a transition function between the M memory states given a signal $s \in S$, which is denoted as $\mathscr{T} : M \times S \to \triangle M$, and a decision rule $d : M \to A$.[8] The set of memory states where the DM chooses action ω is denoted as M^ω. In each period t, as illustrated in Figure 1, the DM starts with some memory state m_t, take action $a_t \sim d(m_t)$, receives a signal s_t, and transit to memory state $m_{t+1} \sim \mathscr{T}(m_t, s_t)$.[9] This paper analyzes the asymptotic learning of the DM, i.e., the DM aims to choose an updating mechanism at period 0 that maximizes his expected long run per-period utility:[10]

$$\lim_{T \to \infty} E_{m_1, \mathscr{T}, d} \left[\frac{1}{T} \sum_{t=1}^{T} u(a_t, \omega) \right].$$

Given state $\omega \in \Omega$, the sequence m_t, together with some specified initial memory state m_1, forms a Markov chain. Denote μ_m^ω as the long-run proportion of time that the DM is in memory state m when the true state of the world is ω, and $\mathbf{Q}^\omega$ as the matrix of transition probabilities, i.e., $\mathbf{Q}^\omega = [\int_s \Pr\{\mathscr{T}(m, s) = m'\} f^\omega(s)\, ds]_{mm'}$. By the Birkhoff–Khinchin theorem, the distribution $\boldsymbol{\mu}^\omega = (\mu_1^\omega, \mu_2^\omega, \cdots, \mu_M^\omega)^T$ solves the following system of equations:

$$\boldsymbol{\mu}^\omega = (\boldsymbol{\mu}^\omega)^T \mathbf{Q}^\omega, \tag{2}$$

By the Brouwer fixed-point theorem, a solution always exists. Moreover, when there are multiple solutions, it implies that the Markov Chain is reducible, and the long-run distribution

[6] The order, i.e., whether the DM receives a signal before or after taking an action in each period, does not affect the result.

[7] To ease exposition, I assume that signals follow a continuous distribution, but the results hold with more general probability measures.

[8] Note that without loss of generality, I restrict attention to deterministic decision rules unless stated otherwise.

[9] Switching between multiple M memory state automatons requires more than M memory states, and the current setting allows switching between smaller automatons, as illustrated in Online Appendix A.

[10] An alternative is to maximize the discounted sum of utility as in Wilson (2014). As shown in Online Appendix B, the results in this paper hold qualitatively when the discount factor is close to 1.

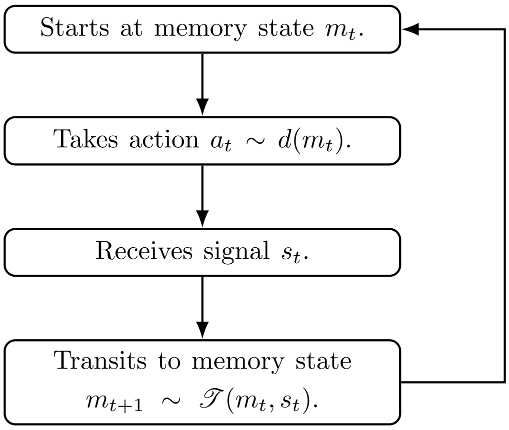

Figure 1: Timeline at period t given an updating mechanism $(\mathscr{T}, d)$.

is uniquely pinned down by the initial memory state.

The asymptotic utility, or the long-run per-period utility, of an updating mechanism $(m_1, \mathscr{T}, d)$ is equal to:

$$U(m_1, \mathscr{T}, d) = \sum_{\omega=1}^{N} \left[p^{\omega} \left(\sum_{m=1}^{M} u\left(d(m), \omega\right) \mu_m^{\omega} \right) \right] \tag{3}$$

and the asymptotic utility loss is equal to:

$$L(m_1, \mathscr{T}, d) = \sum_{\omega=1}^{N} \left[p^{\omega} u(\omega, \omega) \right] - U(m_1, \mathscr{T}, d). \tag{4}$$

The DM maximizes the asymptotic utility or, equivalently, minimizes the asymptotic utility loss. I mostly refer the optimal design of the updating mechanism as the minimization of L. In general, with similar arguments in Hellman and Cover (1970), an optimal mechanism may not exist. Therefore, the rest of the paper focuses on ϵ-optimal updating mechanisms that are defined as follows. Define

$$L_M^* \equiv \inf_{m_1, \mathscr{T}, d} L(m_1, \mathscr{T}, d).$$

An updating mechanism $(m_1, \mathscr{T}, d)$ is ϵ-optimal if and only if $L(m_1, \mathscr{T}, d) \leq L_M^* + \epsilon$. Throughout the paper, I focus on the more interesting case where $L_M^* < \min_a \sum_{\omega=1}^{N} p^{\omega} \left[u(\omega, \omega) - u(a, \omega) \right]$, such that learning strictly improves utility.

In the next section, I compare the characteristics of the ϵ-optimal updating mechanisms in small and big worlds. Roughly speaking, N represents how complicated the world is, and M represents the cognitive resources/ability of the DM. This gives a natural definition of small and big worlds based on relative (or perceived) complexity: an inference problem is a small world when the $\frac{N}{M}$ is small, and is a big world otherwise. Throughout the paper, unless stated otherwise, I focus on the more interesting case that $M \geq N$.

4 Results

4.1 Is behavior close to Bayesian?

In the current setup, a Bayesian individual (almost) perfectly learns the true state of the world asymptotically and achieve asymptotic utility loss close to 0. Thus, we could interpret L_M^* as the distance between the DM's behavior and that of a Bayesian individual. The Proposition presents the formal results and the paragraph after offers interpretations. For ease of exposition, I denote $\boldsymbol{u}^\omega \equiv (u^\omega(a,\omega))_{a=1}^N$.

Proposition 1. *We have the following results regarding L_M^*:*

(i) L_M^* strictly decreases in M;

(ii) For each N and $(\boldsymbol{u}^\omega, p^\omega, f^\omega)_{\omega=1}^N$, there exists some constant $r < 1$ and $K > 0$ such that $L_M^* < Kr^{\lfloor \frac{M-1}{N} \rfloor}$ (also implies $\lim_{M\to\infty} L_M^* = 0$);

(iii) For each N and $(f^\omega)_{\omega=1}^N$, there exists a sequence of updating mechanism $(m_1, \mathscr{T}_M, d_M)$ such that $\lim_{M\to\infty} L(m_1, \mathscr{T}_M, d_M) = 0$ for all $(\boldsymbol{u}^\omega, p^\omega)_{\omega=1}^N$.

The behavioral implications of Proposition 1 are as follows: (i) shows that the DM's behavior gets closer to that of a Bayesian individual as M increases, or equivalently as $\frac{N}{M}$ decreases, i.e., as the world gets smaller.[11] (ii) shows that as $\frac{N}{M}$ converges to 0, the DM's asymptotic decisions are well-approximated by Bayesian, and shows a higher relative complexity of the world decreases the convergence rate. (iii) demonstrates the robustness of perfect learning in very small worlds where $\frac{N}{M}$ is close to 0, as no knowledge of prior or the utility matrix is needed.[12]

Before I describe the the proof, it is important to first define confirmatory signals for each state ω, especially because the model goes beyond the binary state setting. When $N = 2$, the confirmatory signals for state 1 are $S_1 = \{s \in S : f^1(s) > f^2(s)\}$. The set of confirmatory signals for both states defined in this way is non-empty, and because of the binary nature, it is also more likely that the DM receives confirmatory signals for the correct state than for the wrong state, i.e., $\int_{s\in S_1} f^1(s)\, ds > \int_{s\in S_2} f^1(s)\, ds$. However, these observations do not generalize to $N > 2$. For example, consider the following signal structure with $N = 3$ and $S = \{s_1, s_2, s_3\}$:

$$f^1(s_1) = 0.5, f^1(s_2) = 0.4, f^1(s_3) = 0.1;$$
$$f^2(s_1) = 0.4, f^2(s_2) = 0.2, f^2(s_3) = 0.4;$$
$$f^3(s_1) = 0.1, f^3(s_2) = 0.4, f^3(s_3) = 0.5.$$

[11] In the main text, N is fixed, and thus a world is small enough, i.e., $\frac{N}{M}$ is small enough, if and only if M is large enough. In online Appendix C, I analyze the ϵ-optimal updating mechanism of a sequence of inference problems where both $N, M \to \infty$, and show that the behavioral implications depend on the limit of $\frac{N}{M}$ instead of M. In particular, the results of small worlds hold qualitatively when the limit of $\frac{N}{M}$ is 0 and the results of big worlds hold qualitatively when the limit of $\frac{N}{M}$ is strictly greater than 0.

[12] I thank an anonymous referee for pointing this out.

If I define signals supporting state ω as $S_\omega = \{s \in S : \omega = \arg\max_{\omega'} f^{\omega'}(s)\}$, as analogue to $N = 2$, there won't be any signals that support state 2. One could also show that no deterministic definitions of confirmatory signals will guarantee that the DM receives signals supporting the correct state more likely than signals supporting a wrong state in every state of the world. Instead, I stochastically label signals as confirmatory signals for each state. A signal s is labeled as a confirmatory signal for state ω with probability proportional to $f^\omega(s)$, with appropriate normalization such that each signal is labeled as a confirmatory signal for one of the state with probability less than 1, and is labeled as a confirmatory signal for no state with complementary probability. In this example, s_2 is labeled as a confirmatory signal for state 1 with probability 0.4, a confirmatory signal for state 2 with probability 0.2 and a confirmatory signal for state 3 with probability 0.4. s_1 and s_3 are labeled analogously. Thus, in state 2, the DM receives a signal supporting state 2 with probability $0.4^2+0.2^2+0.4^2 = 0.36$, a signal supporting state 1 with probability $0.5*0.4+0.4*0.2+0.1*0.4 = 0.32$, and a signal supporting 3 with probability $0.1*0.4+0.4*0.2+0.5*0.4 = 0.32$. Similarly, in state 1, the DM receives a confirmatory signal for state 1 more likely than a signal supporting state 2 (or 3), and in state 3, the DM receives a confirmatory signal for state 3 more likely than a signal supporting state 1 (or 2). In all ω, the DM receives a signal supporting the correct state more likely than a signal supporting a wrong state.

Now using the definition of confirmatory signals, I briefly describe the proof of (i) of Proposition 1. Consider an ϵ-optimal mechanism $(m_1, \mathscr{T}_M, d_M)$ at memory size M in which $M^1, M^2 \neq \emptyset$ and $M^1 \cup M^2 = M$, i.e., the DM either chooses action 1 or action 2, and suppose the memory size increases to $M+1$. The following construction strictly improves the asymptotic utility of $(m_1, \mathscr{T}_M, d_M)$. First, keep m_1 and d unchanged (for memory state $1, \cdots, M$). Second, add the following transition to $(m_1, \mathscr{T}_M, d_M)$: the DM transits to $M+1$ with some probability δ_1 if he was at a memory state in M^1 and received a signal supporting state 1, and transits out of $M+1$ to, randomly, one of the memory states in M^1 with some probability δ_2 if he received a signal supporting state 2. Last, the DM chooses action 1 in memory state $M+1$. The proof involves choosing the appropriate δ_1 and δ_2 such that the DM chooses action 1 in state 1 with the same probability as before. As it is less likely that the DM will transit to memory state $M+1$ in state 2, he chooses action 1 less likely in state 2 and hence strictly improves his asymptotic utility.

Now to prove (ii) and (iii) of Proposition 1, I construct a simple updating mechanism, illustrated in Figure 2. The mechanism tracks only the DM's favorable action and the corresponding confidence level over time. At any period t, the DM believes one of the N actions or no action is favorable, while his confidence level of his favorable action, if he has one, is an integer between 1 and $\lfloor \frac{M-1}{N} \rfloor$. The memory states could thus be represented by $m_t \in \{0\} \cup \{1, \cdots, N\} \times \{1, \cdots, \lfloor \frac{M-1}{N} \rfloor\}$ where memory state 0 stands for no favorable action. The decision rule is such that he takes the favorable action if he has one, and takes action 1 if he does not have a favorable action.[13] The transition rule is described as follows.

[13]The decision rule when the DM does not have a favorable action does not affect the proof and result.

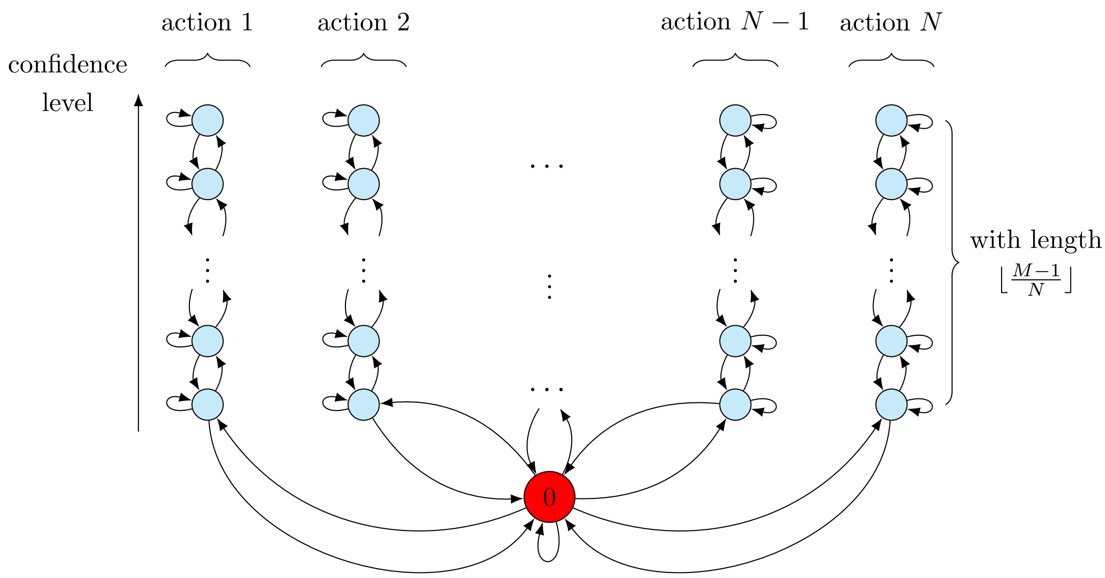

Figure 2: A simple updating mechanism that achieves perfect learning in small worlds for all N, $(p^\omega)_{\omega=1}^N$ and $(f^\omega)_{\omega=1}^N$.

First, the DM starts with no favorable action.[14] If he receives a confirmatory signal for a state ω, he changes his favorable action to action ω with a confidence level 1; if he receives signals that is not confirmatory for any states, he stays in the same memory state 0 in which he has no favorable action. Second, suppose at some period t the DM's favorable action is action ω with confidence level k. If he receives a confirmatory signal for state ω, he revises his confidence level upwards to $k+1$ if k is not already at the maximum $\lfloor \frac{M-1}{N} \rfloor$, and stays in the same memory state if k is at the maximum. Third, if he receives a confirmatory signal for state $\omega' \neq \omega$, he revises his confidence level downwards to $k-1$ with probability $\frac{1}{\delta} < 1$ if $k \geq 2$, transits to the memory state 0 with no favorable action with probability $\frac{1}{\delta} < 1$ if $k = 1$, and stays in the same memory state with probability $1 - \frac{1}{\delta}$. Lastly, if he receives signals that are not supporting any states, he stays in his current memory state with his favorable action and confidence level unchanged. This simple updating mechanism could thus be interpreted as an algorithm that tracks the confidence level of only one state/action at a time, with underreaction to belief-challenging signals (captured by $\frac{1}{\delta} < 1$).

The proof involves choosing a sufficiently large δ such that it is more likely for the DM to adjust his confidence level upwards than to adjust it downwards. This ensures enough "exploitation" that the DM doesn't switch between actions too often. Crucially, when $\frac{M}{N}$ increases, the maximum number of confirmatory signals the DM can "store" increases, and the more likely the DM will be at the correct branch choosing the correct action. When $\frac{M}{N} \to \infty$, the DM almost surely learns perfectly the true state as $t \to \infty$.

Here I point out several noteworthy implications of the analysis. First, as argued above, unlike Cover (1969) and Wilson (2014) with $N = 2$, when $N > 2$, it is in general necessary

[14] The starting memory state has no impact on the long-run distribution over the memory states and does not affect the asymptotic payoff.

to define confirmatory signals for each state stochastically. Thus, a stochastic transition matrix is generally necessary. It implies that upon receiving the same signal, the individual will sometimes regard it as supporting one state, and sometimes regard it as supporting a different state. It resembles empirical evidence that documents heterogeneous interpretations of same pieces of information, especially when the piece of information is imprecise (Gaines et al. (2007)).

Second, as shown in Proposition 1(iii), the simple updating mechanism described in this section approximates perfect learning for all N and $(\boldsymbol{u}^\omega, p^\omega, f^\omega)_{\omega=1}^N$ in small worlds when M is large. Therefore, no knowledge of prior belief $(p^\omega)_{\omega=1}^N$ and utility matrices $(\boldsymbol{u}^\omega)_{\omega=1}^N$ are required. It is in particularly consistent to the result that behavior is close to Bayesian in small worlds as the simple updating mechanism is parsimonious and easy to implement.

Third, perfect learning with the simple updating mechanism in very small worlds is robust to "implementation errors", as shown in the online Appendix E. Roughly speaking, I assume that the DM mistakenly transits to a neighboring memory state with some probability γ in each period regardless of the signal realization s. Such local mistakes could be induced by mistakes in the perception of signals or imperfect tracking (local fluctuation) of memory states. Online Appendix E shows that Proposition 1(iii) hold for all $\gamma \in [0, 1)$, further strengthening the result of perfect learning in very small words.

4.2 Is ignorance optimal?

In this subsection, I present the main results on whether ignorance behavior is optimal in small and big worlds. Ignorance is formally defined as follows: an updating mechanism ignores state ω if the DM almost never chooses action ω no matter what the true state is, i.e.,

$$\lim_{T\to\infty} E_{\omega', m_1 \mathscr{T}, d}\left[\frac{\sum_{t=1}^T \mathbb{1}_{a_t=\omega}}{T}\right] = 0 \text{ for all } \omega'.$$

An updating mechanism is ignorant if it ignores some state. As argued in the introduction, ignorance nests a large set of behavioral biases that depart from the Bayesian model. Note that given the assumption that information strictly improves utility, i.e., $L_M^* < \min_a \sum_{\omega=}^N p^\omega[u(\omega,\omega) - u(a,\omega)]$, when $N = 2$, no ignorant updating mechanism is ϵ-optimal when ϵ is sufficiently small.I present the formal results on ignorance in the following Proposition when $N > 2$, and offer interpretations in the next paragraph.

Proposition 2. *Regarding ignorance behavior:*

(i) For all $N > 2$ and M, there exists some $(\boldsymbol{u}^{\omega'}, p^{\omega'}, f^{\omega'})_{\omega'=1}^N$ and $\bar{\epsilon} > 0$ such that all $\bar{\epsilon}$-optimal updating mechanism ignores state ω for some ω;

(ii) In contrast, take $N > 2$ and $(\boldsymbol{u}^{\omega'}, p^{\omega'}, f^{\omega'})_{\omega'=1}^N$, when M is big enough, all ϵ-optimal update mechanisms are non-ignorant for some small enough ϵ .

First, (i) shows that generally, there exist decision environments such that ignorance is optimal. As $\frac{M}{N}$ is finite, the DM cannot allocate infinite cognitive resources to every state of the world. The DM is bound to make mistakes and faces trades-off between the probability of mistakes in different states of the world. (i) implies that it could be optimal for the DM to ignore some states altogether to improve learning in other states. On the other hand, (ii) shows that the set of decision environments where ignorance is optimal vanishes as M grows large, or equivalently as the world gets smaller. As M increases, the DM can allocate many more memory states to each action. Trade-off between learning in different states is less important. In particular, when $\frac{N}{M}$ converges to 0, the DM learns almost perfectly for all states of the world, and has no incentive to ignore any of the states. Proposition 2 also implies that when $\frac{N}{M}$ is small, the DM will not choose actions that are "safe" but are sub-optimal in every state of the world.

Elaborating on Proposition 2, the following Corollary presents three conditions as examples on when such ignorance happens.

Corollary 1. *When $N > 2$, $u(\omega', \omega) = 0$ for all ω and $\omega' \neq \omega$, there exists some threshold $\xi_p, \xi_u > 0$ and $\overline{F} > 1$ such that if*

(i) $p^\omega < \xi_p$, *or*

(ii) $\frac{u(\omega,\omega)}{\min_{\omega'' \neq \omega} u(\omega'', \omega'')} < \xi_u$, *or*

(iii) $\sup_s \frac{f^\omega(s)}{f^{\omega'}(s)} \times \sup_s \frac{f^{\omega'}(s)}{f^\omega(s)} \leq \overline{F}$ *for some* $\omega' \neq \omega$,

all $\bar{\epsilon}$-optimal mechanisms are ignorant for some $\bar{\epsilon} > 0$.

Intuitively, (i) shows that when the prior probability of a state is low, the DM would rather ignore that state and allocate cognitive resources to learn other states of the world. For example, an individual who is confident about his ability would not update his belief in his but only his teammates' ability, exhibiting persistent overconfidence. Similarly, (ii) shows that when a state is relatively unimportant that the utility of guessing correctly that state is relatively small, the DM would ignore that state to focus on learning others. Lastly, (iii) shows that when it is difficult to distinguish states ω and ω', the DM ignores one of the two states. That is, the DM ignores states that are difficult to identify. To see the intuition, imagine two states ω and ω' where $f^\omega(s) \approx f^{\omega'}(s)$ for all s, and that state ω is a priori less favorable than state ω'. The DM has to receive a large number of signals that support ω against ω' such that he prefers to take action ω instead of ω'. To use many memory states to record signals supporting ω and still be (almost) unsure about choosing ω over ω' is an inefficient use of memory states. Instead, saving those numerous cognitive resources to learn other states of the world, for example, by recording signal supporting ω' against some ω'' and vice versa, he can improve his utility in the other states (ω and ω'') and will be better off. Condition (iii) is particularly applicable in correlation neglect: when it is difficult to distinguish positive or no correlation in the data, and if positive correlation is slightly less likely, the DM ignores the possibility of positive correlation and focuses on learning other states of the world.

Optimal ignorance in symmetric environments

Proposition 2 and Corollary 1 shows that ignorance is optimal in asymmetric environments when $\frac{M}{N}$ is small, *but ignorance is also optimal in some symmetric environments.* I consider the following setup where states and actions are ex-ante identical: There are N possible signal realizations, i.e., $S = \{s^1, \cdots, s^N\}$, and p^ω, $u(\omega, \omega)$, $u(\omega', \omega)$, $\frac{F^\omega(s^\omega)}{F^\omega(s^{\omega'})} > 1$ are invariant across all ω and $\omega' \neq \omega$. Note that for computational simplicity, I consider a discrete signal structure in this Subsection.

The result is shown in Figure 5, where $N = M \in \{4, 6, 8\}$. The y-axis is the asymptotic utility of a symmetric updating mechanism that ignores half of the states, illustrated in Figure 4, minus that of the best non-ignorant updating mechanism, illustrated in Figure 3. Thus, if the y-axis is above 0, all ϵ-optimal mechanisms must be ignorant for small enough ϵ. The x-axis is the informativeness of the signal structure, i.e., $\frac{F^\omega(s^\omega)}{F^\omega(s^{\omega'})}$ where $\omega' \neq \omega$. The analysis and the code is presented in the Appendix B. Figure 5 shows that ignorance is optimal when the informativeness is smaller than some threshold, and the threshold is bigger when N is bigger. In other words, ignorant is optimal when it is difficult to learn, which strengthens the result of Proposition 2.[15] The result also contrasts with the setting when $N = 2$, in which the optimal updating mechanism is symmetric and non-ignorant in such symmetric environments.

The intuition of the result is as follows. To consider all states, the DM allocates one memory state to each action. It is thus easy for the DM to alternate between different actions and unavoidably make mistakes. Put differently, the updating mechanism is "noisy". This is especially true when N is large as one memory state constitutes only a small part of the automaton or when signals are very noisy. If, in contrast, the DM ignores half of the actions, he allocates two memory states to each of the actions that he considers and he switches between actions less frequently. This improves his decision making among the smaller set of states that he considers. When N is large or when $\bar{l}^{\omega\omega'}$ is small, the improvement outweighs the loss he incurs among the states that he ignores, because the loss is small to begin with, i.e., the asymptotic utility of a "noisy" updating mechanism that considers all states is small.

4.3 Is disagreement persistent?

Lastly, I turn to the question of whether disagreement could persist asymptotically in small and big worlds. Consider two individuals A and B who have different utility matrices $(\boldsymbol{u}_A^\omega)_{\omega=1}^N$ and $(\boldsymbol{u}_B^\omega)_{\omega=1}^N$, and/or different prior beliefs $(p_A^\omega)_{\omega=1}^N$ and $(p_B^\omega)_{\omega=1}^N$, and/or different abilities

[15]The magnitude of the thresholds also suggests that ignorance is not an extreme event. For example, when $N = 3$, ignorance is better when the likelihood ratio of the signals is smaller than (around) 3.

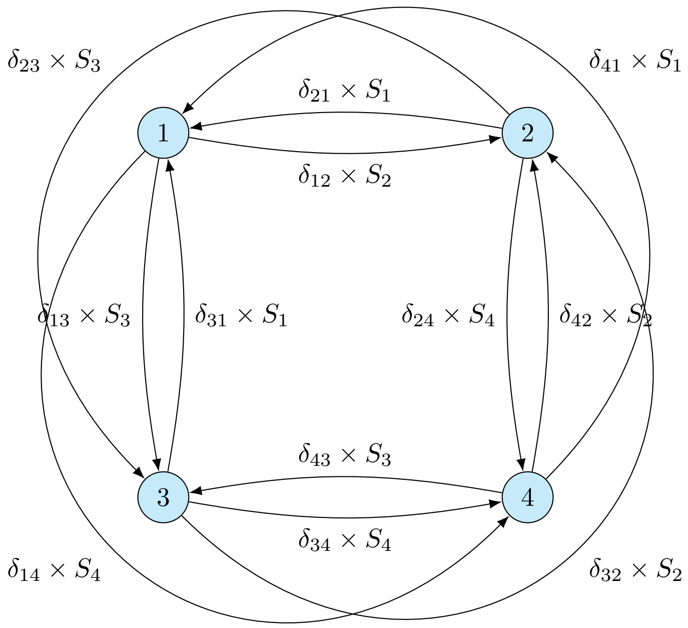

Figure 3: The optimal non-ignorant updating mechanism that considers all states, with $N = M = 4$. The number in the node denotes the action that the DM takes when he is this memory state. Moreover, in memory state ω, and upon receiving a signal that supports state $\omega' \neq \omega$, the DM transits to memory state ω' with probability $\delta_{\omega\omega'} < 1$ and stays in his current memory state otherwise.

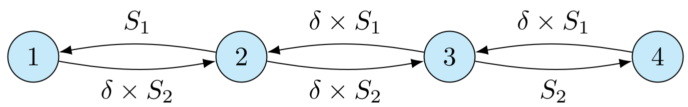

Figure 4: An example of an updating mechanism that ignore two states, with $N = M = 4$. The DM takes action 1 in memory states 1 and 2, and takes action 2 in memory states 3 and 4. In memory state 3, if the DM receives a signal supporting state 2, he transits to memory state 4; if the DM receives a signal supporting state 1, he transits to memory state 2 with probability δ and stays in his current memory state otherwise. In memory state 4, if the DM receives a signal state 1, he transits to memory state 3 with probability δ and stays in his current memory state otherwise. The transition function in memory state 1 and 2 are defined accordingly. δ is chosen to be close to 0 to maximize the asymptotic utility.

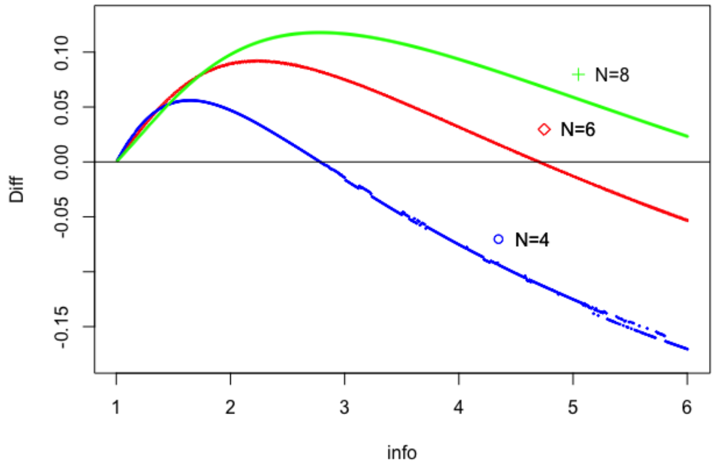

Figure 5: The y-axis is the (estimated) asymptotic utility of the ignorant mechanism illustrated in Figure 4 minus that of the non-ignorant mechanism illustrated in Figure 3, while the x-axis is $\bar{l}^{\omega\omega'}$. The ignorant mechanism outperforms the non-ignorant mechanism when $\bar{l}^{\omega\omega'}$ is smaller than some threshold. Moreover, the threshold is larger as N increases from 4 to 6, and from 6 to 8.

of information acquisition captured by $(f^{\omega}_A)^N_{\omega=1}$ and $(f^{\omega}_B)^N_{\omega=1}$.[16] Their updating mechanisms, $(m_{A1}, \mathscr{F}_A, d_A)$ and $(m_{B1}, \mathscr{F}_B, d_B)$, induce a (random) sequence of actions over time. To define disagreement, I sample one action from each individual A and B and define the "disagreement between $(m_{A1}, \mathscr{F}_A, d_A)$ and $(m_{B1}, \mathscr{F}_B, d_B)$ in state ω" as the probability that the two sampled actions are different.[17] For a given ϵ, the disagreement between individual A and B in state ω is the supremum of disagreement between all pairs of ϵ-optimal updating mechanisms in state ω. Lastly, the two individuals almost always disagree (resp. agree) with each other for a given ϵ if their disagreement equals 1 (resp. 0) in all states.

Corollary 2. *Regarding disagreement:*

(i) For all $N > 2$, M_A and M_B, there exists some $(\boldsymbol{u}^{\omega}_A, p^{\omega}_A, f^{\omega}_A)^N_{\omega=1}$ and $(\boldsymbol{u}^{\omega}_B, p^{\omega}_B, f^{\omega}_B)^N_{\omega=1}$ such that individual A and B almost always disagree as $\epsilon \to 0$.

(ii) In contrast, take $(\boldsymbol{u}^{\omega}_A, p^{\omega}_A, f^{\omega}_A)^N_{\omega=1}$ and $(\boldsymbol{u}^{\omega}_B, p^{\omega}_B, f^{\omega}_B)^N_{\omega=1}$, there exists some $K_1, K_2 > 0$ and $r < 1$ such that the disagreement between individual A and B is bounded above by $\left(K_1 r^{\lfloor \frac{\min\{M_A, M_B\}-1}{N} \rfloor} + K_2\epsilon\right)$ in all states.

[16] For example, individual A could receive noisier signals than individual B, i.e., $f^{\omega}_A = \gamma + (1-\gamma) f^{\omega}_B$ for some $\gamma \in (0,1)$; or individual A could have different learning advantages in identifying some states better but other states worse than individual B, i.e., $\sup_s f^{\omega}_A(s)/f^{\omega'}_A(s) > \sup_s f^{\omega}_B(s)/f^{\omega'}_B(s)$ but $\sup_s f^{\omega''}_A(s)/f^{\omega'''}_A(s) < \sup_s f^{\omega''}_B(s)/f^{\omega'''}_B(s)$ for some $\omega, \omega', \omega'', \omega'''$.

[17] For example, if both individuals alternate between action 1 and 2, their disagreement is $\frac{1}{2}$. An alternative definition is to measure the proportion of time t where $a^A_t \neq a^B_t$: If both individuals alternate between action 1 and 2, their disagreement is 0 if they both start with the same action and is 1 otherwise. Corollary 2 (i) holds with this definition, while the limit result of (ii) holds, i.e., disagreement approaches 0 as M goes to infinite.

(i) shows that in general there exists decision environments such that disagreement almost always happens and is persistent. However, as shown in (ii), when the world gets smaller ($\frac{N}{M}$ goes to 0), such set of decision environments with disagreement vanishes: different individuals with different prior beliefs and/or information acquisition abilities who adopt (almost) optimal updating mechanisms are bound to agree with each other.

The intuition of the result is as follows: as different individuals with different prior beliefs and utilities would adopt different updating mechanisms, they could focus their learning on different subsets of states of the world when they have limited memory. In particular, consider an example with $N = 4$, if individual A ignores state 1 and 2 and individual B ignores state 3 and 4, they will never choose the same action and thus disagree with certainty.[18] However, as $\frac{N}{M}$ goes to 0, decisions gets closer to Bayesian, and the two individuals almost always choose the same actions.

The following result further shows that disagreement could be driven solely by differences in cognitive ability. For simplicity, I allow M to be smaller than N in the following result.

Corollary 3. *There exists $M_A \neq M_B$ and $(\boldsymbol{u}_A^\omega, p_A^\omega, f_A^\omega)_{\omega=1}^N = (\boldsymbol{u}_B^\omega, p_B^\omega, f_B^\omega)_{\omega=1}^N$ such that individual A and B almost always disagree as $\epsilon \to 0$.*

I prove the Corollary with an example where $M_A = 1 < M_B = 2$, $N = 3$, and $u(\omega, \omega) = 1 > u(\omega, \omega') = 0$ for all ω and $\omega' \neq \omega$. Importantly, state 1 is a priori more likely than state 2 and 3, while it is easier to distinguish state 2 and 3 than to distinguish state 1 and state 2, or state 1 and state 3.[19] In this example, individual A always chooses action 1 as he does not have sufficient cognitive resources to learn. On the other hand, if p^1 is not too large compared to p^2 and p^3, individual B ignores state 1. He takes advantages of the informative information structure to distinguish state 2 and 3, such that he can be confident that he doesn't take action 2 in state 3 and vice versa. Thus, as long as state 1 is not a priori too likely, ignoring it would yields a higher utility.[20] To sum up, as individual A has a lower cognitive ability, his learning and actions would depend a lot on the prior belief, while individual B has the ability to take advantage of the information structure.[21] As a result, they adopt different updating mechanisms which leads to disagreement.

[18]Note that the result of disagreement continues to hold in a framework where individuals observe each other's actions. More specifically, one could re-define the signal structures in the current setting to incorporate the information conveyed by the actions taken by the two individuals.

[19]Mathematically, $\sup_s \frac{f^3(s)}{f^2(s)} \sup_s \frac{f^2(s)}{f^3(s)} > \sup_s \frac{f^1(s)}{f^2(s)} \sup_s \frac{f^2(s)}{f^1(s)}$.

[20]One may argue that after seeing individual B choosing action 2 or 3, individual A should change his action. However, this is not possible as he has only one unit of memory capacity $M = 1$ and thus has to effectively commit to one action. In particular, one can generalize this framework to which the two individuals also see each others' actions as signals and Proposition 3 would still hold.

[21]Note that although this example imposes strong assumptions in particular on the size of bounded memory of individual A, it generates a strong form of disagreement in which the two individuals disagree asymptotically with certainty. Similar intuition implies that even when the assumption is relaxed, the difference in M would lead to asymptotic disagreement at least probabilistically.

5 Discussion and Conclusion

Heterogeneity of learning and heuristics under different environments or across different individuals This paper explains a wide range of behavioral anomalies with the same mechanism, i.e., efficient allocation of limited cognitive resources in light of complexity. Importantly, the comparison of small and big worlds illustrates a link between the degree of (relative) complexity of the inference problem and the aforementioned non-Bayesian learning behaviors, which is supported by the experimental results in Enke and Zimmermann (2019) and Graeber (2019). Enke and Zimmermann (2019) shows that correlation neglect negatively correlates with the cognitive ability of subjects and "an extreme reduction in the environment's complexity eliminates the bias", while Graeber (2019) shows that a reduction in the complexity of the problem by removing a decipher stage of signals reduces inattentive learning behavior.

On the other hand, Enke and Zimmermann (2019) and Graeber (2019) also show that simply reminding subjects about the neglected variables reduces inattentive learning and improves inference. This "reminder effect" can be reconciled in the current setup via an effect of a change in the state space. Consider the behavior of inattentive inference in Graeber (2019). The author shows that when subjects are asked to guess the realization of a variable A, they often ignore the effect of another variable B on the signal distribution. Applying to the setting in this paper, consider that before being reminded about the ignored variable, the state space is $\text{supp}(A)\times\text{supp}(B)\times\{B \text{ affects the signal distribution}, B \text{ does not affect the signal distribution}\}$, in which subjects might ignore the states that say "B affects the signal distribution". After being reminded about the effect of B, the set of states of the world is effectively reduced to $\text{supp}(A) \times \text{supp}(B) \times \{B \text{ affects the signal distribution}\}$, the complexity decreases, and subjects adopt another learning mechanism that does not involve ignorant learning.

Future research directions The mechanism mentioned in the previous paragraph brings forth an open question that is not answered in this paper. In reality, individuals face different (sets of) inference problems and are likely endowed with different learning mechanisms for different sets of states of the world. Like in the example mentioned in the previous paragraph, upon receiving new information, individuals could revise the state space and transit from one learning mechanism to another. This is also related to the question of how individuals construct the state space given an inference problem. Arguably, there are infinitely many variables that might affect the signal distributions, and their realizations could be incorporated in the set of possible states. Roughly speaking, the result of ignorance seems to suggest that individuals may only include the most "important" or "a priori probable" states, while the "reminder effect" suggests that the construct of the state space also depends on the information received by the individual. Moving forward, I believe that the question of how individuals construct their perceived state space and the corresponding prior belief deserves more in-depth and careful analysis as it is fundamental to individuals' learning behavior.

References

Acemoglu, Daron, Victor Chernozhukov, and Muhamet Yildiz (2016) "Fragility of asymptotic agreement under Bayesian learning," *Theoretical Economics*, Vol. 11, No. 1, pp. 187–225.

Banovetz, James and Opera Ryan (2020) "Complexity and Procedure Choice," *Working Paper*.

Basu, Pathikrit and Kalyan Chatterjee (2015) "On interim rationality, belief formation and learning in decision problems with bounded memory," *Working Paper*.

Berk, Robert H (1966) "Limiting behavior of posterior distributions when the model is incorrect," *The Annals of Mathematical Statistics*, Vol. 37, No. 1, pp. 51–58.

Börgers, Tilman and Antonio J Morales (2004) "Complexity constraints and adaptive learning: An example," in *Advances in Understanding Strategic Behaviour*: Springer, pp. 9–27.

Caplin, Andrew, Mark Dean, and John Leahy (2019) "Rational inattention, optimal consideration sets, and stochastic choice," *The Review of Economic Studies*, Vol. 86, No. 3, pp. 1061–1094.

Chatterjee, Kalyan and Tai-Wei Hu (2021) "Learning with Limited Memory: Bayesianism vs Heuristics," *Working Paper*.

Chatterjee, Kalyan and Hamid Sabourian (2020) "Game theory and strategic complexity," *Complex Social and Behavioral Systems: Game Theory and Agent-Based Models*, pp. 639–658.

Chauvin, Kyle P (2019) "Euclidean properties of Bayesian updating," *Working paper*.

Compte, Olivier and Andrew Postlewaite (2012) "Belief formation," *Unpublished manuscript, University of Pennsylvania*.

Cover, Thomas M (1969) "Hypothesis testing with finite statistics," *The Annals of Mathematical Statistics*, Vol. 40, No. 3, pp. 828–835.

Enke, Benjamin and Florian Zimmermann (2019) "Correlation neglect in belief formation," *The Review of Economic Studies*, Vol. 86, No. 1, pp. 313–332.

Freedman, David A (1963) "On the asymptotic behavior of Bayes' estimates in the discrete case I," *The Annals of Mathematical Statistics*, Vol. 34, No. 4, pp. 1386–1403.

——— (1965) "On the asymptotic behavior of Bayes estimates in the discrete case II," *The Annals of Mathematical Statistics*, Vol. 36, No. 2, pp. 454–456.

Gaines, Brian J, James H Kuklinski, Paul J Quirk, Buddy Peyton, and Jay Verkuilen (2007) "Same facts, different interpretations: Partisan motivation and opinion on Iraq," *The Journal of Politics*, Vol. 69, No. 4, pp. 957–974.

Gilboa, Itzhak, Larry Samuelson, and David Schmeidler (2020) "Learning (to Disagree?) in Large Worlds," *Working Paper*.

Graeber, Thomas (2019) "Inattentive inference," *Working Paper*.

Heidhues, Paul, Botond Kőszegi, and Philipp Strack (2018) "Unrealistic expectations and misguided learning," *Econometrica*, Vol. 86, No. 4, pp. 1159–1214.

Hellman, Martin E and Thomas M Cover (1970) "Learning with finite memory," *The Annals of Mathematical Statistics*, Vol. 41, No. 3, pp. 765–782.

Hoffman, Mitchell and Stephen V Burks (2017) "Worker overconfidence: Field evidence and implications for employee turnover and returns from training,"Technical report, National Bureau of Economic Research.

Jehiel, Philippe and Jakub Steiner (2020) "Selective sampling with information-storage constraints," *The Economic Journal*, Vol. 130, No. 630, p. 1753–1781, DOI: https://doi.org/10.1093/ej/uez068.

Kahneman, Daniel, Paul Slovic, and Amos Tversky (1982) *Judgment under uncertainty: Heuristics and biases*: Cambridge University Press.

Kahneman, Daniel and Amos Tversky (1979) "Prospect Theory: An Analysis of Decision under Risk," *Econometrica*, Vol. 47, No. 2, pp. 263–292.

Leung, Benson Tsz Kin (2020) "Limited cognitive ability and selective information processing," *Games and Economic Behavior*, Vol. 120, pp. 345–369.

Mailath, George J and Larry Samuelson (2020) "Learning under Diverse World Views: Model-Based Inference," *American Economic Review*, Vol. 110, No. 5, pp. 1464–1501.

Matějka, Filip and Alisdair McKay (2015) "Rational inattention to discrete choices: A new foundation for the multinomial logit model," *American Economic Review*, Vol. 105, No. 1, pp. 272–98.

Monte, Daniel and Maher Said (2014) "The value of (bounded) memory in a changing world," *Economic Theory*, Vol. 56, No. 1, pp. 59–82.

Morris, Stephen (1994) "Trade with heterogeneous prior beliefs and asymmetric information," *Econometrica*, pp. 1327–1347.

Oprea, Ryan (2020) "What Makes a Rule Complex?" *American Economic Review*, p. forthcoming.

Rabin, Matthew and Joel L Schrag (1999) "First impressions matter: A model of confirmatory bias," *The Quarterly Journal of Economics*, Vol. 114, No. 1, pp. 37–82.

Savage, Leonard J (1972) *The foundations of statistics*: Courier Corporation.

Sims, Christopher A (2003) "Implications of rational inattention," *Journal of monetary Economics*, Vol. 50, No. 3, pp. 665–690.

Steiner, Jakub and Colin Stewart (2016) "Perceiving prospects properly," *American Economic Review*, Vol. 106, No. 7, pp. 1601–31.

Wilson, Andrea (2014) "Bounded memory and biases in information processing," *Econometrica*, Vol. 82, No. 6, pp. 2257–2294.

A Proofs

A.1 Proof of Proposition 1

I first define "confirmatory signals" for each state. A signal s supports state ω with probability $G^\omega(s)$ and supports no state with probability $1 - \sum_{\omega=1}^{N} G^\omega(s)$, where

$$G^\omega(s) \propto \frac{f^\omega(s)}{\sqrt{\int (f^\omega(s))^2 ds}} \tag{A.1}$$

with normalization such that $\sum_{\omega=1}^{N} G^\omega(s) \leq 1$. With some abuse of notations, I use F^ω to denote the probability of receiving a signal supporting state ω', or equivalently "a signal $G^{\omega'}$", in state ω, i.e., $F^\omega(G^{\omega'}) \equiv \int f^\omega(s) G^{\omega'}(s)\, ds$. Importantly, Equation A.1 implies it is more likely that the DM receives a signal supporting a correct state than a signal supporting a wrong state:

$$F^\omega(G^\omega) = \frac{\int (f^\omega(s))^2 ds}{\sqrt{\int (f^\omega(s))^2 ds}} = \sqrt{\int (f^\omega(s))^2 ds} > \frac{\int f^\omega(s) f^{\omega'}(s) ds}{\sqrt{\int (f^{\omega'}(s))^2 ds}} = F^\omega(G^{\omega'}),$$

implied by the Cauchy–Schwarz inequality.

Proof of Proposition 1 (i). Consider a ϵ-optimal mechanism with memory size M denoted as $(m_1, \mathscr{T}, d)$. I now construct a mechanism with memory size $M+1$ denoted as $(m_1', \mathscr{T}', d')$ that delivers a strictly higher asymptotic utility than $(m_1, \mathscr{T}, d)$. The result follows when ϵ goes 0. More specifically, pick an $\tilde{m} \in M^1$ and without loss suppose $\arg\min_a u(a, N) = 1$. With some constants $(\delta_i)_{i=0}^{N-1}$ that I will describe later, $(m_1', \mathscr{T}', d')$ follows:

$$\begin{aligned}
m_1' &= m_1. \\
d(m) &= d'(m) \text{ for all } m = 1, \cdots, M. \\
d(M+1) &= 1. \\
\Pr[\mathscr{T}'(m, \cdot) = m'] &\propto \Pr[\mathscr{T}(m, \cdot) = m'] \text{ for all } m \notin \{\tilde{m}, M+1\} \text{ and all } m' \neq m, M+1. \\
\Pr[\mathscr{T}'(\tilde{m}, s) = m] &\propto \begin{cases} \delta_0 \Pr[\mathscr{T}(\tilde{m}, s) = m] \text{ if } m \neq M+1.; \\ \sum_{i=1}^{N} \delta_i G^i(s) \text{ if } m = M+1. \end{cases} \\
\Pr[\mathscr{T}(M+1, s) = m] &\propto \begin{cases} c \text{ if } m = \tilde{m}; \\ 0 \text{ if } m \neq \tilde{m}, M+1. \end{cases}
\end{aligned}$$

for some constant c and with appropriate normalization such that $\sum_{m' \neq m} \Pr[\mathscr{T}'(m, s) = m'] \leq 1$ for all m.[22] $(\delta_i)_{i=1}^{N}$ is chosen such that they satisfy the following system of linear equations for some chosen $\Delta > 1$:

[22] Note that the long-run distribution is invariant if the transition matrix $\Pr[\mathscr{T}'(m, s) = m']$ is scaled up by a common factor.

$$
\delta_1 \begin{bmatrix} \int G^1(s) f^1(s)\, ds - \Delta \int G^N(s) f^1(s)\, ds \\ \int G^1(s) f^2(s)\, ds - \Delta \int G^N(s) f^2(s)\, ds \\ \vdots \\ \int G^1(s) f^{N-1}(s)\, ds - \Delta \int G^N(s) f^{N-1}(s)\, ds \end{bmatrix} + \delta_2 \begin{bmatrix} \int G^2(s) f^1(s)\, ds - \Delta \int G^N(s) f^1(s)\, ds \\ \int G^2(s) f^2(s)\, ds - \Delta \int G^N(s) f^2(s)\, ds \\ \vdots \\ \int G^2(s) f^{N-1}(s)\, ds - \Delta \int G^N(s) f^{N-1}(s)\, ds \end{bmatrix} + \cdots +
$$
$$
\delta_{N-1} \begin{bmatrix} \int G^{N-1}(s) f^1(s)\, ds - \Delta \int G^N(s) f^1(s)\, ds \\ \int G^{N-1}(s) f^2(s)\, ds - \Delta \int G^N(s) f^2(s)\, ds \\ \vdots \\ \int G^{N-1}(s) f^{N-1}(s)\, ds - \Delta \int G^N(s) f^{N-1}(s)\, ds \end{bmatrix} + \delta_N \begin{bmatrix} \int G^N(s) f^1(s)\, ds - \Delta \int G^N(s) f^1(s)\, ds \\ \int G^N(s) f^2(s)\, ds - \Delta \int G^N(s) f^2(s)\, ds \\ \vdots \\ \int G^N(s) f^{N-1}(s)\, ds - \Delta \int G^N(s) f^{N-1}(s)\, ds \end{bmatrix} = 0 \tag{A.2}
$$

Such $(\delta_i)_{i=1}^N$ exists as there are more variables than number of Equations. Now, denote $(\mu_m^{'\omega})_{m=1}^{M+1}$ as the long-run distribution of $(m_1', \mathscr{T}', d')$ and $(\mu_m^\omega)_{m=1}^M$ as the long-run distribution of $(m_1, \mathscr{T}, d)$. Equation (A.2) ensures that

$$
\frac{\mu_{M+1}^{'\omega}}{\mu_{\tilde{m}}^{'\omega}} / \frac{\mu_{M+1}^{'N}}{\mu_{\tilde{m}}^{'N}} = \frac{\sum_{i=1}^N \int G^i(s) f^\omega(s)\, ds}{\sum_{i=1}^N \int G^i(s) f^N(s)\, ds} = \Delta > 1
$$

for all $\omega = 1, \cdots, N-1$. Next, set $\delta_0 = 1 + \frac{\sum_{i=1}^N \int G^i(s) f^N(s)\, ds}{c} \Delta$, we have for $\omega = 1, \cdots, N-1$,

$$
\begin{aligned}
\mu_m^{'\omega} &= \mu_m^\omega \text{ for all } m \neq \tilde{m}, M+1, \\
\mu_{\tilde{m}}^{'\omega} &= \frac{1}{\delta} \mu_{\tilde{m}}^\omega \\
\mu_{\tilde{m}}^{'\omega} + \mu_{M+1}^{'\omega} &= \mu_{\tilde{m}}^{'\omega} \left(1 + \frac{\sum_{i=1}^N \int G^i(s) f^N(s)\, ds}{c} \Delta \right) = \mu_{\tilde{m}}^\omega.
\end{aligned}
$$

One the other hand, in state N,

$$
\begin{aligned}
\mu_m^{'\omega} &= \frac{1}{1 - \frac{\sum_{i=1}^N \int G^i(s) f^N(s)\, ds}{c} (\Delta - 1) \mu_{\tilde{m}}^\omega} \mu_m^\omega \text{ for all } m \neq \tilde{m}, M+1, \\
\mu_{\tilde{m}}^{'\omega} &= \frac{1}{1 - \frac{\sum_{i=1}^N \int G^i(s) f^N(s)\, ds}{c} (\Delta - 1) \mu_{\tilde{m}}^\omega} \left(\frac{1}{\delta} \mu_{\tilde{m}}^\omega \right) \\
\mu_{\tilde{m}}^{'\omega} + \mu_{M+1}^{'\omega} &= \mu_{\tilde{m}}^{'\omega} \left(1 + \frac{\sum_{i=1}^N \int G^i(s) f^N(s)\, ds}{c} \right).
\end{aligned}
$$

Pick $\Delta > 1$ (but not too big), the DM chooses action 1 with a lower probability in state N, and other actions with a higher probability (by a factor of $\frac{1}{1 - \frac{\sum_{i=1}^N \int G^i(s) f^N(s)\, ds}{c} (\Delta - 1) \mu_{\tilde{m}}^\omega}$). As $\arg\min_a u(a, N) = 1$, the result follows. □

Proof of Proposition 1 (ii) and (iii). I formally describe the simple updating mechanism described in the main text. Denote $\lambda = \lfloor (M-1)/N \rfloor$, relabel the memory states as $0, 11, \cdots, 1\lambda, 21, \cdots, 2\lambda, \cdots, N\lambda$ and denote the unused memory states as $N\lambda + 1, N\lambda + 2, \cdots, M - 1$.

The decision rule is as follows:

$$\begin{aligned} d(0) &= 1; \\ d(\omega k) &= \omega \text{ for all } \omega = 1, \ldots, N \text{ and } k = 1, \cdots, \lambda; \\ d(m) &= 1 \text{ for all } m > N\lambda. \end{aligned}$$

The transition function between the memory states is defined below for some big enough $\delta > 1$ such that

$$\frac{\delta F^{\omega}(G^{\omega'})}{\sum_{\omega'' \neq \omega'} F^{\omega}(G^{\omega''})} > 1 \text{ for all } \omega, \omega' \in \Omega.$$

Suppose the DM receives some signal s, he follows the following transition rule:

$$\mathscr{T}(0, s) = \begin{cases} \omega 1 & \text{with probability } G^{\omega}(s) \text{ for all } \omega = 1, \cdots, N; \\ 0 & \text{with probability } 1 - \sum_{\omega=1}^{N} G^{\omega}(s). \end{cases}$$

$$\mathscr{T}(\omega 1, s) = \begin{cases} \omega 2 & \text{with probability } G^{\omega}(s); \\ 0 & \text{with probability } \sum_{j \in \Omega \setminus \{\omega\}} \frac{G^j(s)}{\delta}; \\ \omega 1 & \text{with probability } 1 - \sum_{j \in \Omega \setminus \{\omega\}} \frac{G^j(s)}{\delta} - G^{\omega}(s). \end{cases}$$

$$\mathscr{T}(\omega \lambda, s) = \begin{cases} \omega(\lambda - 1) & \text{with probability } \sum_{j \in \Omega \setminus \{\omega\}} \frac{G^j(s)}{\delta}; \\ \omega \lambda & \text{with probability } 1 - \sum_{j \in \Omega \setminus \{\omega\}} \frac{G^j(s)}{\delta}. \end{cases}$$

while for $k = 2, 3, \cdots, \lambda - 1$,

$$\mathscr{T}(\omega k, s) = \begin{cases} \omega(k + 1) & \text{with probability } G^{\omega}(s); \\ \omega(k - 1) & \text{with probability } \sum_{j \in \Omega \setminus \{\omega\}} \frac{G^j(s)}{\delta}; \\ \omega k & \text{with probability } 1 - \sum_{j \in \Omega \setminus \{\omega\}} \frac{G^j(s)}{\delta} - G^{\omega}(s). \end{cases}$$

Finally, for $m > N\lambda$, $\mathscr{T}(m, s) = m$ for all s. By restricting the initial memory state to one of $0, 11, 12, \cdots, 1\lambda, 21, 22, \cdots, 2\lambda, \cdots, N\lambda$, the DM will never transit to memory states $m > N\lambda$.

Note that this updating mechanism does not depend on p^{ω} nor $\boldsymbol{u}^{\omega}$. Now I compute the long-run distribution $\boldsymbol{\mu}^{\omega}$ and the utility loss $L(m_1, \mathscr{T}, d)$. Fix a state ω, at the two extreme memory states in branch ω', i.e., memory states $\omega'\lambda$ and $\omega'(\lambda - 1)$,

$$\begin{aligned} \mu^{\omega}_{\omega'(\lambda-1)} F^{\omega}(G^{\omega'}) &= \mu^{\omega}_{\omega'\lambda} \frac{1}{\delta} \sum_{\omega'' \neq \omega'} F^{\omega}(G^{\omega''}) \\ \mu^{\omega}_{\omega'(\lambda-1)} &= \mu^{\omega}_{\omega'\lambda} \left[\frac{\delta F^{\omega}(G^{\omega'})}{\sum_{\omega'' \neq \omega'} F^{\omega}(G^{\omega''})} \right]^{-1} \end{aligned}$$

for all ω'. Next, at memory state $\omega'(\lambda-1)$,

$$\mu^{\omega}_{\omega'\lambda}\frac{1}{\delta}\sum_{\omega''\neq\omega'}F^{\omega}(G^{\omega''})+\mu^{\omega}_{\omega'(\lambda-2)}F^{\omega}(G_{\omega'})=\mu^{\omega}_{\omega'(\lambda-1)}\left[\frac{1}{\delta}\sum_{\omega''\neq\omega'}F^{\omega}(G^{\omega''})+F^{\omega}(G^{\omega'})\right]$$

$$\mu^{\omega}_{\omega'(\lambda-2)}=\mu^{\omega}_{\omega'(\lambda-1)}\left[\frac{\delta F^{\omega}(G^{\omega'})}{\sum_{\omega''\neq\omega'}F^{\omega}(G^{\omega''})}\right]^{-1}$$

Repeating the same procedures implies that for all $k=1,\cdots,\lambda$ and $\omega'=1,\cdots,N$

$$\mu^{\omega}_{\omega'k}=\mu^{\omega}_{\omega'\lambda}\left[\frac{\delta F^{\omega}(G^{\omega'})}{\sum_{\omega''\neq\omega'}F^{\omega}(G^{\omega''})}\right]^{-(\lambda-k)}\tag{A.3}$$

and

$$\mu^{\omega}_{\omega'k}=\mu^{\omega}_{0}\left[\frac{\delta F^{\omega}(G^{\omega'})}{\sum_{\omega''\neq\omega'}F^{\omega}(G^{\omega''})}\right]^{k}\tag{A.4}$$

As $\sum_{\omega'=1}^{N}\sum_{k=1}^{\lambda}\mu^{\omega}_{\omega'k}+\mu^{\omega}_{0}=1$, and denote $\frac{\delta F^{\omega}(S^{\omega'})}{\sum_{\omega''\neq\omega'}F^{\omega}(S^{\omega''})}$ by $r^{\omega\omega'}$, we have

$$\mu^{\omega}_{\omega\lambda}\sum_{k=1}^{\lambda}\left[r^{\omega\omega}\right]^{-(\lambda-k)}+\mu^{\omega}_{\omega\lambda}\left[r^{\omega\omega}\right]^{-\lambda}+\mu^{\omega}_{\omega\lambda}\left[r^{\omega\omega}\right]^{-\lambda}\sum_{\omega'\neq\omega}\sum_{k=1}^{\lambda}\left[r^{\omega\omega'}\right]^{k}=1$$

$$\Rightarrow\quad \mu^{\omega}_{\omega\lambda}=\frac{1}{\sum_{k=1}^{\lambda}\left[r^{\omega\omega}\right]^{-(\lambda-k)}+\left[r^{\omega\omega}\right]^{-\lambda}+\left[r^{\omega\omega}\right]^{-\lambda}\sum_{\omega'\neq\omega}\sum_{k=1}^{\lambda}\left[r^{\omega\omega'}\right]^{k}}$$

and the probability of choosing actions in $\Omega\setminus\omega$ in state ω is smaller than $\sum_{m\notin\{\omega 1,\cdots,\omega\lambda\}}\mu^{\omega}_{m}$ which is as follows:

$$\begin{aligned}
&\frac{\left[r^{\omega\omega}\right]^{-\lambda}+\left[r^{\omega\omega}\right]^{-\lambda}\sum_{\omega'\neq\omega}\sum_{k=1}^{\lambda}\left[r^{\omega\omega'}\right]^{k}}{\sum_{k=1}^{\lambda}\left[r^{\omega\omega}\right]^{-(\lambda-k)}+\left[r^{\omega\omega}\right]^{-\lambda}+\left[r^{\omega\omega}\right]^{-\lambda}\sum_{\omega'\neq\omega}\sum_{k=1}^{\lambda}\left[r^{\omega\omega'}\right]^{k}}\\
&<\left[r^{\omega\omega}\right]^{-\lambda}+\left[r^{\omega\omega}\right]^{-\lambda}\sum_{\omega'\neq\omega}\sum_{k=1}^{\lambda}\left[r^{\omega\omega'}\right]^{k}\\
&=\left[r^{\omega\omega}\right]^{-\lambda}+\sum_{\omega'\neq\omega}\frac{r^{\omega\omega'}}{r^{\omega\omega'}-1}\left[\left(\frac{r^{\omega\omega'}}{r^{\omega\omega}}\right)^{\lambda}-\left[r^{\omega\omega}\right]^{-\lambda}\right]\\
&<\left[r^{\omega\omega}\right]^{-\lambda}+\sum_{\omega'\neq\omega}\frac{r^{\omega\omega'}}{r^{\omega\omega'}-1}\left[\left(\frac{r^{\omega\omega'}}{r^{\omega\omega}}\right)^{\lambda}\right]\\
&<\left[r^{\omega\omega}\right]^{-\lambda}+(N-1)\frac{\max_{\omega'\neq\omega}r^{\omega\omega'}}{\max_{\omega'\neq\omega}r^{\omega\omega'}-1}\left[\max_{\omega'\neq\omega}\left(\frac{r^{\omega\omega'}}{r^{\omega\omega}}\right)^{\lambda}\right]\\
&<\left[(N-1)\frac{\max_{\omega'\neq\omega}r^{\omega\omega'}}{\max_{\omega'\neq\omega}r^{\omega\omega'}-1}+1\right]\max\left\{\left[r^{\omega\omega}\right]^{-\lambda},\max_{\omega'\neq\omega}\left(\frac{r^{\omega\omega'}}{r^{\omega\omega}}\right)^{\lambda}\right\}
\end{aligned}\tag{A.5}$$

The first inequality of Equation (A.5) is implied by the fact that the denominator is strictly greater than 1. Now, using Equation (A.5) and denote $K^\omega = \left[(N-1)\frac{\max_{\omega'\neq\omega} r^{\omega\omega'}}{\max_{\omega'\neq\omega} r^{\omega\omega'}-1}+1\right]$ we can compute the upper bound of the utility loss $L(m_1, \mathscr{T}, d)$

$$\begin{aligned}
L(m_1,\mathscr{T},d) &\leq \sum_{\omega=1}^{N} p^\omega \max_{a\neq\omega}\{u(\omega,\omega)-u(a,\omega)\}\left[\sum_{m\notin\{\omega1,\cdots,\omega\lambda\}}\mu_m^\omega\right] \\
&< \sum_{\omega=1}^{N} p^\omega \max_{a\neq\omega}\{u(\omega,\omega)-u(a,\omega)\}K^\omega \max\left\{[r^{\omega\omega}]^{-\lambda}, \max_{\omega'\neq\omega}\left(\frac{r^{\omega\omega'}}{r^{\omega\omega}}\right)^\lambda\right\} \\
&< \max_\omega\left[\max_{a\neq\omega}\{u(\omega,\omega)-u(a,\omega)\}K^\omega \max\left\{[r^{\omega\omega}]^{-\lambda}, \max_{\omega'\neq\omega}\left(\frac{r^{\omega\omega'}}{r^{\omega\omega}}\right)^\lambda\right\}\right]
\end{aligned} \tag{A.6}$$

As $[r^{\omega\omega}]^{-1}$ and $\max_{\omega'\neq\omega}\left(\frac{r^{\omega\omega'}}{r^{\omega\omega}}\right)$ are strictly smaller than 1 for all ω, Proposition 1 (ii) follows. Moreover, as $[r^{\omega\omega}]^{-}1$ and $\max_{\omega'\neq\omega}\left(\frac{r^{\omega\omega'}}{r^{\omega\omega}}\right)$ are strictly smaller than 1 for all ω, the right-hand side of Equation (A.5) converges to 0 as λ goes to ∞. This proves Proposition 1 (iii). □

A.2 Proof of Proposition 2 and Corollary 1

Proof of Proposition 2 (ii). First, an ignorant updating mechanism induces utility loss weakly greater than $\min_{\omega\in\Omega}\min_{a\neq\omega}[p^\omega(u(\omega,\omega)-u(a,\omega))] > 0$. On the other hand, as shown in Proposition 1, L_M^* converges to 0 as $M\to\infty$. The monotonicity of L_M^* implies there exists some big enough $\bar{M}$ such that for $M\geq\bar{M}$, $L_M^*\leq L_{\bar{M}}^* < \min_{\omega\in\Omega}\min_{a\neq\omega}[p^\omega(u(\omega,\omega)-u(a,\omega))]$. The result follows. □

Proof of Proposition 2 (i) and Corollary 1. First note that Corollary 1 implies Proposition 2 (i). I first prove Corollary 1(i). Consider an updating mechanism $(m_1,\mathscr{T},d)$ such that $\mu_m^{\tilde{\omega}} > 0$ for some $\tilde{\omega}$ and m where $d(m)=\omega$. Now I prove that there exists a different updating mechanism that yields higher asymptotic utility if p^ω is small enough.

Consider state ω and ω', recall that the long-run distribution in state ω and ω' are the solution of the following fixed point equations:

$$\begin{aligned}
\boldsymbol{\mu}^\omega &= (\boldsymbol{\mu}^\omega)^T\mathbf{Q}^\omega; \\
\boldsymbol{\mu}^{\omega'} &= (\boldsymbol{\mu}^{\omega'})^T\mathbf{Q}^{\omega'}.
\end{aligned} \tag{A.7}$$

Equation (A.7) shows that, given an updating mechanism $(m_1,\mathscr{T},d)$, f^ω and $f^{\omega'}$, the long-run distribution $\boldsymbol{\mu}^\omega$, $\boldsymbol{\mu}^{\omega'}$ and the ratio $(\frac{\mu_m^\omega}{\mu_m^{\omega'}})_{m=1}^M$ are invariant of N, $(p^{\omega''})_{\omega''=1}^N$ and $(f^{\omega''})_{\omega''\neq\omega,\omega'}$. Therefore, we can apply Theorem 2 of Hellman and Cover (1970) where $N=2$ such that:

$$\frac{\max_m \frac{\mu_m^{\omega'}}{\mu_m^\omega}}{\min_m \frac{\mu_m^{\omega'}}{\mu_m^\omega}} \leq (\bar{l}^{\omega\omega'}\bar{l}^{\omega'\omega})^{(M-1)}. \tag{A.8}$$

where $\bar{l}^{\omega\omega'} = \sup_s \frac{f^\omega(s)}{f^{\omega'}(s)}$. First, $\min_m \frac{\mu_m^{\omega'}}{\mu_m^\omega} > 0$ for all $\omega'\neq\omega$. Suppose to the contrary that for some ω', $\min_m \frac{\mu_m^{\omega'}}{\mu_m^{\omega}} = 0$, Equation (A.8) implies that $\max_m \frac{\mu_m^{\omega'}}{\mu_m^{\omega}} = 0$ and $\mu_m^{\omega'} = 0$ for all m. It in turn implies that $\max_m \frac{\mu_m^{\omega}}{\mu_m^{\omega'}}$ is unbounded and $\min_m \frac{\mu_m^{\omega}}{\mu_m^{\omega'}} = 0$ and it contradicts

$$\frac{\max_m \frac{\mu_m^{\omega}}{\mu_m^{\omega'}}}{\min_m \frac{\mu_m^{\omega}}{\mu_m^{\omega'}}} \leq (\bar{l}^{\omega'\omega}\bar{l}^{\omega\omega'})^{(M-1)}.$$

Now since $\min_m \frac{\mu_m^{\omega'}}{\mu_m^{\omega}} > 0$ for all ω', denote $u^{\omega} = u(\omega,\omega)$, we have

$$\begin{aligned}
\min_m \frac{\mu_m^{\omega'}}{\mu_m^{\omega}} &\geq (\bar{l}^{\omega\omega'}\bar{l}^{\omega'\omega})^{-(M-1)} \max_m \frac{\mu_m^{\omega'}}{\mu_m^{\omega}} \\
\min_m \frac{\mu_m^{\omega'}}{\mu_m^{\omega}} &\geq \varsigma^{2(M-1)} \max_m \frac{\mu_m^{\omega'}}{\mu_m^{\omega}} \\
\min_m \frac{\mu_m^{\omega'}}{\mu_m^{\omega}} &\geq \varsigma^{2(M-1)} \\
\frac{u^{\omega'}p^{\omega'}\mu_m^{\omega'}}{u^{\omega}p^{\omega}\mu_m^{\omega}} &\geq \varsigma^{2(M-1)} \frac{u^{\omega'}p^{\omega'}}{u^{\omega}p^{\omega}} \text{ for all } m.
\end{aligned}$$

The second inequality follows Equation (1) while the third inequality is implied by the fact that $\sum_m \mu_m^{\omega'} = \sum_m \mu_m^{\omega} = 1$. Now, denote $\overline{u} = \max_{\omega\in\Omega} \min_{a\neq\omega}[u(\omega,\omega) - u(a,\omega)]$ and $\underline{u} = \min_{\omega\in\Omega} \min_{a\neq\omega}[u(\omega,\omega) - u(a,\omega)]$, as $\min \max_{\omega'} \frac{u^{\omega'}p^{\omega'}}{u^{\omega}p^{\omega}} = \frac{\underline{u}}{\overline{u}} \frac{\frac{1-p^{\omega}}{N-1}}{p^{\omega}}$, there exists some ω' such that

$$\frac{u^{\omega'}p^{\omega'}\mu_m^{\omega'}}{u^{\omega}p^{\omega}\mu_m^{\omega}} \geq \varsigma^{2(M-1)} \min \max_{\omega'} \frac{u^{\omega'}p^{\omega'}}{u^{\omega}p^{\omega}} \geq \varsigma^{2(M-1)} \frac{\underline{u}}{\overline{u}} \frac{\frac{1-p^{\omega}}{N-1}}{p^{\omega}} \text{ for all } m \in M. \tag{A.9}$$

When $p^{\omega} < \varsigma^{2(M-1)} \frac{\underline{u}}{\overline{u}} \frac{1-p^{\omega}}{N-1}$, or equivalently, $\frac{p^{\omega}}{1-p^{\omega}} < \varsigma^{2(M-1)} \frac{\underline{u}}{\overline{u}} \frac{1}{N-1}$, Equation (A.9) implies that $u^{\omega'}p^{\omega'}\mu_m^{\omega'} > u^{\omega}p^{\omega}\mu_m^{\omega}$, and the DM is better off choosing ω' than choosing ω for all $m \in M$. Setting $\bar{\epsilon} = \frac{u^{\omega'}p^{\omega'}\mu_m^{\omega'} - u^{\omega}p^{\omega}\mu_m^{\omega}}{\sum_{\omega''} p^{\omega''}\mu_m^{\omega''}}$ proves the first bullet point of Corollary 1. Similar argument proves Corollary 1(ii). Check!

I now prove Corollary 1(iii). Again, consider an updating mechanism $\mathscr{T}$ such that $\mu_m^{\tilde{\omega}} > 0$ for some $\tilde{\omega}$ and m where $d(m) = \omega$. Similar to the proof of the first part of Corollary 1,

$$\begin{aligned}
\min_m \frac{\mu_m^{\omega'}}{\mu_m^{\omega}} &\geq (\bar{l}^{\omega\omega'}\bar{l}^{\omega'\omega})^{-(M-1)} \max_m \frac{\mu_m^{\omega'}}{\mu_m^{\omega}} \\
\min_m \frac{\mu_m^{\omega'}}{\mu_m^{\omega}} &\geq \overline{F}^{-(M-1)} \max_m \frac{\mu_m^{\omega'}}{\mu_m^{\omega}} \\
\min_m \frac{\mu_m^{\omega'}}{\mu_m^{\omega}} &\geq \overline{F}^{-(M-1)} \\
\frac{u^{\omega'}p^{\omega'}\mu_m^{\omega'}}{u^{\omega}p^{\omega}\mu_m^{\omega}} &\geq \overline{F}^{-(M-1)} \frac{u^{\omega'}p^{\omega'}}{u^{\omega}p^{\omega}} \text{ for all } m.
\end{aligned} \tag{A.10}$$

As $u^{\omega'}p^{\omega'} > u^{\omega}p^{\omega}$, for $\overline{F} < \left[\frac{u^{\omega}p^{\omega}}{u^{\omega'}p^{\omega'}}\right]^{\frac{1}{M-1}}$, Equation (A.10) implies that $u^{\omega'}p^{\omega'}\mu_m^{\omega'} > u^{\omega}p^{\omega}\mu_m^{\omega}$, i.e., the DM is better to choose action ω' instead of ω. Setting $\bar{\epsilon} = \frac{u^{\omega'} p^{\omega'} \mu_m^{\omega'} - u^{\omega} p^{\omega} \mu_m^{\omega}}{\sum_{\omega''} p^{\omega''} \mu_m^{\omega''}}$, the result follows. $\square$

A.3 Proof of Corollary 2

Proof. First, (i) is directly implied by Corollary 1. If p_A^ω is small enough for all $\omega \in N_A \subset N$ and p_B^ω is small enough for all $\omega \in N_B = N \setminus N_A$, individual A (almost) never picks action ω for all $\omega \in N_A$ and individual B (almost) never picks action ω for all $\omega \in N \setminus N_A$. Therefore, they must disagree with each other.

To prove (ii), note that by Proposition 1, we know that if an individual adopts an ϵ-optimal updating mechanism, his utility loss is bounded above by $Kr^{\lfloor \frac{M-1}{N} \rfloor} + \epsilon$ for some $K > 0$ and $r < 1$. Thus, for all ω, the probability that the individual chooses a "wrong" action $\omega' \neq \omega$ is strictly smaller than $\left[\frac{Kr^{\lfloor \frac{M-1}{N} \rfloor} + \epsilon}{\min_\omega p^\omega} \right]$. For all ω, the probability that both individuals chooses action ω and thus agree with each other is greater than

$$\begin{aligned}
&\left[1 - \left[\frac{K_A r_A^{\lfloor \frac{M_A-1}{N} \rfloor} + \epsilon}{\min_\omega p_A^\omega} \right] \right] \left[1 - \left[\frac{K_B r_B^{\lfloor \frac{M_B-1}{N} \rfloor} + \epsilon}{\min_\omega p_B^\omega} \right] \right] \\
>& 1 - \left[\frac{1}{\min_\omega p_A^\omega} + \frac{1}{\min_\omega p_B^\omega} \right] (\max\{K_A, K_B\} (\max\{r_A, r_B\})^{\lfloor \frac{\min\{M_A, M_B\}-1}{N} \rfloor} + \epsilon)
\end{aligned}$$

for all ϵ-optimal updating mechanisms of individual A and B. The result follows.

$\square$

A.4 Proof of Corollary 3

Proof. I prove the Corollary using the following simple example. Consider a setting with $N = 3$ and two individuals, A and B, who share the same prior belief and the same objective signal structure:

$$\begin{aligned}
p^1 &= \frac{1}{3} + 2\nu \\
p^2 = p^3 &= \frac{1}{3} - \nu
\end{aligned} \tag{A.11}$$

$$\begin{aligned}
\sup_s \frac{f^1(s)}{f^n(s)} &= \sup_s \frac{f^n(s)}{f^1(s)} = \sqrt{1+\tau} \text{ for } n = 2, 3 \\
\sup_s \frac{f^2(s)}{f^3(s)} &= \sup_s \frac{f^3(s)}{f^2(s)} = \sqrt{1+\Psi} \text{ where } \Psi > \tau.
\end{aligned} \tag{A.12}$$

with $1 + \tau \geq \frac{\frac{1}{3}+2\nu}{\frac{1}{3}-\nu}$.[23] Moreover, assume that $u(1,1) = u(2,2) = u(3,3) = 1$ and $u(\omega, \omega') = 0$ for all $\omega \neq \omega'$. The only difference the two individuals have is their levels of cognitive ability,

[23]It ensures that if $M \geq 2$, the DM never chooses action 1 with probability 1 and he can achieve a strictly lower utility loss compared to the benchmark of no information.

where $M_A = 1$ and $M_B = 2$.

First, as $M = 1$ for individual A, his action is constant in all periods for all signal realizations and thus $d(1) = 1$. Individual A always take action 1. Now I characterize the ϵ-optimal updating mechanism of individual B. With some abuse of notations, denote $L_2^*(nn')$ as optimal utility loss where the DM chooses action n in memory state 1 and action n' in memory state 2. We have, obviously,

$$L_2^*(11) = \frac{2}{3} - 2\nu,$$
$$L_2^*(22) = L_2^*(33) = \frac{2}{3} + \nu.$$

Moreover, as argued in the proof of Proposition 2 (i), suppose $d(1) = n$ and $d(2) = n' \neq n$, we have

$$\frac{\mu_1^n}{1-\mu_2^{n'}} / \frac{1-\mu_1^n}{\mu_2^{n'}} = \frac{\mu_1^n}{\mu_1^{n'}} / \frac{\mu_2^n}{\mu_2^{n'}} \leq \sup_s \frac{f^n(s)}{f^{n'}(s)} \sup_s \frac{f^{n'}(s)}{f^n(s)}.$$

The upper bound of $L^*(nn')$ is given by the following minimization problem:

$$\min_{\mu_1^n, \mu_2^{n'}} 1 - p^n - p^{n'} + p^n(1-\mu_1^n) + +p^{n'}(1-\mu_2^{n'})$$
$$\text{given } \frac{\mu_1^n}{1-\mu_2^{n'}} / \frac{1-\mu_1^n}{\mu_2^{n'}} \leq \sup_s \frac{f^n(s)}{f^{n'}(s)} \sup_s \frac{f^{n'}(s)}{f^n(s)},$$

Hellman and Cover (1970) shows that the upper bound is tight with the updating mechanism where $\mathscr{T}(1, s') = 2$ if and only if $\frac{f^{n'}(s')}{f^n(s')}$ is close to $\sup_s \frac{f^{n'}(s)}{f^n(s)}$, and $\mathscr{T}(2, s') = 1$ if and only if $\frac{f^n(s')}{f^{n'}(s')}$ is close to $\sup_s \frac{f^n(s)}{f^{n'}(s)}$. Therefore,

$$L_2^*(12) = L_2^*(13) = \frac{1}{3} - \nu + \frac{2\sqrt{(1+\tau)\left(\frac{1}{3}+2\nu\right)\left(\frac{1}{3}-\nu\right)} - \left(\frac{2}{3}-\nu\right)}{\tau},$$
$$L_2^*(23) = \frac{1}{3} + 2\nu + \frac{2\left(\frac{1}{3}-\nu\right)\sqrt{1+\Psi} - \left(\frac{2}{3}-2\nu\right)}{\Psi}.$$

where $1 + \tau \geq \frac{\frac{1}{3}+2\nu}{\frac{1}{3}-\nu}$ implies $L_2^*(22) = L_2^*(33) > L_2^*(11) \geq L_2^*(12) = L_2^*(13)$. In the following, I prove $L_2^*(12) > L_2^*(23)$ if ν is small enough which implies that $a_t^B = 2$ or 3 for all t in all ϵ-optimal updating mechanism when ϵ is smaller than $L_2^*(12) - L_2^*(23)$. Now, $L_2^*(12) > L_2^*(23)$ if and only if

$$L_2^*(12) - L_2^*(23) = 3\nu + \frac{2(\frac{1}{3}-\nu)\sqrt{1+\Psi}}{\Psi} - \frac{\frac{2}{3}-2\nu}{\Psi} - \frac{2\sqrt{(1+\tau)(\frac{1}{3}+2\nu)(\frac{1}{3}-\nu)}}{\tau} + \frac{\frac{2}{3}+\nu}{\tau} < 0.$$

When $\nu = 0$,

$$L_2^*(12) - L_2^*(23) = \frac{2}{3}\left(\frac{\sqrt{1+\Psi}}{\Psi} - \frac{\sqrt{1+\tau}}{\tau}\right) - \frac{2}{3}\left(\frac{1}{\Psi} - \frac{1}{\tau}\right).$$

As both $\frac{\sqrt{1+x}}{x}$ and $\frac{1}{x}$ decreases in x, $\Psi > \tau$ implies that $L_2^*(12) - L_2^*(23) < 0$ when $\nu = 0$. The result follows by continuity.

$\square$

B An example of ignorance with uniform prior belief and symmetric signal structures

In this section, I consider a case where $N = M \geq 4$ and states of the world are a priori uniformly distributed, i.e., $p^\omega = \frac{1}{N}$ for all $\omega = \{1, \cdots, N\}$. Moreover, $u(\omega, \omega) = 1 > u(\omega', \omega) = 0$ for all ω and $\omega' \neq \omega$. For simplicity, consider a class of "symmetric" discrete signal structures where $S = \{s^1, \cdots .s^N\}$ and s^ω is a signal that supports state ω. More specifically,

$$F^\omega(s^\omega) = \mathscr{I} F^\omega(s^{\omega'}) \text{ for all } \omega \text{ and } \omega' \neq \omega \text{ where } \mathscr{I} > 1.$$

Thus, under all states of the world, it is $\mathscr{I}$ times more likely to receive a signal that supports the true state than a signal that supports one of the other states.

In such a symmetric environment, there seems to be no reason to ignore any of the states. However, I will present an example that shows that it is beneficial to ignores some states when N is large or $\mathscr{I}$ is small. First, consider a simple "symmetric" updating mechanism that ignores no states, illustrated in Figure 3 with an example of $N = 4$. As the DM ignores no state, he allocates one memory state to each action. Without loss of generality, assume he takes action ω in memory state ω. When the DM is in memory state $m = \omega$, upon receiving a signal $s^{\omega'}$, i.e., a signal that supports state ω', he transits to memory state ω' with some probability $\delta_{m\omega} \leq 1$, and stays in his current memory state otherwise. Formally, the transition function is as follows:

$$\mathscr{T}(m, s^\omega) = \delta_{m\omega} \times \{\omega\} + (1 - \delta_{m\omega}) \times \{m\} \text{ for all } m \text{ and } \omega.$$

In the following I show that such updating mechanism is optimal among the class of all non-ignorant mechanism.

Suppose for some non-ignorant mechanism that the DM chooses action ω in memory state ω, and the DM transits from memory state $m \neq \omega$ to memory state ω upon receiving signal $s^{\omega'}$ with strictly positive probability where $\omega' \neq \omega$. Similar to the proof of Proposition 1 (i), I construct an updating mechanism that strictly improves asymptotic utility. First, decrease the probability of transiting from memory state m to ω upon receiving signal $s^{\omega'}$ by p. Second, increase the probability of transiting from memory state m to ω upon receiving signal s^ω by p (re-normalize if necessary to ensure the probability of transiting from memory state m to ω upon receiving signal s^ω is not greater than 1). These two steps do not change the utility in all state $\omega'' \neq \omega, \omega'$ because $F^{\omega''}(s^{\omega'}) = F^{\omega''}(s^\omega)$, increase the utility in state ω but potentially decrease the utility in state ω'. Last, scale up all the transitions out of memory state ω (again, re-normalize if necessary) such that the utility in state ω is the same before the first step of construction. As it is less likely to receive s^{ω} in state ω' than in state ω, the last step "over-scale-up" the transition out of memory state ω in ω', and thus increases the utility in state ω' such that it is higher before the construction. It also increases the utility in state $\omega'' \neq \omega, \omega'$ such that it is higher before the construction.

Now consider an ignorant mechanism that follows the similar idea of the non-ignorant updating mechanism illustrated in Figure 3, but ignores half of the states of the worlds. For simplicity, assume that N is plural. The ignorant mechanism is illustrated in Figure 4, with an example of $N = 4$. By ignoring half of the states, the DM allocates two memory states to each action that he does not ignore. Without loss of generality, assume that the DM takes action ω in memory states $2\omega - 1$ and 2ω for $\omega \leq \frac{N}{2}$, and ignores all actions $\omega' > \frac{N}{2}$. In the "more confident" memory state 2ω, upon receiving a signal supporting state $\omega' \neq \omega$ where $\omega' \leq \frac{N}{2}$, the DM transits to state $2\omega - 1$ with probability $\delta < 1$ and stays in his current memory state otherwise. In the "less confident" memory state $2\omega - 1$, upon receiving a signal that supports state ω, he transits to the "more confident" memory state 2ω; upon receiving a signal that supports state $\omega' \neq \omega$ where $\omega' \leq \frac{N}{2}$, the DM transits to state $2\omega' - 1$ with probability $\delta < 1$ and stays in his current memory state otherwise.

Formally, the transition function is as follows:

$$
\begin{aligned}
\mathscr{T}(2\omega, s^{\omega}) &= 2\omega && \text{for all } \omega \leq \frac{N}{2}, \\
\mathscr{T}(2\omega, s^{\omega'}) &= \begin{cases} 2\omega - 1 \text{ with probability } \delta \\ 2\omega \text{ with probability } 1 - \delta \end{cases} && \text{for all } \omega \leq \frac{N}{2}, \omega' \neq \omega \text{ and } \omega' \leq \frac{N}{2}, \\
\mathscr{T}(2\omega - 1, s^{\omega}) &= 2\omega && \text{for all } \omega \leq \frac{N}{2}, \\
\mathscr{T}(2\omega - 1, s^{\omega'}) &= \begin{cases} 2\omega' - 1 \text{ with probability } \delta \\ 2\omega \text{ with probability } 1 - \delta \end{cases} && \text{for all } \omega \leq \frac{N}{2}, \omega' \neq \omega \text{ and } \omega' \leq \frac{N}{2}, \\
\mathscr{T}(m, s^{\omega'}) &= m && \text{for all } \omega > \frac{N}{2}.
\end{aligned}
$$

Suppose the true state is 1, In the stationary distribution, we must have

$$
\begin{aligned}
\delta \mu^1_{2\omega} \sum_{\omega' \neq \omega, \omega' \leq \frac{N}{2}} F^1(s^{\omega'}) &= \mu^1_{2\omega-1} F^1(s^{\omega}) \text{ for all } \omega \leq \frac{N}{2}, \\
\delta \mu^1_{2\omega-1} F^1(s^{\omega'}) &= \delta \mu^1_{2\omega'-1} F^1(s^{\omega}) \text{ for all } \omega, \omega' \leq \frac{N}{2} \text{ and } \omega \neq \omega'.
\end{aligned} \tag{B.1}
$$

From the first Equation of Equation (B.1), we know that when δ is close to 0, $\mu^1_{2\omega-1}$ is close to 0 for all $\omega \leq \frac{N}{2}$. Moreover, we have

$$\begin{aligned}
\mu_2^1 &= \frac{F^1(s^1)}{\delta \sum_{\omega' \neq 1, \omega' \leq \frac{N}{2}} F^1(s^{\omega'})} \times \frac{\delta F^1(s^1)}{\delta F^1(s^\omega)} \times \frac{\delta \sum_{\omega' \neq \omega, \omega' \leq \frac{N}{2}} F^1(s^{\omega'})}{F^1(s^\omega)} \times \mu_{2\omega}^1 \\
&= \frac{\mathscr{I}}{\frac{N}{2} - 1} \times \mathscr{I} \times \left(\frac{N}{2} - 2 + \mathscr{I} \right) \times \mu_{2\omega}^1 \\
&= \frac{\mathscr{I}^2(N + 2\mathscr{I} - 4)}{N - 2} \mu_{2\omega}^1
\end{aligned}$$

for all $\omega \neq 1$ and $\omega \leq \frac{N}{2}$. We have

$$\begin{aligned}
\mu_2^1 + \sum_{\omega \neq 1, \omega \leq \frac{N}{2}} \mu_{2\omega}^1 &= 1 \\
\mu_2^1 + \frac{N - 2}{\mathscr{I}^2(N + 2\mathscr{I} - 4)} \left(\frac{N}{2} - 1 \right) \mu_2^1 &= 1 \\
\mu_2^1 &= \frac{2\mathscr{I}^2(N + 2\mathscr{I} - 4)}{2\mathscr{I}^2(N + 2\mathscr{I} - 4) + (N - 2)^2}.
\end{aligned}$$

By repeating the same computation for all $\omega \leq \frac{N}{2}$, the asymptotic utility equals:

$$\sum_{\omega=1}^{\frac{N}{2}} \frac{2\mathscr{I}^2(N + 2\mathscr{I} - 4)}{2\mathscr{I}^2(N + 2\mathscr{I} - 4) + (N - 2)^2} \times \frac{1}{N} = \frac{\mathscr{I}^2(N + 2\mathscr{I} - 4)}{2\mathscr{I}^2(N + 2\mathscr{I} - 4) + (N - 2)^2}. \tag{B.2}$$

Then the R code on `https://sites.google.com/site/ltkbenson/research` produces Figure 5 in the main text.