Mislearning from Censored Data: The Gambler's Fallacy in Optimal-Stopping Problems

Kevin He*

First version: March 21, 2018
This version: November 20, 2018

Abstract

I study endogenous learning dynamics for people expecting systematic reversals from random sequences — the "gambler's fallacy." Biased agents face an optimal-stopping problem, such as managers conducting sequential interviews. They are uncertain about the underlying distribution (e.g. talent distribution in the labor pool) and must learn its parameters from previous agents' histories. Agents stop when early draws are deemed "good enough," so predecessors' histories contain negative streaks but not positive streaks. Since biased learners understate the likelihood of consecutive below-average draws, histories induce pessimistic beliefs about the distribution's mean. When early agents decrease their acceptance thresholds due to pessimism, later learners will become more surprised by the lack of positive reversals in their predecessors' histories, leading to even more pessimistic inferences and even lower acceptance thresholds — a positive-feedback loop. Agents who are additionally uncertain about the distribution's variance believe in fictitious variation (exaggerated variance) to an extent depending on the severity of data censoring. When payoffs are convex in the draws (e.g. managers can hire previously rejected interviewees), variance uncertainty provides another channel of positive feedback between past and future thresholds.

Latest version of this paper: https://scholar.harvard.edu/files/kevin/files/gambler.pdf
Online Appendix: https://scholar.harvard.edu/files/kevin/files/gambler_oa.pdf

*Department of Economics, Harvard University. Email: hesichao@gmail.com. I am indebted to Drew Fudenberg, Matthew Rabin, Tomasz Strzalecki, and Ben Golub for their guidance and support. I thank Isaiah Andrews, Ruiqing Cao, In-Koo Cho, Martin Cripps, Krishna Dasaratha, Jetlir Duraj, Ben Enke, Ignacio Esponda, Jiacheng Feng, Tristan Gagnon-Bartsch, Ashvin Gandhi, Oliver Hart, Alice Hsiaw, Yizhou Jin, Michihiro Kandori, Max Kasy, Shengwu Li, Jonathan Libgober, Matthew Lilley, George Mailath, Eric Maskin, Xiaosheng Mu, Andy Newman, Joshua Schwartzstein, Roberto Serrano, Philipp Strack, Elie Tamer, Michael Thaler, Linh T. Tô, Maria Voronina, and various seminar participants for their insightful comments.

1 Introduction

The gambler's fallacy is widespread. Many people believe that a fair coin has a higher chance of landing on tails after landing on heads three times in a row, think a son is "due" to a woman who has given birth to consecutive daughters, and, in general, expect too much reversal from sequential realizations of independent random events. Studies have documented the gambler's fallacy in lottery games where the bias is strictly costly (Terrell, 1994; Suetens, Galbo-Jørgensen, and Tyran, 2016) and in incentivized lab experiments (Benjamin, Moore, and Rabin, 2017). Recent analysis of field data by Chen, Moskowitz, and Shue (2016) shows this bias also affects experienced decision-makers in high-stakes decisions, such as judges in asylum courts. Section 1.3 surveys more of this empirical literature.

This paper highlights novel implications of the gambler's fallacy in optimal-stopping problems when agents are uncertain about the underlying distributions. As a running example, consider an HR manager (Alice) recruiting for a job opening, sequentially interviewing candidates. In deciding whether to hire a candidate, Alice needs to form a belief about the labor pool with regard to the distribution of potential future applicants should she keep the position open. She consults with other managers who have recruited for similar positions to learn about the distribution of talent in the labor pool, then decides on a stopping strategy for her own hiring problem. Suppose all managers believe in the gambler's fallacy — that is, they exaggerate how unlikely it is to get consecutive above-average or consecutive below-average applicants (relative to the labor pool mean). This error stems from the same psychology that leads people to exaggerate how unlikely it is to get consecutive heads or consecutive tails when tossing a fair coin. What are the implications of this bias for the managers' beliefs and behavior over time?

In this example and other natural optimal-stopping problems, agents tend to stop when early draws are deemed "good enough," leading to an asymmetric censoring of experience. When a manager discovers a very strong candidate early in the hiring cycle, she stops her recruitment efforts and future managers do not observe what alternative candidates she would have found for the same job opening with a longer search. This endogenous *censoring effect* interacts with the gambler's fallacy bias and leads to pessimistic inference about the labor pool. Suppose Alice's predecessors held the correct beliefs about the labor pool, and the qualities of different candidates are objectively independent. Predecessors with below-average early interviewees continue searching, but they are systematically surprised because their subsequent interviewees also turn out to be below-average half of the time, contrary to their (false) expectations of positive reversals after bad initial outcomes. When these managers communicate their disappointment to Alice, she becomes overly pessimistic about the labor pool. This pessimism informs Alice's stopping strategy and affects the kind of (censored) experience that she communicates to future managers in turn.

This paper examines the endogenous learning dynamics of a society of agents believing in the gambler's fallacy. All agents face a common stage game: an optimal-stopping problem with draws in different periods independently generated from fixed yet unknown distributions.[1] They take turns playing the stage game, with the game's outcome determining each agent's payoff. Agents are Bayesians, except for the statistical bias. That is, they start with a prior belief over a class of *feasible models* about the joint distribution of draws. Feasible models are Gaussian distributions indexed by different unconditional means of the draws (the *fundamentals*). Reflecting a mistaken belief in reversals, all feasible models specify the same negative correlation between draws. Biased agents dogmatically believe that worse earlier draws lead to better distributions of later draws, conditional on the fundamentals. Before playing her own stage game, each agent observes the stage-game histories of her predecessors, then applies Bayes' rule to update her beliefs about the fundamentals. This inference procedure amounts to misspecified Bayesian learning in the class of feasible models, a class that excludes the true draw-generating distribution.

I consider two social-learning environments. When agents play the stage game one at a time, the stochastic processes of their beliefs and behavior almost surely converge globally to a unique steady state in which agents are *over-pessimistic* about the fundamentals and *stop too early* relative to the objectively optimal strategy. This result formalizes the intuition about how the gambler's fallacy interacts with the censoring effect to produce pessimistic inference.

When agents arrive in large generations, with everyone in the same generation playing simultaneously, society converges to the same steady state as the previous environment. This large-generations model features deterministic learning dynamics and illustrates a positive-feedback cycle between distorted beliefs and distorted stopping strategies. More severely censored datasets lead to more pessimistic beliefs, while more pessimistic beliefs lead to earlier stopping and, as a consequence, heavier history censoring. Mapping back to the hiring example, suppose a firm appoints HR managers in cohorts. Upon arrival, each junior manager learns the recruiting experience of all previous managers. If managers in the first cohort start with correct beliefs about labor-market conditions, then the average hiring outcome monotonically deteriorates across all successive cohorts. After Alice and others in her cohort consult with their predecessors and end up with over-pessimistic inferences, their beliefs lead them to be less "choosy" when hiring and only keep searching if their first interviewees prove to be truly unsatisfactory. On average, managers in Alice's cohort who reject their

[1] Using panel data, Conlon, Pilossoph, Wiswall, and Zafar (2018) find that unemployed workers make inferences about the wage distribution in the labor market using previously received (and rejected) offers, an empirical example of agents using histories to learn about the draw-generating distributions of an optimal-stopping problem. Interestingly, their paper also documents overinference from small samples, suggesting workers exhibit a "law of small numbers" psychology that underlies the gambler's fallacy.

early applicants under these newly lowered acceptance standards become disappointed by the quality of their later interviewees, even given these managers' already pessimistic beliefs about the labor pool. This is because biased agents expect more positive reversal following worse early outcomes, so for a fixed realization of later interviewee's quality, managers experience greater disappointment following worse earlier applicants. The pessimism that managers in Alice's generation started with thus becomes amplified in the next generation, leading to a further lowering of acceptance thresholds and a further decrease in the average quality of the hired candidate.

The endogenous-data setting leads to novel comparative statics predictions about how the payoff parameters of the stage game affect learning outcomes under the gambler's fallacy. For instance, suppose managers become more impatient, incurring a larger waiting cost when they decide to continue searching. If data is exogenous or if agents are correctly specified, then learning outcomes are independent of the details of the decision problem. When agents believing in the gambler's fallacy learn from endogenously censored histories, however, lower patience in the stage game leads to more distorted long-run beliefs about the fundamentals. This result is another expression of the positive feedback between actions and beliefs. Impatient managers use lower acceptance thresholds, so they tend to be more disappointed by the lack of reversals in the search process compared to their more patient (and therefore choosier) counterparts. This implies that more impatient agents also become more pessimistic about the fundamentals, thus compounding their initial lowering of the cutoff threshold (due to impatience) with a further change in the same direction (due to beliefs).

Finally, I expand the set of feasible models and consider agents who are uncertain about both the means and variances of the draw-generating distributions. I show that in this joint estimation, agents make the same misinference about the means as in the baseline model. However, they exaggerate the variances in a way that depends on the censoring of histories in their dataset. In the hiring example, this exaggeration corresponds to Alice believing that applicants for different vacancies come from different labor pools that vary in average quality, when in reality applicants for all vacancies originate from the same pool with a fixed quality distribution. Alice's belief in vacancy-specific fixed effects helps her explain the experience of her predecessors who had consecutive below-average interviewees, reasoning that it must have been especially difficult to find good candidates for these particular job openings. The severity of censoring in Alice's dataset determines her belief about the variance in average candidate quality across different vacancies, a belief that also influences her stopping strategy. I derive two results that illustrate how this belief in *fictitious variation* interacts with endogenous learning. First, when the stage-game payoff function is convex in draws (such as when previously rejected candidates can be recalled with some probability in the sequential interviewing problem), the positive-feedback cycle of the baseline environment strengthens. This is because a more severely censored dataset not only makes future agents more pessimistic about the fundamentals due to the usual censoring effect, but also decreases their belief in fictitious variation. Due to the convexity of the optimal-stopping problem, both forces encourage earlier stopping, leading to even heavier data censoring in the future. Second, a society where agents are uncertain about the variances can end up with a different long-run belief about the means than another society where agents know the correct variances. This is despite the fact that agents in both societies would make the same (mis)inference about the means given the same dataset of histories.

While I focus on (misspecified) Bayesian agents estimating parameters of a Gaussian model, my main results remain robust to a range of alternative specifications. These include a non-Bayesian method-of-moments inference procedure and general distributional assumptions.

1.1 Key Contributions

This work contributes to two strands of literature: the behavioral economics literature on inference mistakes for biased learners, and the theoretical literature on the dynamics of misspecified endogenous learning.

As a contribution to behavioral economics, I highlight a novel channel of misinference for behavioral agents — the interaction between psychological bias and data censoring. In many natural environments, agents learn from censored data. The economics literature has recently focused on the learning implications of selection neglect in these settings, where agents act as if their dataset is not censored.[2] My work points out that other well-documented behavioral biases can also interact with data censoring in interesting ways, producing important and novel implications in these environments. Mislearning in my model stems precisely from this interaction, not from either censored data or the gambler's fallacy alone. If agents were correctly specified (i.e. they do not suffer from the statistical bias), then they would correctly learn the fundamentals even from a dataset of censored histories. On the other hand, consider an "uncensored" environment where agents observe what their predecessors would have drawn in each period of the optimal-stopping problem, regardless of the actual stopping decisions of these predecessors. In such a counterfactual scenario, even biased agents would learn the fundamentals correctly. The intuition is that the gambler's fallacy is a "symmetric" bias. The "asymmetric" outcome of over-pessimism only occurs when the bias interacts with an (endogenous) asymmetric censoring mechanism that tends to produce data containing negative streaks but not positive streaks.

[2] See, for example, Enke (2017) and Jehiel (2018).

As a theoretical contribution, I prove convergence of beliefs and behavior in a non-self-confirming misspecified setting with a continuum of states of the world. Economists study many kinds of misspecifications with the property that even the "best-fitting" feasible belief does not match data exactly — that is to say, no feasible self-confirming belief exists. The gambler's fallacy falls into this family, since all feasible models imply a negative correlation absent in the data. I analyze the stochastic processes of belief and behavior under this statistical bias, proving their global almost-sure convergence to a unique steady state. This is a technically challenging problem. For biased agents facing a large dataset of histories generated from a fixed censoring threshold, the "best-fitting" feasible belief[3] depends on the threshold. But in the environment where agents act in a sequence, histories of predecessors are generated one at a time based on their ex-ante random and correlated stopping strategies, which may start arbitrarily far away from the steady-state stopping behavior. In related work, Heidhues, Koszegi, and Strack (2018) study learning dynamics under a different bias: overconfidence about one's own ability. Despite being biased, agents in their setting always have some feasible belief that exactly rationalizes data, and so their learning steady-state is a self-confirming equilibrium. By contrast, the steady state in my paper is not self-confirming. In addition, I prove my convergence result in a setting with multiple dimensions of uncertainty (the distributional parameters for different periods of the stage game), whereas Heidhues, Koszegi, and Strack (2018) consider convergence of misspecified learning with one-dimensional uncertainty. Fudenberg, Romanyuk, and Strack (2017) study a continuous-time model of active learning under misspecification, but their learning problem has an even more restricted state space. The agent's belief is binary, that is to say her prior is supported on exactly two feasible models. In my setting, agents' prior belief about each distributional parameter is supported on a continuum of feasible values.

As another contribution to the theoretical literature on misspecified learning dynamics, my project studies a new mechanism of endogeneity: the censoring effect in a dynamic stage game. A dynamic stage game is essential for studying learning under the gambler's fallacy, a behavioral bias concerning the serial correlation of data. The censoring effect relies on the dynamic structure of the decision problem and has no analog in the static stage-games of Heidhues, Koszegi, and Strack (2018) and Fudenberg, Romanyuk, and Strack (2017). In my setting, the *type* of data that an agent generates depends on her beliefs. To understand the distinction from the existing literature, consider the classic paper in this area, Nyarko (1991), who studies a monopolist setting a price on each date and observing the resulting sales. No matter what action the monopolist takes, she observes the same *type* of data: quantity sold. Similarly, the agent in Fudenberg, Romanyuk, and Strack (2017) always observes payoffs

[3] More precisely, the feasible model whose implied history distribution has the minimum (though still positive) Kullback-Leibler divergence relative to the observed history distribution, given the history censoring threshold.

and the agent in Heidhues, Koszegi, and Strack (2018) always observes output levels, after any action. Endogenous learning in these other papers takes the form of agents attributing different meanings to the same type of data, when interpreted through the lenses of different actions. On the other hand, we may think of stage-game histories censored with different thresholds as different types of data about the fundamentals. The distinction is that these different types of data, by themselves, lead to different beliefs about the fundamentals for biased learners. Actions play no role in inference except to generate these different types of data, since the likelihood of a (feasible) history does not depend on the censoring threshold that produced it.

1.2 Other Related Theoretical Work

Rabin (2002) and Rabin and Vayanos (2010) are the first to study the inferential mistakes implied by the gambler's fallacy. Except for an example in Rabin (2002), discussed below, all such investigations focus on passive inference, whereby learners observe an exogenous information process. By contrast, I examine an endogenous learning setting where the actions of predecessors censor the dataset of future learners. This setting allows me to ask whether the feedback loop between learners' actions and biased beliefs will attenuate or exaggerate the distortions caused by the fallacy over the course of learning. In addition, relative to this existing literature, the present paper uniquely focuses on the dynamics of mislearning under the gambler's fallacy. I prove that the stochastic process of beliefs and behavior almost surely converges when biased agents act one at a time, and I trace out the exact trajectory of beliefs and behavior when agents act in generations.

Section 7 of Rabin (2002) discusses an example of endogenous learning under a finite-urn model of the gambler's fallacy. The nature of Rabin (2002)'s endogenous data, however, is unrelated to the censoring effect central to my paper.[4] In Appendix E, I modify that example to induce the censoring effect. I find a misinference result in his finite-urn model of the gambler's fallacy, similar to what I find in the continuous Gaussian model of this paper. This exercise shows the robustness of my results within different modeling frameworks of the same statistical bias.

My steady state corresponds to Esponda and Pouzo (2016)'s Berk-Nash equilibrium. Rather than focusing only on equilibrium analysis, however, I focus on non-equilibrium learning dynamics and prove global convergence. That is, in the environment with agents

[4]In Rabin (2002)'s example, biased agents (correctly) believe that the part of the data which is always observable is independent of the part of the data which is sometimes missing. However, what I term the "censoring effect" is about misinference resulting from agents wrongly believing in negative correlation between the early draws that are always observed and the later draws that may be censored, depending on the realizations of the early draws. Therefore, my central mechanism highlights a novel interaction between data censoring and the gambler's fallacy bias that is absent in the previous literature.

acting one at a time, society converges to the steady state for all prior beliefs satisfying regularity conditions. In the environment where agents act in large generations, society converges for all initial conditions of the first generation. The large-generations environment also allows me to study how the positive feedback between beliefs and stopping strategies leads to monotonic convergence.

Although my learning framework involves short-lived agents learning from predecessors' histories, the social-learning aspect of my framework is not central to the results. In fact, the environment where a sequence of short-lived agents acts one at a time is equivalent to the environment where a single long-lived agent plays the stage game repeatedly, myopically maximizing her expected payoff in each iteration of the stage game. In the growing literature on social learning with misspecified Bayesians (e.g., Eyster and Rabin (2010); Gaurino and Jehiel (2013); Bohren (2016); Bohren and Hauser (2018); Frick, Iijima, and Ishii (2018)), agents observe their predecessors' actions but make errors when inverting these actions to deduce said predecessors' information about the fundamentals. This kind of action inversion does not take place in my framework: later agents observe all the information that their predecessors have seen, so actions of predecessors are uninformative.

The econometrics literature has also studied data-generating processes with censoring — for example, the Tobit model and models of competing risks.[5] This literature has primarily focused on the issue of model identification from censored data (Cox, 1962; Tsiatis, 1975; Heckman and Honoré, 1989). In my setting, there is no identifiability problem for correctly specified agents, since censored histories can identify the mean and the covariance matrix of the draws. Instead, I study how agents make wrong inferences from censored data when they have a family of misspecified models. Another contrast is that the econometrics literature has focused on exogenous data-censoring mechanisms, but censoring is endogenous in my setting and depends on the beliefs of previous agents. As discussed before, this endogeneity is central to my results.

1.3 Empirical Evidence on the Gambler's Fallacy

Bar-Hillel and Wagenaar (1991) review classical psychology studies on the gambler's fallacy. The earliest lab evidence involves two types of tasks. In "production tasks," subjects are asked to write down sequences using a given alphabet, with the goal of generating sequences that resemble the realizations of an i.i.d. random process. Subjects tend to produce sequences with too many alternations between symbols, as they attempt to locally balance out symbol frequencies. In "judgment tasks" where subjects are asked to identify which sequence of binary symbols appears most like consecutive tosses of a fair coin, subjects find sequences

[5] References can be found in Amemiya (1985) and Crowder (2001).

with an alternation probability of 0.6 more random than those with an alternation probability of 0.5. While most of these studies are unincentivized, Benjamin, Moore, and Rabin (2017) have found the gambler's fallacy with strict monetary incentives, where a bet on a fair coin continuing its streak pays strictly more than the bet on the streak reversing. Barron and Leider (2010) have shown that experiencing a streak of binary outcomes one at a time exacerbates the gambler's fallacy, compared with simply being told the past sequence of outcomes all at once.

Other studies have identified the gambler's fallacy using field data on lotteries and casino games. Unlike in experiments, agents in field settings are typically not explicitly told the underlying probabilities of the randomization devices. In state lotteries, players tend to avoid betting on numbers that have very recently won. This under-betting behavior is strictly costly for the players when lotteries have a pari-mutuel payout structure (as in the studies of Terrell (1994) and Suetens, Galbo-Jørgensen, and Tyran (2016)), since it leads to a larger-than-average payout per winner in the event that the same number is drawn again the following week. Using security video footage, Croson and Sundali (2005) show that roulette gamblers in casinos bet more on a color after a long streak of the opposite color. Narayanan and Manchanda (2012) use individual-level data tracked using casino loyalty cards to find that a larger recent win has a negative effect on the next bet that the gambler places, while a larger recent loss increases the size of the next bet. Finally, using field data from asylum judges, loan officers, and baseball umpires, Chen, Moskowitz, and Shue (2016) show that even very experienced decision-makers show a tendency to alternate between two decisions across a sequence of randomly ordered decision problems. This can be explained by the gambler's fallacy, as the fallacy leads to the belief that the objectively "correct" decision is negatively auto-correlated across a sequence of decision problems. The authors rule out several other explanations, including contrast effect and quotas.

As Rabin (2002) and Rabin and Vayanos (2010) have argued, someone who dogmatically believes in the gambler's fallacy must attribute the lack of reversals in the data to the fundamental probabilities of the randomizing device, leading to overinference from small samples. This overinference can be seen in the field data. Cumulative win/loss (as opposed to very recent win/loss) on a casino trip is positively correlated with the size of future bets (Narayanan and Manchanda, 2012). A player who believes in the gambler's fallacy rationalizes his persistent good luck on a particular day by thinking he must be in a "hot" state, where his fundamental probability of winning in each game is higher than usual. In a similar vein, a number that has been drawn more often in the past six weeks, excluding the most recent past week, gets more bets in the Denmark lottery (Suetens, Galbo-Jørgensen, and Tyran, 2016). This kind of overinference resulting from small samples persists even in a market setting where participants have had several rounds of experience and feedback (Camerer, 1987). In line with these studies, the model I consider involves agents who dogmatically believe in the gambler's fallacy and misinfer some parameter of the world as a result — though the misinference mechanism in my model is further complicated by the presence of endogenous data censoring.

2 Overview

This section presents the basic elements of the model, previews my main results, and provides intuition for how the censoring effect drives my conclusions. I describe a class of optimal-stopping problems serving as the (single-player) *stage game*. Agents are uncertain about the distribution of draws in the stage game. They entertain a prior belief over a family of distributions that they find plausible, the *feasible models* of how draws are generated. All feasible models specify the same negative correlation between the draws, even though draws are objectively independent: an error reflecting the gambler's fallacy. Sections 3 and 4 embed these model elements into social-learning environments. In each environment, a society of agents takes turns playing the stage game, making inferences over feasible models using others' stage-game histories. Section 5 contains a number of extensions that verify the robustness of my main results with regard to different specifications.

2.1 Basic Elements of the Model

2.1.1 Optimal-Stopping Problem as a Dynamic Stage Game

The stage game is a two-period optimal-stopping problem. In the first period, the agent draws $x_1 \in \mathbb{R}$ and decides whether to stop. If she stops at x_1, her payoff is $u_1(x_1)$ and the stage game ends. Otherwise, she continues to the second period, where she draws $x_2 \in \mathbb{R}$. The stage game then ends with the agent getting payoffs $u_2(x_1, x_2)$.

The payoff functions $u_1 : \mathbb{R} \to \mathbb{R}$ and $u_2 : \mathbb{R}^2 \to \mathbb{R}$ satisfy some regularity conditions to be introduced in Assumption 1. The following example satisfies Assumption 1 and will be used to illustrate my results throughout this paper.

Example 1 (search with q probability of recall)**.** Many industries have a regular hiring cycle each year. Consider a firm in such an industry and its HR manager, who must fill a job opening during this year's cycle. In the early phase of the hiring cycle, she finds a candidate who would bring net benefit x_1 to the firm if hired. She must decide between hiring this candidate immediately or waiting. Choosing to wait means she will continue searching in the late phase of the hiring cycle, finding another candidate who would bring benefit x_2 to the organization. Waiting, however, carries the risk that the early candidate accepts an offer from a different firm in the interim. Suppose there is $0 \leq q < 1$ probability that the early candidate will remain available in the late hiring phase. This situation then has payoff functions $u_1(x_1) = x_1$ and $u_2(x_1, x_2) = q \cdot \max(x_1, x_2) + (1-q)x_2$. That is, in the late phase, there is q probability the manager gets payoff equal to the higher of the two candidates' qualities, and $1-q$ probability that she only has the option to hire the second candidate.

I now present regularity conditions on the payoff functions that define the class of optimal-stopping problems I study.

Assumption 1 (regularity conditions)**.** *The payoff functions satisfy :*

(a) For $x_1^{'} > x_1^{''}$ and $x_2^{'} > x_2^{''}$, $u_1(x_1^{'}) > u_1(x_1^{''})$ and $u_2(x_1^{'}, x_2^{'}) > u_2(x_1^{'}, x_2^{''})$.

(b) For $x_1^{'} > x_1^{''}$ and any $\bar{x}_2$, $u_1(x_1^{'}) - u_1(x_1^{''}) > |u_2(x_1^{'}, \bar{x}_2) - u_2(x_1^{''}, \bar{x}_2)|$.

(c) There exist $x_1^g, x_2^b, x_1^b, x_2^g \in \mathbb{R}$ so that $u_1(x_1^g) - u_2(x_1^g, x_2^b) > 0$, while $u_1(x_1^b) - u_2(x_1^b, x_2^g) < 0$.

(d) u_1, u_2 are continuous. Also, for any $\bar{x}_1 \in \mathbb{R}$, $x_2 \mapsto u_2(\bar{x}_1, x_2)$ is absolutely integrable with respect to any Gaussian distribution on $\mathbb{R}$.

Assumption 1(a) says u_1, u_2 are strictly increasing in the draws in their respective periods. Assumption 1(b) says a higher realization of the early draw increases first-period payoff more than it changes second-period payoff. Under Assumption 1(a), Assumption 1(b) is satisfied whenever u_2 is not a function of x_1, as in optimal-stopping problems where stopping in period k gives payoff only depending on the k-th draw. More generally, Assumption 1(b) is satisfied when $u_2(x_1, x_2) = z_{2,1}(x_1) + z_{2,2}(x_2)$ is separable across the draws of the two periods with $|z_{2,1}^{'}(x_1)| < u_1^{'}(x_1)$ at all $x_1 \in \mathbb{R}$. Assumption 1(c) says there exist a good enough realization x_1^g and bad enough realization x_2^b, so that the agent prefers stopping in period 1 after x_1^g than continuing when she knows for sure that her second draw will be x_2^b. Conversely, there are x_1^b, x_2^g so that she prefers continuing after x_1^b if she knows she will get x_2^g in the second period for sure. Assumption 1(d) is a technical condition. The absolute integrability requirement ensures that the expected payoff from choosing to continue is always well-defined. These conditions are satisfied by my recurrent example.[6]

Claim 1. Example 1 satisfies Assumption 1.

I now define strategies and histories of the stage game.

Definition 1. A *strategy* is a function $S : \mathbb{R} \to$ {Stop, Continue} that maps the realization of the first-period draw $X_1 = x_1$ into a stopping decision.

Without loss I only consider pure strategies, because there always exists a payoff-maximizing pure strategy under any belief about the distribution of draws.

[6] Omitted proofs from the main text can be found in Appendix A.

Definition 2. The *history* of the stage game is an element $h \in \mathbb{H} := \mathbb{R} \times (\mathbb{R} \cup \{\varnothing\})$. If an agent decides to stop after $X_1 = x_1$, her history is $(x_1, \varnothing)$. If the agent continues after $X_1 = x_1$ and draws $X_2 = x_2$ in the second period, her history is (x_1, x_2).

The symbol $\varnothing$ is a *censoring indicator*, emphasizing that the hypothetical second-period draw is unobserved when an agent does not continue into the second period. In Example 1, if the HR manager hires the first candidate, she stops her recruitment efforts early and the counterfactual second candidate that she would have found had she kept the position open remains unknown.

2.1.2 Feasible Models and the Objective Model

Objectively, draws X_1, X_2 in the stage game are independently distributed with Gaussian distributions $X_1 \sim \mathcal{N}(\mu_1^\bullet, \sigma^2)$ and $X_2 \sim \mathcal{N}(\mu_2^\bullet, \sigma^2)$ for some $\sigma^2 > 0$. The parameters $\mu_1^\bullet, \mu_2^\bullet \in \mathbb{R}$ are fixed and called *true fundamentals*. In Example 1, $\mu_1^\bullet$ and $\mu_2^\bullet$ stand for the underlying qualities of the two applicant pools in the early and late phases of the hiring season.

Agents are uncertain about the distribution of (X_1, X_2). The next definition provides a language to discuss the set of distributions that a gambler's fallacy agent deems plausible.

Definition 3. The set of *feasible models* $\{\Psi(\mu_1, \mu_2; \gamma) : (\mu_1, \mu_2) \in \mathcal{M}\}$ is a family of joint distributions of (X_1, X_2) indexed by *feasible fundamentals* $(\mu_1, \mu_2) \in \mathcal{M} \subseteq \mathbb{R}^2$, for some *bias parameter* $\gamma > 0$. Here $\Psi(\mu_1, \mu_2; \gamma)$ refers to the *subjective model*

$$\begin{aligned} X_1 &\sim \mathcal{N}(\mu_1, \sigma^2) \\ (X_2 | X_1 = x_1) &\sim \mathcal{N}(\mu_2 - \gamma(x_1 - \mu_1), \sigma^2), \end{aligned}$$

where $X_2 | (X_1 = x_1)$ is the conditional distribution of X_2 given $X_1 = x_1$.

Every feasible model has the property that $\mathbb{E}[X_2 \mid X_1 = x_1]$ decreases in x_1, which reflects the gambler's fallacy. Conditional on the fundamentals, if the realization of X_1 is higher than expected, then the agent believes bad luck is due in the near future and the second draw is likely below average.[7] Conversely, an exceptionally bad early draw likely portends above-average luck in the next period. This expected luck reversal is more obvious in the following

[7] I study gambler's fallacy for continuous random variables, where the magnitude of X_1 affects the agent's prediction about X_2. Chen, Moskowitz, and Shue (2016)'s analysis of baseball umpire data provides support for the continuous version of the statistical bias. They find that an umpire is more likely to call the current pitch a ball after having called the previous pitch a strike, controlling for the actual location of the pitch. Crucially, the effect size is larger after more obvious strikes, where "obviousness" is based on the distance of the pitch to the center of the regulated strike zone. This distance can be thought of as a continuous measure of the "quality" of each pitch.

equivalent formulation of $\Psi(\mu_1, \mu_2; \gamma)$:

$$X_1 = \mu_1 + \epsilon_1$$
$$X_2 = \mu_2 + \epsilon_2$$

where $\epsilon_1 \sim \mathcal{N}(0, \sigma^2)$ and $\epsilon_2 | \epsilon_1 \sim \mathcal{N}(-\gamma\epsilon_1, \sigma^2)$. The zero-mean terms ϵ_1, ϵ_2 represent the idiosyncratic factors, or "luck," that determine how the realizations of X_1 and X_2 deviate from their unconditional means μ_1 and μ_2 in the model $\Psi(\mu_1, \mu_2; \gamma)$. The subjective model stipulates reversal of luck, since ϵ_1, ϵ_2 are negatively correlated. Larger $\gamma > 0$ implies greater magnitude in these expected reversals and thus more bias.

The set of feasible models is indexed by the set of feasible fundamentals, which correspond to the unconditional means[8] of X_1 and X_2. Therefore, the agent's prior belief over the feasible models is given by a prior belief supported on the feasible fundamentals.

Remark 1. I will consider several specifications of $\mathcal{M}$ in this paper. I list them here and provide interpretations below.

(a) $\mathcal{M} = \mathbb{R}^2$. The agent thinks all values $(\mu_1, \mu_2) \in \mathbb{R}^2$ are possible.

(b) $\mathcal{M} = \lozenge$, where $\lozenge$ is a bounded parallelogram in $\mathbb{R}^2$ whose left and right edges are parallel to the y-axis, whose top and bottom edges have slope $-\gamma$. The agent is uncertain about both μ_1 and μ_2, but her uncertainty has bounded support.[9]

(c) $\mathcal{M} = \{\mu_1^\bullet\} \times [\underline{\mu}_2, \bar{\mu}_2]$. The agent has a correct, dogmatic belief about μ_1, but has uncertainty about μ_2 supported on a bounded interval.

(d) $\mathcal{M} = \{(\mu, \mu) : \mu \in \mathbb{R}\}$. The agent is convinced that the first-period and second-period fundamentals are the same, but is uncertain what this common parameter is.

While the agent can freely update her belief about the fundamentals on $\mathcal{M}$, she holds a dogmatic belief about $\gamma > 0$.[10] This implies that the set of feasible models excludes the true model, $\Psi^\bullet = \Psi(\mu_1^\bullet, \mu_2^\bullet; 0)$, so the support of the agent's prior belief is misspecified. I maintain this misspecification to match the field evidence of Chen, Moskowitz, and Shue (2016), where even very experienced decision-makers continue to exhibit a non-negligible amount of the gambler's fallacy in high-stakes settings. Another reason why agents may never question their misspecified prior is that the misspecification is "attentionally stable" in the sense of

[8]Section 5.2 discusses the extension where agents are also uncertain about the variances and jointly estimate means and variances from censored histories.

[9]Any prior belief over fundamentals (μ_1, μ_2) supported on a bounded set in $\mathbb{R}^2$ can be arbitrarily well-approximated by a prior belief over a large enough $\lozenge$.

[10]Section 5.3 studies the extension where agents are uncertain about γ, but the support of their prior belief about γ lies to the left of 0 and is bounded away from it.

Gagnon-Bartsch, Rabin, and Schwartzstein (2018). Under the theory that the true model falls within the feasible models, an agent finds it harmless to coarsen her dataset by only paying attention to certain "summary statistics." In large datasets, the statistics extracted by the limited-attention agent do not lead her to question the validity of her theory. I discuss this further in Appendix F.

I write $\mathbb{E}_\Psi$ and $\mathbb{P}_\Psi$ throughout for expectation and probability with respect to the subjective model Ψ. When $\mathbb{E}$ and $\mathbb{P}$ are used without subscripts, they refer to expectation and probability under the true model $\Psi^\bullet$.

Before stating my main results, I first establish a result about the optimal stage-game strategy under any feasible model, which will motivate a slight strengthening of Assumption 1 that I need for some results. For $c \in \mathbb{R}$, write S_c for the *cutoff strategy* such that $S_c(x_1) =$ Stop if and only if $x_1 > c$. That is, S_c accepts all early draws above a cutoff threshold c.

Proposition 1. *Under Assumption 1 and for $\gamma > 0$,*

- *Under each subjective model $\Psi(\mu_1, \mu_2; \gamma)$, there exists a cutoff threshold $C(\mu_1, \mu_2; \gamma) \in \mathbb{R}$ such that it is strictly optimal to continue whenever $x_1 < C(\mu_1, \mu_2; \gamma)$ and strictly optimal to stop whenever $x_1 > C(\mu_1, \mu_2; \gamma)$.*
- *For every $\mu_1 \in \mathbb{R}$, $\mu_2 \mapsto C(\mu_1, \mu_2; \gamma)$ is strictly increasing.*
- *For every $\mu_1 \in \mathbb{R}$, $\mu_2 \mapsto C(\mu_1, \mu_2; \gamma)$ is Lipschitz continuous with Lipschitz constant $1/\gamma$.*

The content of this lemma is threefold.

First, it shows that the best strategy for the class of optimal-stopping problems I study takes a cutoff form, regardless of the underlying distributions. This is because a higher x_1 both increases the payoff to stopping and, under the gambler's fallacy, predicts worse draws in the next period. Both forces push in the direction of stopping. The optimality of cutoff strategies leads to an endogenous, asymmetric censoring of histories, formalizing the idea that agents stop after "good enough" draws.

Second, holding fixed μ_1, the cutoff threshold is higher when μ_2 is higher. In other words, the definition of a "good enough" early draw x_1 increases with μ_2. This is because agents can afford to be choosier in the first period when facing improved prospects in the second period.

The third statement about Lipschitz continuity, on the other hand, gives a bound on how quickly $\mu_2 \mapsto C(\mu_1, \mu_2; \gamma)$ increases. To understand why it holds, suppose that one agent believes draws are generated according to $\Psi(\mu_1, \mu_2; \gamma)$, while another agent believes they are generated according to $\Psi(\mu_1, \mu_2 + 1; \gamma)$. Under any feasible model, when X_1 increases by $1/\gamma$, the predicted conditional mean of X_2 falls by $(1/\gamma) \cdot \gamma = 1$. Therefore, the indifference condition of the first agent at cutoff c implies the second agent prefers stopping after $X = c + \frac{1}{\gamma}$, since the expected reversal cancels out the relative optimism of the second agent about the unconditional distribution of X_2.

The Lipschitz constant $1/\gamma$ is guaranteed for every optimal-stopping problem satisfying Assumption 1 and for every $\gamma > 0$. But, it may not be the best Lipschitz constant. My results use the slightly stronger condition that $\mu_2 \mapsto C(\mu_1^\bullet, \mu_2; \gamma)$ has a Lipschitz constant *strictly* smaller than $1/\gamma$. Intuitively this should be easy to satisfy, but instead of assuming it directly, I consider the following condition on primitives that implies the desired infinitesimally stronger Lipschitz continuity. It is a joint restriction on γ and the stage game.

Assumption 2 (ℓ-Lipschitz continuity)**.** *There exists* $0 < \ell < \frac{1}{\gamma}$ *so that for every* $x_1, x_2 \in \mathbb{R}$ *and* $d > 0$,

$$u_1(x_1 + \ell d) - u_1(x_1) \geq u_2(x_1 + \ell d, x_2 + (1 - \gamma\ell)d) - u_2(x_1, x_2)$$

This condition is satisfied for search with q probability of recall.

Claim 2. Example 1 satisfies Assumption 2 with $\ell = \frac{1}{1+\gamma}$ for every probability of recall $0 \leq q < 1$ and every bias $\gamma > 0$.

Assumption 2 strengthens Assumption 1(b), which already implies $u_1(x_1 + \ell d) - u_1(x_1) > u_2(x_1 + \ell d, x_2) - u_2(x_1, x_2)$. For any $0 < \ell < \frac{1}{\gamma}$, $(1 - \gamma\ell)d > 0$, which makes the inequality harder to satisfy as adding a positive term to the second argument of u_2 makes the RHS larger.

2.2 Main Results

I now state my two main results, which concern learning dynamics under the gambler's fallacy in two different social-learning environments. I defer precise details of these environments to later sections.

In the first environment, short-lived agents arrive one per round, $t = 1, 2, 3, ...$ All agents start with the same full-support prior density $g : \Diamond \to \mathbb{R}_{>0}$, where $\Diamond$ is a bounded parallelogram in $\mathbb{R}^2$ as in Remark 1(b). Agent t observes the stage-game histories of all predecessors, updates her prior g to a posterior density $\tilde{g}_t$, then chooses a cutoff threshold $\tilde{C}_t$ to maximize her expected payoff based on this posterior belief. In this environment, the sequences of cutoffs $(\tilde{C}_t)$ and posterior beliefs $(\tilde{g}_t)$ are stochastic processes whose randomness derives from the randomness of draws. Draws are objectively independent, both between the two periods in the same round of the stage game and across different rounds.

Theorem 1. *Suppose Assumptions 1 and 2 hold and* $\frac{\partial g}{\partial \mu_1}$ *and* $\frac{\partial g}{\partial \mu_2}$ *are continuous on* $\Diamond$*. There exists a unique steady state* $\mu_2^\infty, c^\infty \in \mathbb{R}$ *not dependent on* g*, so that provided* $(\mu_1^\bullet, \mu_2^\infty) \in \Diamond$, *almost surely* $\lim_{t\to\infty} \tilde{C}_t = c^\infty$ *and* $\lim_{t\to\infty} \mathbb{E}_{(\mu_1,\mu_2)\sim\tilde{g}t}[|\mu_1 - \mu_1^\bullet| + |\mu_2 - \mu_2^\infty|] = 0$. *The steady state satisfies* $\mu_2^\infty < \mu_2^\bullet$ *and* $c^\infty < c^\bullet$, *where* $c^\bullet$ *is the objectively optimal cutoff threshold.*

In other words, almost surely behavior and belief converge in the society, and this steady state is independent of the prior over fundamentals (provided its support is large enough). In the steady state, agents hold overly pessimistic beliefs about the fundamentals and stop too early, relative to the objectively optimal strategy.

In the second environment, short-lived agents arrive in generations, $t = 0, 1, 2, ...$, with a continuum of agents per generation. Agents' prior belief about the fundamentals is given by a full-support density on $\mathbb{R}^2$, as in Remark 1(a). Each agent observes the stage-game histories of all predecessors from all past generations to make inferences about the fundamentals. Due to the large generations, cutoffs and beliefs are deterministic in generations $t \geq 1$, which I denote as $c_{[t]}$ and $\mu_{[t]} = (\mu_{1,[t]}, \mu_{2,[t]})$ respectively. The society is initialized at an arbitrary cutoff strategy $S_{c_{[0]}}$ in the 0th generation, the *initial condition*.

Theorem 2. *Suppose Assumption 1 holds. Starting from any initial condition and any* g, *cutoffs* $(c_{[t]})_{t\geq1}$ *and beliefs* $(\mu_{[t]})_{t\geq1}$ *form monotonic sequences across generations. When Assumption 2 also holds, there exists a unique steady state* $\mu_2^\infty, c^\infty \in \mathbb{R}$ *so that* $c_{[t]} \to c^\infty$ *and* $(\mu_{1,[t]}, \mu_{2,[t]}) \to (\mu_1^\bullet, \mu_2^\infty)$ *monotonically, regardless of the initial condition and* g. *These steady states are the same as those in Theorem 1.*

The monotonicity of beliefs and cutoffs across generations reflects a positive-feedback loop between changes in beliefs and changes in behavior. Suppose generation t is more pessimistic than generation $t-1$ about the second-period fundamental, $\mu_{2,[t]} < \mu_{2,[t-1]}$. The monotonicity result implies that beliefs move in the same direction again in generation $t+1$, that is $\mu_{2,[t+1]} < \mu_{2,[t]}$. The information of generation $t+1$ differs from that of generation t only in that agents in generation $t+1$ observe all stage-game histories of generation t. This means generation t's stopping behavior differs from that of generation $t-1$ in such a way as to generate histories that amplify, not dampen, the initial change in beliefs from generation $t-1$ to generation t.

2.3 Intuition for the Main Results

In the learning environments I study, each agent censors the data of future agents through her stopping strategy, where the strategy choice depends on her beliefs. To build intuition for how this censoring effect relates to the two main theorems, I first consider a biased agent with feasible fundamentals $\mathcal{M} = \mathbb{R}^2$, facing a large sample of histories all censored according to some cutoff threshold $c \in \mathbb{R}$. I characterize her inference about fundamentals when the sample size grows and analyze how her inference depends on the cutoff threshold c.

For a cutoff strategy S_c and a subjective model Ψ, $\mathcal{H}(\Psi; c)$ refers to the distribution of histories when draws are generated by Ψ and histories censored according to S_c, or more precisely:

Definition 4. For $c \in \mathbb{R}$ and Ψ a subjective model, $\mathcal{H}(\Psi; c) \in \Delta(\mathbb{H})$ is the distribution of histories given by

$$\begin{aligned} \mathcal{H}(\Psi; c)[E_1 \times E_2] &:= \quad \mathbb{P}_\Psi\left[(E_1 \cap (c, \infty)) \times E_2\right] \quad \text{for } E_1, E_2 \in \mathcal{B}(\mathbb{R}) \\ \mathcal{H}(\Psi; c)[E_1 \times \{\varnothing\}] &:= \quad \mathbb{P}_\Psi[(E_1 \cap (-\infty, c]) \times \mathbb{R}] \quad \text{for } E_1 \in \mathcal{B}(\mathbb{R}), \end{aligned}$$

where $\mathcal{B}(\mathbb{R})$ is the collection of Borel subsets of $\mathbb{R}$.

I abbreviate $\mathcal{H}(\Psi^\bullet; c)$ as simply $\mathcal{H}^\bullet(c)$, the true distribution of histories under the true model of draws and the cutoff threshold c. The next definition gives a measure of how much the implied distribution of histories under the feasible model with fundamentals (μ_1, μ_2) differs from the true distribution of histories, both generated with the same censoring.

Definition 5. For $c, \mu_1, \mu_2 \in \mathbb{R}$, the *Kullback-Leibler (KL) divergence* from $\mathcal{H}^\bullet(c)$ to $\mathcal{H}(\Psi(\mu_1, \mu_2; \gamma); c)$, denoted by $D_{KL}(\mathcal{H}^\bullet(c) \mid\mid \mathcal{H}(\Psi(\mu_1, \mu_2; \gamma); c)\ )$, is

$$\int_c^\infty \phi(x_1; \mu_1^\bullet, \sigma^2) \cdot \ln\left(\frac{\phi(x_1; \mu_1^\bullet, \sigma^2)}{\phi(x_1; \mu_1, \sigma^2)}\right) dx_1$$
$$+ \int_{-\infty}^c \left\{ \int_{-\infty}^\infty \phi(x_1; \mu_1^\bullet, \sigma^2) \cdot \phi(x_2; \mu_2^\bullet, \sigma^2) \cdot \ln\left[\frac{\phi(x_1; \mu_1^\bullet, \sigma^2) \cdot \phi(x_2; \mu_2^\bullet, \sigma^2)}{\phi(x_1; \mu_1, \sigma^2) \cdot \phi(x_2; \mu_2 - \gamma(x_1 - \mu_1), \sigma^2)}\right] dx_2 \right\} dx_1$$

where $\phi(x; a, b^2)$ is Gaussian density with mean a and variance b^2.

The minimizers of KL divergence with respect to cutoff c,

$$(\mu_1^*, \mu_2^*) \in \underset{\mu_1, \mu_2 \in \mathbb{R}}{\arg\min}\ D_{KL}(\mathcal{H}^\bullet(c) \mid\mid \mathcal{H}(\Psi(\mu_1, \mu_2; \gamma); c)\ ),$$

are called the *pseudo-true fundamentals* with respect to c.

To interpret, the likelihood of the history $h = (x_1, x_2)$ with $x_1 \le c$ is $\phi(x_1; \mu_1^\bullet, \sigma^2) \cdot \phi(x_2; \mu_2^\bullet, \sigma^2)$ under the true model $\Psi^\bullet$, $\phi(x_1; \mu_1, \sigma^2) \cdot \phi(x_2; \mu_2 - \gamma(x_1 - \mu_1), \sigma^2)$ under the feasible model $\Psi(\mu_1, \mu_2; \gamma)$. The likelihood of the history $h = (x_1, \varnothing)$ with $x_1 > c$ is $\phi(x_1; \mu_1^\bullet, \sigma^2)$ under the true model, $\phi(x_1; \mu_1, \sigma^2)$ under the feasible model. The likelihoods of all other histories are 0 under both models. So the KL divergence expression in Definition 5 is the expected log-likelihood ratio of the history under the true model versus under the feasible model with fundamentals (μ_1, μ_2), where expectation over histories is taken under the true model. In general, this optimization objective depends on the cutoff threshold c that determines how histories are censored. I will therefore denote the pseudo-true fundamentals as $\mu_1^*(c)$, $\mu_2^*(c)$ to emphasize this dependence.

The pseudo-true fundamentals correspond to the biased agent's inference about the fundamentals in large samples (hence the name). More precisely, suppose an agent starts with a prior belief on the fundamentals supported on $\mathbb{R}^2$ and observes a large but finite sample of histories drawn from $\mathcal{H}^\bullet(c)$. Proposition A.1 in Appendix B shows that as the sample size grows, her posterior belief almost surely converges in L^1 to the point mass on $(\mu_1^*(c), \mu_2^*(c))$.

The next proposition explicitly solves the pseudo-true fundamentals in a simple closed-form expression.[11] This result makes much of the later analysis tractable and contains the key intuition behind the two main theorems.

Proposition 2. *For $c \in \mathbb{R}$, the pseudo-true fundamentals are $\mu_1^*(c) = \mu_1^\bullet$ and*

$$\mu_2^*(c) = \mu_2^\bullet - \gamma\left(\mu_1^\bullet - \mathbb{E}\left[X_1 \mid X_1 \leq c\right]\right).$$

So $\mu_2^*(c) < \mu_2^\bullet$ for all $c \in \mathbb{R}$ and $\mu_2^*(c)$ strictly increases in c.

The directional data censoring where histories only contain X_2 following low values of X_1 leads to over-pessimism, $\mu_2^*(c) < \mu_2^\bullet$ for all c. In every feasible model of draws $\Psi(\mu_1, \mu_2; \gamma)$, the realization of X_2 depends on two factors: the second-period fundamental μ_2, and a reversal effect based on the realization of X_1. Under the correct or over-optimistic belief about μ_2, a biased agent would be systematically disappointed by realizations of X_2 in her dataset. This is because X_2 is only uncensored when X_1 is low enough, a contingency where the agent expects positive reversal on average.[12] Over-pessimism can therefore be thought of as "two wrongs making a right," as the biased agent's pessimism about the unconditional X_2 mean counteracts her false expectation of positive reversals in the dataset of censored histories.

This mechanism explains the long-run pessimism in Theorem 1 and Theorem 2. In fact, in the large-generations setting of Theorem 2, every generation $t \geq 1$ holds strictly

[11] This result shows the pseudo-true fundamentals have a method-of-moments interpretation. Suppose that instead of minimizing KL divergence, agents find $\mu_1^M, \mu_2^M \in \mathbb{R}$ so that $\mathcal{H}(\Psi(\mu_1^M, \mu_2^M; \gamma); c)$ matches $\mathcal{H}^\bullet(c)$ in terms of two moments: the means of the first- and second-period draws in the distribution of censored histories. We can show that in fact, $\mu_1^M(c) = \mu_1^*(c)$ and $\mu_2^M(c) = \mu_2^*(c)$ for all $c \in \mathbb{R}$. This provides an alternative, non-Bayesian foundation for agents' inference behavior. In Appendix C, I study the large-generations learning dynamics for agents who apply this kind of method-of-moments inference to a family of general, non-Gaussian feasible models of draws.

[12] This intuition presumes that agents understand selection in the dataset. Selection neglect is unlikely in this environment due to the salience of censoring. In large datasets, agents observe both censored histories with length 1 and uncensored histories with length 2, so the presence of censoring is highly explicit in the data. By contrast, both intuition about selection neglect and experiments documenting it (e.g., Enke (2017)) have focused on settings where the dataset does not contain "reminders" about censoring and could be reasonably mistaken as a dataset without selection. Interestingly, Enke (2017) finds that the simple hint "*Also think about the players whom you do not communicate with!*" reduces the fraction of selection neglecters by 40%. This suggests the salience of censoring in my setting should mitigate selection neglect even further. In Online Appendix OA 3.2, I show that the presence of a fraction of selection neglecters in the population moderates the pessimism of the baseline gambler's fallacy agents, but does not eliminate it.

pessimistic beliefs, so over-pessimism is also a short-run phenomenon provided there are enough predecessors per generation. The idea that asymmetric data censoring combined with the gambler's fallacy leads to pessimistic inference is highly robust. It continues to hold when the feasible fundamentals reflect agents' knowledge that $\mu_1 = \mu_2$ as in Remark 1(d) (Section 5.4), when agents are uncertain about variances (Section 5.2), under a joint relaxation of Bayesian inference and Gaussian models (Appendix C), when the stage game has more than two periods (Appendix D), under additional behavioral biases in inference (Online Appendices OA 3.2 and OA 3.3), when higher draws bring worse payoffs (Online Appendix OA 3.1), and with high probability after observing a finite dataset containing just 100 censored histories (Online Appendix OA 4.1).

Not only are the pseudo-true fundamentals always too pessimistic, the severity of censoring also increases pessimism. To understand the intuition, consider two datasets of histories from the distributions $\mathcal{H}^{\bullet}(c^{'})$ and $\mathcal{H}^{\bullet}(c^{''})$, where $c^{''} < c^{'}$. The acceptance threshold is lower in the second dataset, implying that uncensored values of X_2 are preceded by worse values of X_1 there, as the second draw is only observed when the first draw falls below the threshold. A biased agent expects a greater amount of positive reversal in the dataset with distribution $\mathcal{H}^{\bullet}(c^{''})$ than the one with distribution $\mathcal{H}^{\bullet}(c^{'})$, but in truth uncensored X_2 has the same distribution in both datasets, since X_1 and X_2 are objectively independent. For a fixed realization $X_2 = x_2$, an agent is more disappointed when she expects more positive reversal, so inference about μ_2 is pessimistic in the more heavily censored dataset.

The comparative static $\frac{d\mu_2^*}{dc} > 0$ is central to the positive-feedback loop from Theorem 2. In the large-generations model, Generation 1 observes a large dataset of histories drawn from $\mathcal{H}^{\bullet}(c_{[0]})$ and chooses a cutoff $c_{[1]}$. Generation 2 then observes histories from all predecessor generations, that is histories drawn from both $\mathcal{H}^{\bullet}(c_{[0]})$ and $\mathcal{H}^{\bullet}(c_{[1]})$. If $c_{[1]} < c_{[0]}$, then Generation 2's dataset features (on average) more severe censoring than Generation 1's dataset. Thus, Generation 2 comes to a more pessimistic inference about the second-period fundamental. By Proposition 1, this leads to a further lowering of the cutoff threshold, $c_{[2]} < c_{[1]}$, and the pattern continues.

3 Convergence, Over-Pessimism, and Early Stopping

In this section, I study a social-learning environment where biased agents act one at a time, inferring fundamentals from predecessors' histories. I begin by defining the *steady state* of the stage game for biased agents. The steady state depends on the optimal-stopping problem, the true fundamentals $(\mu_1^{\bullet}, \mu_2^{\bullet})$, and the bias parameter $\gamma > 0$, but is independent of the details of the agents' prior density over fundamentals. I prove existence and uniqueness of the steady state and show it features over-pessimism about fundamentals and early stopping. Then, I turn to the stochastic process of beliefs and behavior in the social-learning environment, showing that this process almost surely converges to the steady state I defined.

3.1 Steady State: Existence, Uniqueness, and Other Properties

A steady state is a triplet consisting of fundamentals $(\mu_1^\infty, \mu_2^\infty) \in \mathbb{R}$ and a cutoff threshold $c^\infty \in \mathbb{R}$ that endogenously determine each other. The cutoff strategy with acceptance threshold c^∞ maximizes expected payoff under the subjective model $\Psi(\mu_1^\infty, \mu_2^\infty; \gamma)$, while the fundamentals are the pseudo-true fundamentals under data censoring with threshold c^∞. More precisely,

Definition 6. A *steady state* consists of $\mu_1^\infty, \mu_2^\infty, c^\infty \in \mathbb{R}$ such that:

1. $c^\infty = C(\mu_1^\infty, \mu_2^\infty; \gamma)$.
2. $\mu_1^\infty = \mu_1^*(c^\infty)$ and $\mu_2^\infty = \mu_2^*(c^\infty)$.

Steady states correspond to Esponda and Pouzo (2016)'s Berk-Nash equilibria for an agent whose prior is supported on the feasible models with feasible fundamentals $\mathcal{M} = \mathbb{R}^2$, under the restriction that equilibrium belief puts full confidence in a single fundamental pair. The set of steady states depends on γ, since the severity of the bias changes both the optimal cutoff thresholds under different fundamentals and inference about fundamentals from stage-game histories.

The terminology "steady state" will soon be justified, as I will show the "steady state" defined here almost surely characterizes the long-run learning outcome in the society where biased agents act one by one. This convergence does not follow from Esponda and Pouzo (2016), for their results only imply local convergence from prior beliefs sufficiently close to the equilibrium beliefs, and only in a "perturbed game" environment where learners receive idiosyncratic payoff shocks to different actions. I will show global convergence of the stochastic processes of beliefs and behavior without payoff shocks.

Like almost all examples of Berk-Nash equilibrium in Esponda and Pouzo (2016), my steady state generates data with positive KL divergence relative to the implied data distribution under the steady-state beliefs. That is, $\mathcal{H}^\bullet(c^\infty) \neq \mathcal{H}(\Psi(\mu_1^\infty, \mu_2^\infty; \gamma); c^\infty)$, so the steady state is not a self-confirming equilibrium.[13] This is because for every censoring threshold c

[13] For example, under the history distribution $\mathcal{H}^\bullet(c^\infty)$,

$$\mathbb{E}[h_2 | c^\infty - 1 \leq h_1 \leq c^\infty] = \mathbb{E}[h_2 | c^\infty - 2 \leq h_1 \leq c^\infty - 1]$$

since draws are objectively independent. However, under the history distribution driven by the steady-state feasible model $\Psi(\mu_1^\infty, \mu_2^\infty; \gamma)$, we must have

$$\mathbb{E}[h_2 | c^\infty - 1 \leq h_1 \leq c^\infty] < \mathbb{E}[h_2 | c^\infty - 2 \leq h_1 \leq c^\infty - 1]$$

(and in particular for $c = c^\infty$), the KL divergences of the true history distribution to the implied history distributions under different feasible models is bounded away from 0.

To prove the existence and uniqueness of steady state, I define the following belief iteration map on the second-period fundamental.

Definition 7. For $\gamma > 0$, the **iteration map** $\mathcal{I} : \mathbb{R} \to \mathbb{R}$ is given by

$$\mathcal{I}(\mu_2; \gamma) := \mu_2^\bullet - \gamma \left(\mu_1^\bullet - \mathbb{E}\left[X_1 \mid X_1 \le C(\mu_1^\bullet, \mu_2; \gamma)\right]\right).$$

Given the explicit expression of the pseudo-true fundamentals in Proposition 2, it is not difficult to see that all steady states must have correct belief about μ_1, and that steady-state beliefs about μ_2 are in bijection with fixed points of $\mathcal{I}$.

Proposition 3. *Under Assumptions 1 and 2, $\mathcal{I}$ is a contraction mapping with contraction constant $0 < \ell\gamma < 1$. Therefore, a unique steady state exists.*

As hinted at in Section 2.1.2 the contraction mapping property of $\mathcal{I}$ comes from the Lipschitz continuity of the indifference threshold implied by Assumption 2.

Lemma 1. *Under Assumptions 1 and 2, $\mu_2 \mapsto C(\mu_1^\bullet, \mu_2; \gamma)$ is Lipschitz continuous with Lipschitz constant ℓ.*

Even under Assumption 1 alone, the basic regularity conditions we maintain throughout, it turns out $\mathcal{I}$ is "almost" a contraction mapping for any $\gamma > 0$, in the sense that $|\mathcal{I}(\mu_2^{'}) - \mathcal{I}(\mu_2^{''})| < |\mu_2^{'} - \mu_2^{''}|$ for every $\mu_2^{'}, \mu_2^{''} \in \mathbb{R}$. But, there is no guarantee of a uniform contraction constant strictly less than 1. The slight strengthening in Assumption 2 ensures such a uniform contraction constant exists, providing the crucial step needed for existence and uniqueness of a steady state.

Since $\mu_2^*(c) < \mu_2^\bullet$ for all $c \in \mathbb{R}$ by Proposition 2, this shows steady state belief about μ_2 is exhibits over-pessimism. From the same Proposition, $\mu_1^\infty = \mu_1^\bullet$.

Proposition 4. *Every steady state satisfies $\mu_2^\infty < \mu_2^\bullet$, $\mu_1^\infty = \mu_1^\bullet$.*

I now show the steady-state stopping threshold always features stopping too early. For every $\mu_1^\bullet, \mu_2^\bullet \in \mathbb{R}$, the objectively optimal stopping strategy takes the form of a cutoff $c^\bullet \in \mathbb{R} \cup \{\pm\infty\}$, where $c^\bullet = -\infty$ means always stopping and $c^\bullet = \infty$ means never stopping.[14] I show that $c^\bullet > c^\infty$ for every steady-state cutoff c^∞. (This result only requires Assumption 2 and does not require uniqueness of steady states.)

since $\gamma > 0$.

[14]This follows from Lemma A.2 in the Appendix, which shows even when $\gamma = 0$, the difference between stopping payoff at x_1 and expected continuation payoff after x_1 is strictly increasing and continuous in x_1.

This result does not directly follow from over-pessimism. In fact, short of the steady state, there is an intuition that a biased agent may stop later than a rational agent, not earlier. For a concrete illustration, consider Example 1 with $q = 0$, so there is no probability of recall. Suppose the true fundamentals are $\mu_1^\bullet \gg \mu_2^\bullet$, meaning the late applicant pool is much worse than the early applicant pool. If a biased agent has the correct beliefs about the fundamentals, she perceives a greater continuation value after $X_1 = \mu_2^\bullet$ than a rational agent with the same correct beliefs, since the former holds a false expectation of positive reversals after a bad early draw. Even though $c^\bullet = \mu_2^\bullet$ and the rational agent chooses to stop, the biased agent chooses to continue and has an indifference threshold strictly above $c^\bullet$. By continuity, the biased agent's cutoff threshold remains strictly above $c^\bullet$ even under slightly pessimistic beliefs about μ_2.

Nevertheless, the next Proposition shows that in the steady state, it is unambiguous that the biased agent stops too early relative to the objectively optimal threshold.

Proposition 5. *Under Assumption 2, every steady-state stopping threshold c^∞ is strictly lower than the objectively optimal threshold, $c^\bullet$.*

The early-stopping result strengthens the over-pessimism result. In the steady state, agents must be sufficiently pessimistic as to overcome the opposite intuition about late stopping that I discussed before.

3.2 Social Learning with Agents Acting One by One

This section shows the "steady state" defined and studied earlier warrants its name — it corresponds to the long-run learning outcome for a society of biased agents acting one at a time. I outline the convergence proof for a simpler variant of Theorem 1, where agents start off knowing $\mu_1^\bullet$ and only entertain uncertainty over μ_2. That is, the feasible fundamentals are given by Remark 1(c) rather than Remark 1(b). This simplification is without much loss: even when agents are initially uncertainty about μ_1, they will almost surely learn it in the long run regardless of the stochastic process of their predecessors' stopping strategies. Intuitively, this is because X_1 can never be censored, so no belief distortion in μ_1 is possible.[15] Once agents have learned $\mu_1^\bullet$, the rest of the argument proceeds much like the case where $\mu_1^\bullet$ is known from the start. In the next section I comment on the key steps in extending the proof to the case uncertainty over two-dimensional fundamentals (μ_1, μ_2), but will defer the details to Online Appendix OA 2.

In the learning environment, time is discrete and partitioned into rounds[16] $t = 1, 2, 3, ...$ One short-lived agent arrives per round. Agent t observes the stage-game histories of all

[15] This is similar to the intuition for why $\mu_1^*(c) = \mu_1^\bullet$ for every c.

[16] I use the term "rounds" to refer to different iterations of the stage game, reserving the term "periods" for the dynamic aspect within the stage game.

predecessors[17] and forms a posterior belief $\tilde{G}_t$ about the fundamentals using Bayes' rule. Next, agent t chooses a cutoff threshold $\tilde{C}_t$ maximizing expected payoff based on expected utility, plays the stage game, and exits. Her stage-game history, $\tilde{H}_t \in \mathbb{H}$, then becomes part of the dataset for all future agents.

The sequences $(\tilde{G}_t), (\tilde{C}_t), (\tilde{H}_t)$ are stochastic processes whose randomness stem from randomness of the stage-game draws realizations in different rounds. The convergence theorem is about the almost sure convergence of processes $(\tilde{G}_t)$ and $(\tilde{C}_t)$. To define the probability space formally, consider the $\mathbb{R}^2$-valued stochastic process $(X_t)_{t\geq 1} = (X_{1,t}, X_{2,t})_{t\geq 1}$, where X_t and $X_{t'}$ are independent for $t \neq t^{'}$. Within each t, $X_{1,t} \sim \mathcal{N}(\mu_1^\bullet, \sigma^2)$, $X_{2,t} \sim \mathcal{N}(\mu_2^\bullet, \sigma^2)$ are also independent. Interpret X_t as the pair of potential draws in the t-th round of the stage game. Clearly, there exists a probability space $(\Omega, \mathcal{A}, \mathbb{P})$, with sample space $\Omega = (\mathbb{R}^2)^\infty$ interpreted as paths of the process just described, $\mathcal{A}$ the Borel σ-algebra on Ω, and $\mathbb{P}$ the measure on sample paths so that the process $X_t(\omega) = \omega_t$ has the desired distribution. The term "almost surely" means "with probability 1 with respect to the realization of the infinite sequence of all (potential) draws", i.e. $\mathbb{P}$-almost surely. The processes $(\tilde{G}_t), (\tilde{C}_t), (\tilde{H}_t)$ are defined on this probability space and adapted to the filtration $(\mathcal{F}_t)_{t\geq 1}$, where $\mathcal{F}_t$ is the sub-σ-algebra generated by draws up to round t, $\mathcal{F}_t = \sigma((X_s)_{s=1}^t)$.

Under Assumptions 1 and 2, by Proposition 3 there exists a unique steady state $(\mu_1^\bullet, \mu_2^\infty, c^\infty)$. Following the specification of feasible models in Remark 1(c), let feasible fundamentals be $\mathcal{M} = \{\mu_1^\bullet\} \times [\underline{\mu}_2, \times \bar{\mu}_2]$ and suppose agents' prior belief over fundamentals is given by a common prior density $g : [\underline{\mu}_2, \bar{\mu}_2] \to \mathbb{R}_{>0}$. Theorem 1′ shows that, provided the support of g contains μ_2^∞ and $g^{'}$ is continuous, the stochastic processes $(\tilde{C}_t)$ and $(\tilde{G}_t)$ almost surely converge to the steady state. This is a global convergence result since the bounded interval $[\underline{\mu}_2, \bar{\mu}_2]$ can be arbitrarily large and the prior density g can assign arbitrarily small probability to neighborhoods around μ_2^∞.

Theorem 1′. *Suppose Assumptions 1 and 2 hold, $\underline{\mu}_2 \leq \mu_2^\infty \leq \bar{\mu}_2$ where μ_2^∞ is the unique steady-state belief, and agents have prior density $g : [\underline{\mu}_2, \bar{\mu}_2] \to \mathbb{R}_{>0}$ with $g^{'}$ continuous. Almost surely, $\lim_{t\to\infty} \tilde{C}_t = c^\infty$ and $\lim_{t\to\infty} \mathbb{E}_{\mu_2 \sim \tilde{G}_t} |\mu_2 - \mu_2^\infty| = 0$, where c^∞ is the unique steady-state cutoff threshold.*

I will now discuss the obstacles to proving convergence and provide the outline of my argument. In each round t, the cutoff choice of the t-th agent determines how history $\tilde{H}_t$ will be censored. We can think of each $c \in \mathbb{R}$ as generating a different "type" of data. As we saw in Proposition 2, different "types" of data (in large samples) lead to different inferences about the fundamentals for biased agents. Yet this cutoff $\tilde{C}_t$ is an endogenous, ex-ante random object that depends on the belief of the t-th agent, which complicates the analysis of learning dynamics.

[17] Results are unchanged if agent t does not know the order in which her predecessors moved.

To be more precise, the log-likelihood of all X_2 data up to the end of round t under fundamental $\mu_2 \in [\underline{\mu}_2, \bar{\mu}_2]$ is the random variable

$$\sum_{s=1}^{t} \ln(\phi(X_{2,s}; \mu_2 - \gamma(X_{1,s} - \mu_1^\bullet), \sigma^2) \cdot \mathbf{1}\{X_{1,s} \leq \tilde{C}_s\}.$$

The s-th summand contains the indicator $\mathbf{1}\{X_{1,s} \leq C_s\}$, referring to the fact that $X_{2,s}$ would be censored if $X_{1,s}$ exceeds the cutoff $\tilde{C}_s$. The cutoff $\tilde{C}_s$ depends on histories in periods $1, 2, ..., s-1$, hence indirectly on $(X_k)_{k<s}$. This makes the summands non-exchangeable: they are correlated and non-identically distributed. So the usual law of large numbers does not apply.

A first step to gaining traction on this problem is use a statistical tool from Heidhues, Koszegi, and Strack (2018), a version of law of large numbers for martingales whose quadratic variation grows linearly.

Proposition 10 from Heidhues, Koszegi, and Strack (2018): *Let $(y_t)_t$ be a martingale that satisfies a.s. $[y_t] \leq vt$ for some constant $v \geq 0$. We have that a.s. $\lim_{t\to\infty} \frac{y_t}{t} = 0$.*

After simplifying the problem with this result, I can establish a pair of mutual bounds on asymptotic behavior and asymptotic beliefs. If we know cutoff thresholds are asymptotically bounded between c^l and c^h, $c^l < c^h$, then beliefs about μ_2 must be asymptotically supported on the interval $[\mu_2^*(c^l), \mu_2^*(c^h)]$. Conversely, if belief is asymptotically supported on the subinterval $[\mu_2^l, \mu_2^h] \subseteq [\underline{\mu}_2, \bar{\mu}_2]$, then cutoff thresholds must be asymptotically bounded between $C(\mu_1^\bullet, \mu_2^l; \gamma)$ and $C(\mu_1^\bullet, \mu_2^h; \gamma)$.

Lemma A.11. *For $c^l \geq C(\mu_1^\bullet, \underline{\mu}_2; \gamma)$, if almost surely $\liminf_{t\to\infty} \tilde{C}_t \geq c^l$, then almost surely*

$$\lim_{t\to\infty} \tilde{G}_t(\ [\underline{\mu}_2, \mu_2^*(c^l))\ ) = 0.$$

Also, for $c^h \leq C(\mu_1^\bullet, \bar{\mu}_2; \gamma)$, if almost surely $\limsup_{t\to\infty} \tilde{C}_t \leq c^h$, then almost surely

$$\lim_{t\to\infty} \tilde{G}_t(\ (\mu_2^*(c^h), \bar{\mu}_2]) = 0.$$

Lemma A.12. *For $\underline{\mu}_2 \leq \mu_2^l < \mu_2^h \leq \bar{\mu}_2$, if $\lim_{t\to\infty} \tilde{G}_t([\mu_2^l, \mu_2^h]) = 1$ almost surely, then $\liminf_{t\to\infty} \tilde{C}_t \geq C(\mu_1^\bullet, \mu_2^l; \gamma)$ and $\limsup_{t\to\infty} \tilde{C}_t \leq C(\mu_1^\bullet, \mu_2^h; \gamma)$ almost surely.*

Applying this pair of Lemmas to $\text{supp}(g) = [\underline{\mu}_2, \bar{\mu}_2]$, we conclude that asymptotically $\tilde{G}_t$ must be supported on the subinterval $[\mathcal{I}(\underline{\mu}_2), \mathcal{I}(\bar{\mu}_2)]$, where $\mathcal{I}$ is the iteration map from Definition 7 first used in analyzing the existence and uniqueness of steady states. Under Assumptions 1 and 2, Proposition 3 implies that $\mathcal{I}$ is a contraction mapping whose iterates converge to μ_2^∞. Therefore by repeatedly applying the pair of Lemmas A.11 and A.12, we can refine the bound on asymptotic beliefs down to the singleton $\{\mu_2^\infty\}$, since $\mu_2^\infty \in [\underline{\mu}_2, \bar{\mu}_2]$. This shows the almost-sure convergence of beliefs to μ_2^∞. The almost-sure convergence of behavior follows easily from Lemma A.12.

3.3 Uncertainty About μ_1

The hypotheses of Theorem 1 differ from those of Theorem 1′ just discussed in that agents start off with uncertainty about μ_1. I now comment on the key step to proving almost-sure convergence of beliefs and behavior in the environment with two-dimensional uncertainty about fundamentals.

The structure of the inference problem in my setting is such that I can separately bound the agents' asymptotic beliefs in two "directions," thus reducing the task of proving a two-dimensional belief bound into a pair of tasks involving one-dimensional belief bounds. To understand why, consider a pair of fundamentals, (μ_1, μ_2) and $(\mu_1^{'}, \mu_2^{'}) = (\mu_1 + d, \mu_2 - \gamma d)$ for some $d > 0$, satisfying $\mu_1, \mu_1^{'} \leq \mu_1^\bullet$. That is, (μ_1, μ_2) and $(\mu_1^{'}, \mu_2^{'})$ lie on the same line with slope $-\gamma$. For any uncensored history $(x_1, x_2) \in \mathbb{R}^2$, the likelihood of second-period draw x_2 is the same under both pairs of fundamentals,

$$\phi(x_2; \mu_2 - \gamma(x_1 - \mu_1); \sigma^2) = \phi(x_2; \mu_2^{'} - \gamma(x_1 - \mu_1^{'}); \sigma^2).$$

This is because the subjective model $\Psi(\mu_1, \mu_2; \gamma)$ has a lower first-period mean but also a higher second-period unconditional mean, compared to the subjective model $\Psi(\mu_1^{'}, \mu_2^{'}; \gamma)$. An agent who believes in the first model feels less disappointed by the draw x_1, since she evaluates it against a lower expectation. This leads a weaker anticipation of positive reversal under the gambler's fallacy, compared to another agent who believes in the second model. But, this difference canceled out by the more optimistic belief about the unconditional distribution of second-period draw, $\mu_2 > \mu_2^{'}$.

This argument shows that both pairs of fundamentals (μ_1, μ_2) and $(\mu_1^{'}, \mu_2^{'})$ explain X_2 data equally well in *all* uncensored histories. This is important as it shows (μ_1, μ_2) and $(\mu_1^{'}, \mu_2^{'})$ always lead to the same likelihood of second-period data regardless of how predecessors have censored their histories. At the same time, $(\mu_1^{'}, \mu_2^{'})$ provides a strictly better fit for X_1 data on average than (μ_1, μ_2), since $|\mu_1^{'} - \mu_1^\bullet| < |\mu_1 - \mu_1^\bullet|$. This means in the long run, fundamentals (μ_1, μ_2) should receive much less posterior probability than $(\mu_1^{'}, \mu_2^{'})$, as the latter better rationalize the data overall.

This heuristic comparison of the asymptotic goodness-of-fit for two feasible models is formalized by computing the *directional* derivative for data log-likelihood along the vector $\begin{pmatrix} 1 \\ -\gamma \end{pmatrix}$ in the space of fundamentals. I establish an (almost-sure) positive lowerbound on this directional derivative to the left of $\mu_1^\bullet$, and an analogous negative upperbound to the right of $\mu_1^\bullet$. This allows me to show the region colored in red receives 0 posterior probability asymptotically, by comparing each point in red with a corresponding point in blue along a line of slope $-\gamma$. By repeating this argument (and applying the symmetric bound to the right of $\mu_1^\bullet$), I show that belief is asymptotically concentrated along an ϵ-width vertical strip containing the steady state beliefs, $(\mu_1^\bullet, \mu_2^\infty)$.

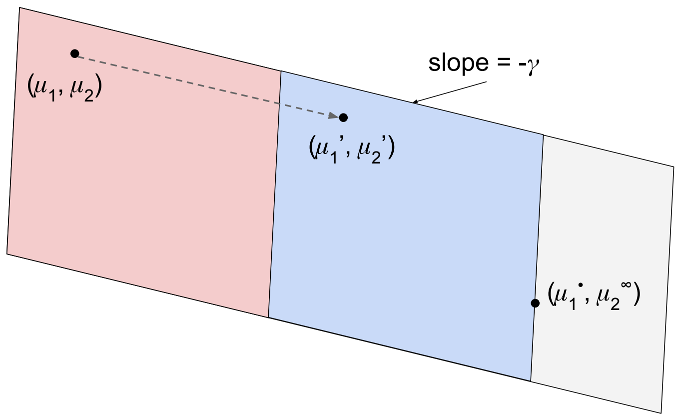

Having restricted the long-run belief to a small vertical strip, we have completed one "direction" of the belief bounds and effectively reduced the dimensionality of uncertainty back to one. The rest of the argument proceeds similarly to the case where agents know $\mu_1^\bullet$ discussed before, iteratively restricting agents' asymptotic behavior and asymptotic belief about μ_2. These restrictions amount to "vertical" belief refinements within the ϵ-strip, so eventually belief is restricted to the single point $(\mu_1^\bullet, \mu_2^\infty)$, the unique steady-state beliefs.

4 The Positive-Feedback Loop

In this section, I turn to my second social-learning environment where agents arrive in large generations and all agents in the same generation act simultaneously. I will prove Theorem 2, fully characterizing the learning dynamics in this environment. I will also discuss the positive-feedback loop between distorted beliefs about fundamentals and distorted stopping behavior.

4.1 Social Learning in Large Generations

There is an infinite sequence of generations, $t \in \{0, 1, 2, ...\}$. Each generation is "large" and will be modeled as a continuum of short-lived agents, $n \in [0, 1]$. In the search problem of Example 1, for instance, different generations refer to cohorts of HR managers working in different hiring cycles. Each agent lives for one generation, so agent n from generation 1 is unrelated to agent n from generation 2. The realizations of draws X_1, X_2 are independent across all stage games, including those from the same generation. Generation 0 agents play some strategy $S_{c_{[0]}}$, where $c_{[0]} \in \mathbb{R}$ is the initial condition of social learning.

Write $h_{\tau,n} \in \mathbb{H}$ for the stage-game history of agent n from generation τ. Before playing her own stage game, each agent in generation $t \geq 1$ observes an infinite dataset consisting of all histories $(h_{\tau,n})_{n\in[0,1]}$ from each predecessor generation, $0 \leq \tau \leq t-1$. If all[18] generation τ predecessors used the stopping strategy S_{c_τ} for some $c_\tau \in \mathbb{R}$, then the sub-dataset $(h_{\tau,n})_{n\in[0,1]}$ has the distribution $\mathcal{H}^\bullet(c_\tau)$. Agents are told the stopping strategies of their predecessors from all past generations[19] and use the entire dataset of histories to infer fundamentals. The space of feasible fundamentals is $\mathcal{M} = \mathbb{R}^2$ as in Remark 1(a), so agents can flexibly estimate the unconditional means of draws from different periods, subject to their dogmatic belief in reversals.

Agents only infer from predecessors' histories, not from their behavior. This is rational as information sets are nested across generations. For $t_2 > t_1$, generation t_2 observes all the social information that generation t_1 saw. In addition, generation t_2's dataset contains a complete record of everything that happened in generation t_1's stage games. Since generation t_1 has no private information that is unobserved by generation t_2, the behavior of these predecessors is uninformative about the fundamentals beyond what generation t_2 can learn from the dataset of histories.

In the large-generations model, generation t agents infer fundamentals $(\mu_{1,[t]}, \mu_{2,[t]})$ that minimize the sum of the KL divergences between the implied history distribution under the feasible model $\Psi(\mu_{1,[t]}, \mu_{2,[t]}; \gamma)$ on the one hand, and the t observed history distributions in generations $0 \leq \tau \leq t-1$ on the other hand. Then, these agents use the stopping strategy optimal for the inferred subjective model. I formally define generation t's minimization objective below.

Definition 8. The *large-generations pseudo-true fundamentals* with respect to cutoff thresholds $(c_\tau)_{\tau=0}^{t-1}$ solve

$$\min_{\mu_1,\mu_2\in\mathbb{R}} \sum_{\tau=0}^{t-1} D_{KL}(\ \mathcal{H}^\bullet(c_\tau)\ ||\ \mathcal{H}(\Psi(\mu_1,\mu_2;\gamma);c_\tau)\ ), \tag{1}$$

where D_{KL} is KL divergence from Definition 5. Denote the minimizers as $\mu_1^*(c_0, ..., c_{t-1})$ and $\mu_2^*(c_0, ..., c_{t-1})$.

I interpret the continuum of agents in each generation as an idealized, tractable modeling device representing a large but finite number of agents. Appendix B provides a finite-population foundation for inference and behavior in the continuum-population model. There,

[18] All generation τ predecessors had the same information about the fundamentals, so all of them would have found the same stopping strategy subjectively optimal.

[19] These stopping rules can also be exactly inferred from the infinite dataset.

I show that when an agent observe t finite sub-datasets of histories drawn from distributions $\mathcal{H}^\bullet(c_\tau)$ for $0 \le \tau \le t-1$, as these datasets grow large her inference and behavior almost surely converge to the infinite-population analogs I study, under some regularity assumptions.

The next lemma relates the large-generations pseudo-true parameters to the pseudo-true parameters of Definition 5.

Lemma 2. *For any $c_0, ..., c_{t-1} \in \mathbb{R}$, the large-generations pseudo-true parameters are*

$$\mu_1^*(c_0, ..., c_{t-1}) = \mu_1^\bullet$$

and

$$\mu_2^*(c_0, ..., c_{t-1}) = \frac{1}{\sum_{\tau=0}^{t-1} \mathbb{P}[X_1 \le c_\tau]} \sum_{\tau=0}^{t-1} \mathbb{P}[X_1 \le c_\tau] \cdot \mu_2^*(c_\tau),$$

where $\mu_2^*(c_\tau)$ is the pseudo-true fundamental associated with the cutoff c_τ.

Intuitively speaking, generation t's inference has to accommodate t sub-datasets of histories, which are censored using t potentially different cutoffs. While a biased agent would draw the same inference about μ_1 using any of these sub-datasets, different sub-datasets lead to different beliefs about μ_2. Her overall inference about μ_2 is a weighted average between these t different beliefs that the different sub-datasets would separately induce. The weight of sub-dataset τ is proportional to $\mathbb{P}[X_1 \le c_\tau]$, the fraction of the observations with an uncensored X_2. That is, in estimating μ_2, the agent puts more weight on those sub-datasets where second-period draws are observed more frequently.

4.2 Learning Dynamics in Large Generations

Now I develop the proof of Theorem 2.

Theorem 2. *Suppose Assumption 1 holds. Starting from any initial condition and any g, cutoffs $(c_{[t]})_{t\ge1}$ and beliefs $(\mu_{[t]})_{t\ge1}$ form monotonic sequences across generations. When Assumption 2 also holds, there exists a unique steady state $\mu_2^\infty, c^\infty \in \mathbb{R}$ so that $c_{[t]} \to c^\infty$ and $(\mu_{1,[t]}, \mu_{2,[t]}) \to (\mu_1^\bullet, \mu_2^\infty)$ monotonically, regardless of the initial condition and g. These steady states are the same as those in Theorem 1.*

Towards a proof, I first consider learning dynamics in an *auxiliary environment*. The auxiliary environment is identical to the large-generations social-learning environment just described, except that agents in each generation $t \ge 1$ only infer from the histories of the immediate predecessor generation, $t-1$. Write $\mu_{[t]}^A$ and $c_{[t]}^A$ for the inference and cutoff threshold in generation t of the auxiliary environment, where the superscript "A" distinguishes them from the corresponding processes of the baseline large-generations environment.

We have $\mu^A_{1,[t]} = \mu^\bullet_1$ for every $t \geq 1$ while $(\mu^A_{2,[t]})_{t\geq 1}$ are iterates of the $\mathcal{I}$ map from Definition 7. The first claim comes from the fact that $\mu^*_1(c) = \mu^\bullet_1$ for every c, from Proposition 2. To see the second claim, start at $\mu^A_{[t]}, c^A_{[t]}$ in generation t. The next generation observes a dataset of histories with the distribution $\mathcal{H}^\bullet(c^A_{[t]})$ and infers

$$\mu^A_{[t+1]} = (\mu^*_1(c^A_{[t]}), \mu^*_2(c^A_{[t]})) = (\mu^\bullet_1, \mu^\bullet_2 - \gamma(\mu^\bullet_1 - \mathbb{E}[X_1 \mid X_1 \leq c^A_{[t]}]),$$

again using Proposition 2. From the optimality of $c^A_{[t]}$, we have $c^A_{[t]} = C(\mu^\bullet_1, \mu^A_{2,[t]}; \gamma)$. So altogether, $\mu^A_{2,[t+1]} = \mathcal{I}(\mu^A_{2,[t]}; \gamma)$.

The pair of comparative statics $\frac{\partial C}{\partial \mu_2} > 0$ and $\frac{d\mu^*_2}{dc} > 0$ have the same sign and lead to the positive-feedback loop in the auxiliary environment. Changes in beliefs across successive generations are amplified, not dampened, by the corresponding changes in cutoff thresholds.

Proposition 6. *Suppose Assumption 1 holds. In the auxiliary environment, starting from any initial condition and any g, cutoffs $(c^A_{[t]})_{t\geq 1}$ and beliefs $(\mu^A_{[t]})_{t\geq 1}$ form monotonic sequences across generations. When Assumption 2 also holds, there exists a unique steady state $\mu^\infty_2, c^\infty \in \mathbb{R}$ so that $c^A_{[t]} \to c^\infty$ and $(\mu^A_{1,[t]}, \mu^A_{2,[t]}) \to (\mu^\bullet_1, \mu^\infty_2)$ monotonically, regardless of the initial condition and g. These steady states are the same as those in Theorem 1.*

The monotonicity and convergence results of Theorem 2 follow from comparing the learning dynamics of the baseline large-generations environment to the dynamics of the auxiliary environment. Suppose two societies in these two different environments start at the same initial condition, $c_{[0]} \in \mathbb{R}$. First, we have $\mu_{2,[1]} = \mu^A_{2,[1]}$. This is because generation 1 only has one generation of predecessors, so observing the histories of all past generations is equivalent to observing the histories of the immediate predecessor generation. Suppose that $\mu^A_{2,[1]} > \mu^\infty_2$, so that by Proposition 6 we have $\mu^A_{2,[t]} \searrow \mu^\infty_2$. Consider the three inferences $\mu^A_{2,[2]}, \mu_{2,[2]}$, and $\mu_{2,[1]}$. The first is based on the history distribution $\mathcal{H}^\bullet(c_{[1]})$, the second based on $\mathcal{H}^\bullet(c_{[1]})$ and $\mathcal{H}^\bullet(c_{[0]})$, and the third based on just $\mathcal{H}^\bullet(c_{[0]})$. We have $c_{[1]} < c_{[0]}$, for otherwise we would have the contradiction $\mu^A_{2,[2]} = \mu^*_2(c_{[1]}) \geq \mu^*_2(c_{[0]}) = \mu^A_{2,[1]}$. From Lemma 2, we may rank the three inferences $\mu^A_{2,[2]} < \mu_{2,[2]} < \mu_{2,[1]}$. By continuing this argument, we can show that beliefs in the baseline large-generations environment $(\mu_{2,[t]})_{t\geq 1}$ decrease in t, but they decrease "more slowly" than beliefs in the auxiliary environment, $(\mu^A_{2,[t]})_{t\geq 1}$. Since $(\mu^A_{2,[t]})_{t\geq 1}$ converges by Proposition 6, $(\mu_{2,[t]})_{t\geq 1}$ must also converge. To see why these two sequences have the same limit, note that as beliefs converge across generations, the stopping thresholds converge as well. For agents in late enough generations of the baseline large-generations environment, most of the histories in their dataset are censored according to stopping thresholds very similar to the limit threshold. So, observing sub-datasets of histories from all past generations induces very similar inferences as observing just one dataset of histories from the immediate predecessor generation.

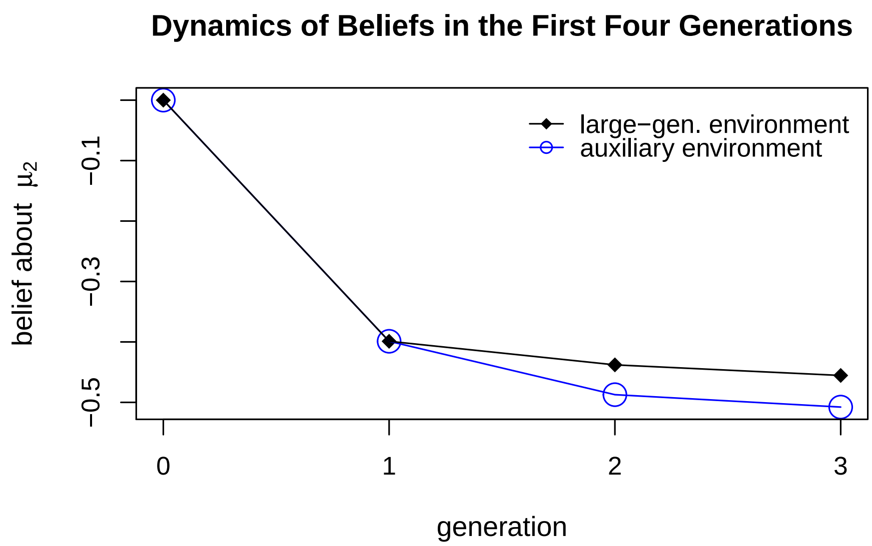

Figure 1: The dynamics of beliefs about μ_2 in the first four generations. The stage game is search (without recall), with true fundamentals are $\mu_1^\bullet = \mu_2^\bullet = 0$, bias parameter $\gamma > 0$, and initial condition is $c_{[0]} = 0$. In both the baseline large-generations environment and the auxiliary environment, beliefs are monotonic across generations, an illustration of Theorem 2 and Proposition 6. Beliefs in both environments converge to the same steady-state beliefs, though the rate of convergence is faster in the auxiliary environment.

While the large-generations environment I set out to study has the same long-run learning outcome as the auxiliary environment, the two environments may differ in their short-run welfare. For example, in settings where learning leads generations further and further astray from the objectively optimal strategy, the auxiliary environment speeds up this harmful learning. This is because the less-censored histories from the earlier generations no longer moderate the society's descent into pessimism when agents only infer from the immediate predecessor generation. In Figure 1, I plot the dynamics of beliefs in the first four generations for a society playing the stage game from Example 1 with $q = 0$. Suppose $\mu_1^\bullet = \mu_2^\bullet = 0$, $\gamma = -0.5$, and the society starts at the objectively optimal cutoff threshold, $c_{[0]} = 0$. Society mislearns monotonically in both the baseline large-generations environment and the auxiliary environment. This mislearning is more exaggerated in the in the auxiliary environment, but Proposition 6 and Theorem 2 imply that both environments lead to the same long-run outcome.

The map $\mathcal{I}(\cdot;\gamma)$ creates an interesting connection between the environment where large generations of agents act simultaneously and the environment where agents act one by one. We can think of $\mathcal{I}(\cdot;\gamma)$ as the one-generation-forward belief map in the auxiliary society, whose belief dynamics are closely related to the belief dynamics of the baseline large-generations environment. There are no large generations at all in the environment where agents at one by one, but there $\mathcal{I}$ still plays a critical role in establishing the long-run convergence of beliefs and behavior. Intuitively, in analyzing the long-run learning outcome in the one-by-one environment, a large dataset containing the histories of one predecessor from each of *many past generations* replaces a large dataset of histories from *many agents* all belonging to the same past generation.

The long-run learning outcome is the same in the environment with large generations and the environment where agents act one by one. This equivalence allows me to combine the asymptotic early-stopping result of Theorem 1 with the monotonic learning dynamics of Theorem 2 to deduce:

Corollary 1. *Suppose Assumptions 1 and 2 hold. In the large-generations environment, if society starts at the objectively optimal initial condition* $c_{[0]} = c^\bullet$*, then expected payoff strictly decreases across all successive generations.*

This stark "monotonic" mislearning result relies crucially on the endogenous-data setting. Each generation uses a lower acceptance threshold relative to their predecessors, a change with the side effect of changing censoring threshold of their successors' data. The new "type" of censored data causes the next generation to become more pessimism about the fundamentals than any past generation.

5 Extensions

In this section I explore a number of alternative model specifications to examine the robustness of my main results. The Online Appendix OA 3 contains additional extensions.

5.1 Comparative Statics

In the first extension, I consider how learning dynamics change with changes in parameters of the stage game. In general, when agents learn from exogenous data, their decision problem does not influence learning outcomes. This observation holds independently of whether agents are misspecified. On the other hand, correctly specified agents in my setting always learn correctly in the long run, so the stage game is again irrelevant. With misspecified learners in an endogenous-data setting, however, changes in the stage game have long-lasting effects on agents' beliefs about the fundamentals.

Definition 9. Given a pair of second-period payoff functions u_2', u_2'', say u_2' *payoff dominates* u_2'' (abbreviated $u_2' \succ u_2''$) if for every $x_1 \in \mathbb{R}$, $u_2'(x_1, x_2) \geq u_2''(x_1, x_2)$ for every $x_2 \in \mathbb{R}$, and also $u_2'(x_1, x_2) > u_2''(x_1, x_2)$ for a positive-measure set of x_2.

For instance, in Example 1, increasing q (the probability of recall) creates a new optimal-stopping problem that payoff dominates the old one. More generally, starting from any stage game with payoff functions u_1 and u_2, we can impose an extra waiting cost $\kappa_{\text{wait}} > 0$ for continuing into the second period. This generates a new stage game with payoff functions u_1 and $u_2^{''}$ with $u_2^{''} = u_2^{'} - \kappa_{\text{wait}}$. The modified stage game is payoff dominated by the unmodified one.

When $u_2^{'} \succ u_2^{''}$, a society facing the problem $(u_1, u_2^{'})$ always uses a higher stopping threshold than a society facing the problem $(u_1, u_2^{''})$, given the same beliefs about fundamentals. To state this formally, let C_{u_1,u_2} be the optimal cutoff threshold function for the stage game (u_1, u_2).

Lemma 3. *Suppose stage games $(u_1, u_2^{'})$ and $(u_1, u_2^{''})$ both satisfy Assumption 1, and $u_2^{'} \succ u_2^{''}$. For all $\mu_1, \mu_2 \in \mathbb{R}$, $\gamma > 0$, $C_{u_1,u_2^{'}}(\mu_1, \mu_2; \gamma) > C_{u_1,u_2^{''}}(\mu_1, \mu_2; \gamma)$.*

The next Proposition shows that when one stage game payoff dominated another in terms of second-period payoffs, the dominated stage game has more pessimistic beliefs and lower cutoff threshold in the steady state.

Proposition 7. *Suppose both $(u_1, u_2^{'})$ and $(u_1, u_2^{''})$ satisfy Assumptions 1 and 2, and that $u_2^{'} \succ u_2^{''}$. The steady state of $(u_1, u_2^{'})$ features strictly more optimistic belief about the second-period fundamental and a strictly higher cutoff threshold than the steady state of $(u_1, u_2^{''})$.*

Combined with my main results on learning dynamics (Theorems 1 and 2), Proposition 7 illustrates the long-run inference implications of a change in the stage game payoff structure. Consider two societies of gambler's fallacy agents with the same bias parameter $\gamma > 0$, facing stage games $(u_1, u_2^{'})$ and $(u_1, u_2^{''})$ respectively, where $u_2^{'}$ payoff dominates $u_2^{''}$. Suppose prior beliefs differ across these two societies and the latter starts with a much more optimistic belief about μ_2. Nevertheless, in the long run the second society ends up with a strictly more pessimistic belief and uses strictly lower cutoff thresholds, both when agents act one at a time and when agents arrive in large generations. Since steady-state beliefs are too pessimistic in both societies, the second society's long-run beliefs are more distorted.

This comparative statics result provides novel predictions about how parameters of the decision problem affect mislearning for gambler's fallacy agents. In the context of hiring from Example 1, this result says when managers are more impatient or when they have a lower chance of recalling previous applicants, then they will end up with more pessimistic beliefs about the labor pool. The direction of the comparative statics is another expression of the positive-feedback cycle between stopping threshold and inference. When managers become more impatient, for instance, they use a lower acceptance threshold as they wish to finish recruiting earlier. The lower cutoff intensifies the censoring effect on histories, leading to more pessimistic inference about the fundamentals. The extra pessimism, in turn, leads future managers to further lower their acceptance threshold, amplifying the initial change in behavior that came from an increase in waiting cost.

5.2 Fictitious Variation and Censored Datasets

So far, I have assumed agents hold dogmatic and correct beliefs about the variance of X_1 and the conditional variance of $X_2|(X_1 = x_1)$. In this extension, I expand the set of feasible models and consider agents who are uncertain about the variances of the draws and jointly estimate means and variances using the histories of their predecessors. I show that agents exaggerate variances in a way that depends on the severity of data censoring, and study how this belief in *fictitious variation* strengthens the positive-feedback cycle between beliefs and behavior.

For $\mu_1, \mu_2 \in \mathbb{R}$, $\sigma_1^2, \sigma_2^2 \geq 0$, and $\gamma \geq 0$, let $\Psi(\mu_1, \mu_2, \sigma_1^2, \sigma_2^2; \gamma)$ refer to the joint distribution

$$X_1 \sim \mathcal{N}(\mu_1, \sigma_1^2)$$
$$(X_2|X_1 = x_1) \sim \mathcal{N}(\mu_2 - \gamma(x_1 - \mu_1), \sigma_2^2).$$

Objectively, X_1, X_2 are independent Gaussian random variables each with a variance of $(\sigma^\bullet)^2 > 0$, so the true joint distribution of (X_1, X_2) is $\Psi^\bullet = \Psi(\mu_1^\bullet, \mu_2^\bullet, (\sigma^\bullet)^2, (\sigma^\bullet)^2; 0)$. Suppose agents have a full-support belief over the class of feasible models

$$\left\{\Psi(\mu_1, \mu_2, \sigma_1^2, \sigma_2^2; \gamma) : \mu_1, \mu_2 \in \mathbb{R}, \sigma_1^2, \sigma_2^2 \geq 0\right\}$$

for a fixed bias parameter $\gamma > 0$. For this extension, "fundamentals" refer to the four parameters $\mu_1, \mu_2, \sigma_1^2, \sigma_2^2$.

Following Definition 5, I write $D_{KL}(\mathcal{H}^\bullet(c) \parallel \mathcal{H}(\Psi(\mu_1, \mu_2, \sigma_1^2, \sigma_2^2; \gamma); c))$ to denote the KL divergence between the true distribution of histories with censoring threshold c and the implied history distribution under the fundamentals $\mu_1, \mu_2, \sigma_1^2, \sigma_2^2$. This divergence is given by

$$\int_c^\infty \phi(x_1; \mu_1^\bullet, (\sigma^\bullet)^2) \cdot \ln\left(\frac{\phi(x_1; \mu_1^\bullet, (\sigma^\bullet)^2)}{\phi(x_1; \mu_1, \sigma_1^2)}\right) dx_1 \tag{2}$$
$$+ \int_{-\infty}^c \left\{\int_{-\infty}^\infty \phi(x_1; \mu_1^\bullet, (\sigma^\bullet)^2) \cdot \phi(x_2; \mu_2^\bullet, (\sigma^\bullet)^2) \cdot \ln\left[\frac{\phi(x_1; \mu_1^\bullet, (\sigma^\bullet)^2) \cdot \phi(x_2; \mu_2^\bullet (\sigma^\bullet)^2)}{\phi(x_1; \mu_1, \sigma_2^2) \cdot \phi(x_2; \mu_2 - \gamma(x_1 - \mu_1), \sigma_2^2)}\right] dx_2\right\} dx_1.$$

The next Proposition characterizes the pseudo-true fundamentals $\mu_1^*, \mu_2^*, (\sigma_1^*)^2, (\sigma_2^*)^2$ that minimize Equation (2) in closed-form expressions.

Proposition 8. *The solutions of*

$$\min_{\mu_1,\mu_2\in\mathbb{R},\sigma_1^2,\sigma_2^2\geq 0} D_{KL}(\mathcal{H}^\bullet(c) \parallel \mathcal{H}(\Psi(\mu_1,\mu_2,\sigma_1^2,\sigma_2^2;\gamma);c))\ )$$

are $\mu_1^* = \mu_1^\bullet$, $\mu_2^* = \mu_2^\bullet - \gamma\left(\mu_1^\bullet - \mathbb{E}\left[X_1 \mid X_1 \leq c\right]\right)$, $(\sigma_1^*)^2 = (\sigma^\bullet)^2$, *and*

$$(\sigma_2^*)^2 = (\sigma^\bullet)^2 + \gamma^2\, Var[X_1 \mid X_1 \leq c].$$

Given a large dataset of histories with censoring threshold c, an agent's inferences about the means remain the same as in the case when she know the variances. This result shows the robustness of the over-pessimism prediction, as the biased inferences about the means are too low even when agents jointly estimate both means and variances.

A biased agent correctly estimate the first-period variance, $(\sigma_1^*)^2 = (\sigma^\bullet)^2$, but her estimate of the second-period variance is too high. The magnitude of this distortion increases in the severity of the gambler's fallacy but decreases with the severity of the censoring, as $\mathrm{Var}[X_1 \mid X_1 \leq c]$ is increasing in c for X_1 is Gaussian.

The intuition for misinferring the second-period conditional variance is the following. Whereas the objective conditional distribution of $X_2|(X_1 = x_1)$ is independent of x_1, the agent entertains a different subjective model for this conditional distribution for each x_1. The agent's best-fitting belief about the second-period fundamental $\mu_2^* < \mu_2^\bullet$ ensures her subjective model about $X_2|X_1 = x_1$ fits the data well following "typical" realizations of x_1 under the restriction $X_1 \leq c$. However, following unusually high X_1 the agent is surprised by high values of X_2, while following unusually low X_1 she is surprised by low values of X_1. To better account for these surprising observations of X_2, the agent increases estimated conditional variance of $X_2|(X_1 = x_1)$ and attributes these surprises to "noise." The extent of variance overestimation increases in $\mathrm{Var}[X_1|X_1 \leq c]$, for the frequency of these surprising observations depends on how much X_1 under the restriction $X_1 \leq c$ tends to deviate from its typical value, $\mathbb{E}[X_1|X_1 \leq c]$. And of course, the extent of overestimation increases in severity of the gambler's fallacy bias, which increases the size of these surprises.

An equivalent formulation of this result helps interpret the distorted $(\sigma_2^*)^2$. We may write the subjective model $\Psi(\mu_1,\mu_2,\sigma_1^2,\sigma_2^2;\gamma)$ with $\sigma_2^2 = \sigma_1^2 + \sigma_\eta^2$, $\sigma_\eta^2 \geq 0$ as

$$\begin{aligned} X_1 &= \mu_1 + \epsilon_1 \\ X_2 &= \mu_2 + \zeta + \epsilon_2 \end{aligned}$$

where $\epsilon_1 \sim \mathcal{N}(0,\sigma_1^2)$, $\epsilon_2|\epsilon_1 \sim \mathcal{N}(-\gamma\epsilon_1,\sigma_1^2)$, and $\zeta \sim \mathcal{N}(0,\sigma_\zeta^2)$, with ζ independent of ϵ_1,ϵ_2. In the context where X_1 and X_2 represent the quality realizations of two candidates from the early and late applicant pools, ζ is a vacancy-specific shock to the average quality of the second-period applicant. A positive σ_ζ^2 means there are some vacancies for which the late applicants are an especially poor fit and some others for which they are especially suitable. Proposition 8 says that in an environment where all jobs are objectively homogeneous with respect to the fit of the late applicants, biased managers who find it possible that jobs are heterogeneous in this dimension will indeed estimate a positive amount of this heterogeneity, $\sigma_\zeta^2 > 0$, from the censored histories of their predecessors. This added heterogeneity allows agents to better rationalize histories (X_1, X_2) where both candidates have unusually high/low qualities as vacancies that happen to be an especially good/bad fit for second-period applicants (i.e. the realization of ζ, a vacancy-specific fixed effect, is far from 0.)

This phenomenon relates to findings in Rabin (2002) and Rabin and Vayanos (2010), who refer to exaggeration of variance under the gambler's fallacy as *fictitious variation*. The key innovation of Proposition 8 is to show, in an endogenous-data setting, how the degree of fictitious variation depends on the severity of the censoring. To highlight this point, I now derive two results focusing on the interplay between fictitious variation and the endogenous censoring. For simplicity, I derive these results using the auxiliary large-generations environment defined in Section 4.1, where agents arrive in continuum generations and only infer from the histories of the immediate predecessor generation.

The first result says that when the second-period payoff $u_2(x_1, x_2)$ is a linear or convex function of x_2, the positive-feedback cycle from Section 4 continues to obtain — cutoffs, beliefs about fundamentals, and beliefs about variances form monotonic sequences across generations. This includes the case of search with recall in Example 1 for any recall probability $0 \le q < 1$. Also, when u_2 is a convex function of x_2, inference about variance provides a new channel for amplifying changes in behavior across generations.

Definition 10. The optimal-stopping problem is *convex* if for every $x_1 \in \mathbb{R}$, $x_2 \mapsto u_2(x_1, x_2)$ is convex with strict convexity for x_2 in a positive-measure set. The optimal-stopping problem is *concave* if for every $x_1 \in \mathbb{R}$, $x_2 \mapsto u_2(x_1, x_2)$ is concave with strict concavity for x_2 in a positive-measure set.

Proposition 9. *Suppose the optimal-stopping problem is convex. Suppose agents start with a full-support prior over* $\{\Psi(\mu_1, \mu_2, \sigma_1^2, \sigma_2^2; \gamma) : \mu_1, \mu_2 \in \mathbb{R}, \sigma_1^2, \sigma_2^2 \ge 0\}$ *and society starts at the initial condition* $c_{[0]} \in \mathbb{R}$. *For* $t \ge 1$, *denote the beliefs of generation* t *as* $(\mu_{1,[t]}, \mu_{2,[t]}, \sigma_{1,[t]}^2, \sigma_{2,[t]}^2)$ *and their cutoff threshold as* $c_{[t]}$. *Then* $\mu_{1,[t]} = \mu_1^\bullet$, $(\sigma_{1,[t]})^2 = (\sigma^\bullet)^2$ *for all* t, *while* $(\mu_{2,[t]})_{t\ge1}$, ${(\sigma_{2,[t]})^2}_{t\ge1}$, *and* $(c_{[t]})_{t\ge1}$ *are monotonic sequences.*

The intuition is straightforward. Suppose generation t uses a more relaxed acceptance threshold $c_{[t]} < c_{[t-1]}$ than generation $t-1$, resulting in a more severely censored dataset. By the usual censoring effect with known variances, generation $t+1$ becomes more pessimistic about second-period mean than generation t. In addition, by Proposition 8 we know that generation $t+1$ suffers less from fictitious variation than generation t. This implies generation $t+1$ agents would perceive less continuation value than generation t agents even if they held the same beliefs about the means, for a larger variance in $X_2|(X_1 = x_1)$ improves the expected payoff when continuing due to the convexity of u_2 in x_2. Combining these two forces, we deduce $c_{[t+1]} < c_{[t]}$.

The intuition just discussed shows that uncertainty about variance strengthens the monotonicity result. To be more precise, suppose $c_{[t]} < c_{[t-1]}$. Consider a hypothetical generation $t+1$ agent who dogmatically adopts generation t's beliefs about variances, $\sigma^2_{1,[t]}$ and $\sigma^2_{2,[t]}$, and infers from the class of models $\{\Psi(\mu_1, \mu_2, \sigma^2_{1,[t]}, \sigma^2_{2,[t]}; \gamma) : \mu_1, \mu_2 \in \mathbb{R}\}$. Based on generation t's histories, this hypothetical agent makes inferences about means and chooses a cutoff threshold, $\hat{\mu}_{1,[t+1]}, \hat{\mu}_{2,[t+1]}, \hat{c}_{[t+1]}$. By comparing Proposition 8 and Proposition 2, $\hat{\mu}_{1,[t+1]} = \mu_{1,[t+1]}$, $\hat{\mu}_{2,[t+1]} = \mu_{2,[t+1]}$, but $c_{[t+1]} < \hat{c}_{[t+1]} < c_{[t]}$. That is, while the cutoff threshold of this hypothetical agent follows the monotonicity pattern in the previous two generations, $\hat{c}_{[t+1]} < c_{[t]} < c_{[t-1]}$, the cutoff adjusts downwards by an even greater amount, $c_{[t+1]} < \hat{c}_{[t+1]}$, when agents are uncertain about variances.

The second result compares the learning dynamics of two societies facing the same optimal-stopping problem and starting at the same initial condition. One society knows the correct variances of X_1 and $X_2|(X_1 = x_1)$. The other society is uncertain about the variances and infers them from data. Proposition 10 shows that in generation 1, the two societies hold the same beliefs about the means of the distributions, $\mu_1^\bullet$ and $\mu_2^\bullet$. But in all later generations $t \geq 2$, the society with uncertainty about variances ends up with a more pessimistic/optimistic belief about the second-period fundamental compared with the society that knows the variances, provided the optimal-stopping problem is convex/concave. This divergence depends crucially on the endogenous-learning setting, for Proposition 8 implies that the two societies make the same inferences about the means when given the same dataset. But, since agents inferring variances end up believing in fictitious variation, they perceive a different continuation value than their peers in the same generation from the society that knows the variances. This causes the variance-inferring agents to use a different cutoff threshold, which affects the dataset that their successors observe. In short, allowing uncertainty on one dimension (variance) ends up affecting society's long-run inference in another dimension (mean).

Formally, consider two societies of agents, A and B. Agents in society A start with a full-support prior over $\{\Psi(\mu_1, \mu_2, (\sigma^\bullet)^2, (\sigma^\bullet)^2; \gamma) : \mu_1, \mu_2 \in \mathbb{R}\}$. Agents in society B start with a full-support prior over $\{\Psi(\mu_1, \mu_2, \sigma_1^2, \sigma_2^2; \gamma) : \mu_1, \mu_2 \in \mathbb{R}, \sigma_1^2, \sigma_2^2 \geq 0\}$. Fix the same generation 0 initial condition $c_{[0]} \in \mathbb{R}$ for both societies. For $t \geq 1$, denote the beliefs of generation t in society $k \in \{A, B\}$ as $(\mu_{1,[k,t]}, \mu_{2,[k,t]}, (\sigma_{1,[k,t]})^2, (\sigma_{2,[k,t]})^2)$ and their cutoff threshold as $c_{[k,t]}$.

Proposition 10. *In the first generation,* $\mu_{1,[A,1]} = \mu_{1,[B,1]}$ *and* $\mu_{2,[A,1]} = \mu_{2,[B,1]}$. *If the* *optimal-stopping problem is convex, then* $\mu_{2,[B,t]} > \mu_{2,[A,t]}$ *and* $c_{[B,t]} > c_{[A,t]}$ *for every* $t \geq 2$. *If the optimal-stopping problem is concave, then* $\mu_{2,[B,t]} < \mu_{2,[A,t]}$ *and* $c_{[B,t]} < c_{[A,t]}$ *for every* $t \geq 2$.

5.3 Objectively Correlated $(X_1 X_2)$ and Uncertainty About γ

So far I have assumed that draws (X_1, X_2) within the stage game are objectively independent, and that agents have a dogmatic $\gamma > 0$, interpreted as the severity of the gambler's fallacy bias. This extension considers a joint relaxation of these two assumptions.

Suppose the true model is $(X_1, X_2) \sim \Psi(\mu_1^\bullet, \mu_2^\bullet; \gamma^\bullet)$, where $\gamma^\bullet \in \mathbb{R}$ is possibly non-zero. Agents jointly estimate $(\mu_1, \mu_2, \gamma) \in \mathbb{R}^3$, with a prior supported on $\mathbb{R} \times \mathbb{R} \times [\underline{\gamma}, \bar{\gamma}]$ where $[\underline{\gamma}, \bar{\gamma}]$ is a finite interval. The next result generalizes Proposition 2. It shows that when $\gamma^\bullet \notin [\underline{\gamma}, \bar{\gamma}]$, the KL-divergence minimizing inference involves γ^* equal to $\hat{\gamma} \in \{\underline{\gamma}, \bar{\gamma}\}$, boundary point of the support of γ that is the closest to $\gamma^\bullet$. Given the estimated pseudo-true parameter $\hat{\gamma}$, the estimates of the first- and second-period fundamentals take similar forms to those in Proposition 2.

Proposition 11. *Suppose* $\gamma^\bullet \notin [\underline{\gamma}, \bar{\gamma}]$. *Let* $\tilde{\gamma} = \bar{\gamma}$ *if* $\gamma^\bullet > \bar{\gamma}$, *otherwise* $\tilde{\gamma} = \underline{\gamma}$ *when* $\gamma^\bullet < \underline{\gamma}$. *The solution of the KL-divergence minimization problem*

$$\min_{\mu_1, \mu_2 \in \mathbb{R}, \gamma \in [\underline{\gamma}, \bar{\gamma}]} D_{KL}(\mathcal{H}(\Psi(\mu_1^\bullet, \mu_2^\bullet; \gamma^\bullet); c) \mid\mid \mathcal{H}(\Psi(\mu_1, \mu_2; \gamma); c))$$

is given by $\mu_1^*(c) = \mu_1^\bullet$, $\mu_2^*(c) = \mu_2^\bullet + (\gamma^\bullet - \tilde{\gamma}) \cdot \left(\mu_1^\bullet - \mathbb{E}_{\Psi(\mu_1^\bullet, \mu_2^\bullet; \gamma^\bullet)}[X_1 | X_1 \leq c]\right)$, $\gamma^*(c) = \tilde{\gamma}$.

Intuitively, we may expect the closest distance (in the KL divergence sense) from the set of subjective models $\{\Psi(\mu_1, \mu_2; \hat{\gamma}) : \mu_1, \mu_2 \in \mathbb{R}\}$ to the objective distribution $\Psi(\mu_1^\bullet, \mu_2^\bullet; \gamma^\bullet)$ to decrease in $|\hat{\gamma} - \gamma^\bullet|$. Proposition 11 confirms this intuition, showing that the pseudo-true model from the set $\{\Psi(\mu_1, \mu_2; \gamma) : \mu_1, \mu_2 \in \mathbb{R}, \gamma \in [\underline{\gamma}, \bar{\gamma}]\}$ lies in the subset $\{\Psi(\mu_1, \mu_2; \gamma) : \mu_1, \mu_2 \in \mathbb{R}, \gamma = \tilde{\gamma}\}$, where $\tilde{\gamma}$ is the closest point (in the Euclidean sense) to $\gamma^\bullet$ in the interval $[\underline{\gamma}, \bar{\gamma}]$.

When $\gamma^\bullet = 0$ and $\bar{\gamma} < 0$, this result shows that over-pessimism in inference is robust to agents learning the correlation of X_1 and X_2, provided the support of their uncertainty about correlation lies to the left of 0 and excludes 0. In this case, it is also easy to see that the learning dynamics in the large-generations auxiliary environment are exactly the same as when agents start with a dogmatic belief in $\gamma = \bar{\gamma}$.

5.4 Inference under the Constraint $\mu_1^\bullet = \mu_2^\bullet$

I now consider the special case where the true fundamentals are time-invariant, $\mu_1^\bullet = \mu_2^\bullet = \mu^\bullet \in \mathbb{R}$. If agents' feasible fundamentals are $\mathcal{M} = \mathbb{R}^2$ as in Remark 1(a), then Proposition 2 continues to apply. But now suppose agents know the fundamentals are time-invariant and only have uncertainty over this common value, so the set of feasible fundamentals is the diagonal $\mathcal{M} = \{(x, x) : x \in \mathbb{R}\}$, as in 1(d). I investigate inference in this setting when agents' prior belief over subjective models is supported on $\{\Psi(\mu, \mu; \gamma) : \mu \in \mathbb{R}\}$.

Let $\mu_\triangle^*(c) \in \mathbb{R}$ stand for the common fundamental that minimizes the KL divergence relative to the history distribution $\mathcal{H}^\bullet(c)$, that is

$$\mu_\triangle^*(c) := \arg\min_{\mu \in \mathbb{R}} D_{KL}(\mathcal{H}^\bullet(c) \parallel \mathcal{H}(\Psi(\mu, \mu; \gamma); c))$$

The next lemma characterizes $\mu_\triangle^*(c)$.

Proposition 12. *$\mu_\triangle^*(c) = \frac{1}{1+\mathbb{P}[X_1 \le c]\cdot(1+\gamma)^2}\mu_1^\circ(c) + \frac{\mathbb{P}[X_1 \le c]\cdot(1+\gamma)^2}{1+\mathbb{P}[X_1 \le c]\cdot(1+\gamma)^2}\mu_2^\circ(c)$, where $\mu_1^\circ(c) = \mu^\bullet$ and $\mu_2^\circ(c) = \mu^\bullet - \frac{\gamma}{1+\gamma}\left(\mu^\bullet - \mathbb{E}[X_1 \mid X_1 \le c]\right)$.*

Agents face two kinds of data about the common fundamental: observations of first-period draws and observations of second-period draws. Subjective models $\Psi(\mu_1^\circ(c), \mu_1^\circ(c); \gamma)$ and $\Psi(\mu_2^\circ(c), \mu_2^\circ(c); \gamma)$ minimize the KL divergence of these two kinds of data, respectively.[20] The overall KL-divergence minimizing estimator is a certain convex combination between these two points. Through the term $\mathbb{P}[X_1 \le c]$, the relative weight given to $\mu_2^\circ(c)$ increases as the cutoff c increases, because the second-period data is observed more often if previous agents have used a more stringent cutoff in the first period.

For any censoring threshold c generating the history distribution, agents underestimates the common fundamental. We have $\mu_2^\circ(c) < \mu^\bullet$ while $\mu_1^\circ(c) = \mu^\bullet$. This shows the robustness of the over-pessimism result from the setting with $\mathcal{M} = \mathbb{R}^2$. However, the extent of over-pessimism about μ_2 is dampened relative to agents who can flexibly estimate different μ_1 and μ_2 for the two periods. Compared with the unconstrained pseudo-true fundamentals from Proposition 2, we have $\mu_2^\circ(c) > \mu_2^*(c)$ since $\frac{\gamma}{1+\gamma} < \gamma$, hence $\mu_\triangle^*(c) > \mu_2^*(c)$. This makes intuitive sense: when unconstrained, agents come to two different beliefs about μ_1 and μ_2, even though they are objectively the same. They hold correct beliefs about μ_1 but pessimistic beliefs about μ_2. When constrained to a common inference across two fundamentals, agents distort their belief about μ_1 downwards and their belief about μ_2 upwards, relative to the unconstrained environment.

[20] Note that $\mu_2^\circ(c)$ differs from the pseudo-true fundamental $\mu_2^*(c)$ from Proposition 2. The estimator $\mu_2^\circ(c)$ minimizes the KL divergence of second-period draws under the constraint that the same fundamental must be inferred for both periods, whereas $\mu_2^*(c)$ minimizes this divergence when first-period fundamental is fixed at its true value, $\mu_1^\bullet$.

6 Concluding Remarks

This paper studies endogenous learning dynamics of misspecified agents. I focus on the gambler's fallacy, a non-self-confirming misspecification where no feasible beliefs of the biased agents can exactly match the data. In natural optimal-stopping problems, agents tend to stop after "good enough" early draws, where the threshold for "good enough" evolves as agents update their beliefs about the underlying distributions. Stopping decisions thus impose an endogenous censoring effect on the data of later agents, who use their predecessors' histories to learn about the distributions. The statistical bias interacts with data censoring, generating over-pessimism about the fundamentals and resulting in stopping too early in the long run. These asymptotic mistakes are driven by a positive-feedback loop between distorted beliefs and distorted behavior.

I have studied a particular behavioral bias (the gambler's fallacy) in a natural environment where censoring happens (histories in optimal-stopping problems), leaving open the interaction of other kinds of behavioral learning with other censoring mechanisms to future work.

References

Amemiya, T. (1985): *Advanced econometrics*, Harvard university press.

Andrews, D. W. (1992): "Generic uniform convergence," *Econometric theory*, 8, 241–257.

Bar-Hillel, M. and W. A. Wagenaar (1991): "The perception of randomness," *Advances in Applied Mathematics*, 12, 428–454.

Barron, G. and S. Leider (2010): "The role of experience in the gambler's fallacy," *Journal of Behavioral Decision Making*, 23, 117–129.

Benjamin, D. J., D. A. Moore, and M. Rabin (2017): "Biased Beliefs About Random Samples: Evidence from Two Integrated Experiments," *Working Paper*.

Berk, R. H. (1966): "Limiting behavior of posterior distributions when the model is incorrect," *Annals of Mathematical Statistics*, 37, 51–58.

Bohren, J. A. (2016): "Informational herding with model misspecification," *Journal of Economic Theory*, 163, 222–247.

Bohren, J. A. and D. Hauser (2018): "Bounded Rationality And Learning: A Framework and A Robustness Result," *Working Paper*.

Bromiley, P. (2003): "Products and Convolutions of Gaussian Probability Density Functions," *Working Paper*.

Bunke, O. and X. Milhaud (1998): "Asymptotic behavior of Bayes estimates under possibly incorrect models," *Annals of Statistics*, 26, 617–644.

Bushong, B. and T. Gagnon-Bartsch (2018): "Learning with Misattribution of Reference Dependence," *Working Paper*.

Camerer, C. F. (1987): "Do biases in probability judgment matter in markets? Experimental evidence," *American Economic Review*, 77, 981–997.

Chen, D. L., T. J. Moskowitz, and K. Shue (2016): "Decision making under the gambler's fallacy: Evidence from asylum judges, loan officers, and baseball umpires," *Quarterly Journal of Economics*, 131, 1181–1242.

Conlon, J. J., L. Pilossoph, M. Wiswall, and B. Zafar (2018): "Labor Market Search With Imperfect Information and Learning," *Working Paper*.

Cox, D. R. (1962): *Renewal Theory*, Methuen.

Croson, R. and J. Sundali (2005): "The gambler's fallacy and the hot hand: Empirical data from casinos," *Journal of Risk and Uncertainty*, 30, 195–209.

Crowder, M. J. (2001): *Classical competing risks*, Chapman and Hall/CRC.

Enke, B. (2017): "What you see is all there is," *Working Paper*.

Esponda, I. and D. Pouzo (2016): "Berk–Nash equilibrium: A framework for modeling agents with misspecified models," *Econometrica*, 84, 1093–1130.

Eyster, E. and M. Rabin (2010): "Naive herding in rich-information settings," *American Economic Journal: Microeconomics*, 2, 221–243.

Frick, M., R. Iijima, and Y. Ishii (2018): "Misinterpreting Others and the Fragility of Social Learning," *Working Paper*.

Fudenberg, D., G. Romanyuk, and P. Strack (2017): "Active learning with a misspecified prior," *Theoretical Economics*, 12, 1155–1189.

Gagnon-Bartsch, T., M. Rabin, and J. Schwartzstein (2018): "Channeled Attention and Stable Errors," *Working Paper*.

Gaurino, A. and P. Jehiel (2013): "Social Learning with Coarse Inference," *American Economic Journal: Microeconomics*, 5, 147–74.

Gumbel, E. J. (1960): "Bivariate exponential distributions," *Journal of the American Statistical Association*, 55, 698–707.

Heckman, J. J. and B. E. Honoré (1989): "The identifiability of the competing risks model," *Biometrika*, 76, 325–330.

Heidhues, P., B. Koszegi, and P. Strack (2018): "Unrealistic expectations and misguided learning," *Econometrica*, 86, 1159–1214.

Jehiel, P. (2018): "Investment strategy and selection bias: An equilibrium perspective on overoptimism," *American Economic Review*, 108, 1582–97.

Narayanan, S. and P. Manchanda (2012): "An empirical analysis of individual level casino gambling behavior," *Quantitative Marketing and Economics*, 10, 27–62.

Nyarko, Y. (1991): "Learning in mis-specified models and the possibility of cycles," *Journal of Economic Theory*, 55, 416–427.

Pinelis, I. (2018): "Exact Bounds on the Inverse Mills Ratio and Its Derivatives," *Complex Analysis and Operator Theory*.

Rabin, M. (2002): "Inference by believers in the law of small numbers," *Quarterly Journal of Economics*, 117, 775–816.

Rabin, M. and D. Vayanos (2010): "The gambler's and hot-hand fallacies: Theory and applications," *Review of Economic Studies*, 77, 730–778.

Suetens, S., C. B. Galbo-Jørgensen, and J.-R. Tyran (2016): "Predicting lotto numbers: a natural experiment on the gambler's fallacy and the hot-hand fallacy," *Journal of the European Economic Association*, 14, 584–607.

Terrell, D. (1994): "A test of the gambler's fallacy: Evidence from pari-mutuel games," *Journal of Risk and Uncertainty*, 8, 309–317.

Tsiatis, A. (1975): "A nonidentifiability aspect of the problem of competing risks," *Proceedings of the National Academy of Sciences*, 72, 20–22.

Appendix

A Omitted Proofs from the Main Text

A.1 Proof of Claim 1

Proof. For Example 1, clearly u_1 and u_2 are strictly increasing functions of x_1 and x_2 respectively. We also have that $|u_2(x_1', \bar{x}_2) - u_2(x_1'', \bar{x}_2)| \leq q(x_1' - x_1'')$ for $x_1' > x_1''$ and any $\bar{x}_2$, while $u_1'(x_1) = 1$. This shows Assumption 1(b) holds. If $x_1 > 0$ and $x_2 < 0$, then $u_2(x_1, x_2) = q \cdot x_1 + (1-q)x_2 < x_1 = u_1(x_1)$, and conversely $x_1 < 0, x_2 > 0$ imply $u_2(x_1, x_2) > u_1(x_1)$. This shows Assumption 1(c) holds. It is clear that u_1, u_2 are continuous. We have $|u_2(\bar{x}_1, x_2)| \leq |\bar{x}_1| + |x_2|$, and $\mathbb{E}(|X_2|)$ exists whenever X_2 is Gaussian. This shows Assumption 1(d) holds. □

A.2 Proofs of Proposition 1 and Lemma 1

I prove three lemmas (A.1, A.3, and A.4) which correspond to the three statements in Proposition 1. Along the way, I will also prove Lemma 1.

A.2.1 The Optimal Strategy Has a Cutoff Form

In the first part, I prove the lemma:

Lemma A.1. Under Assumption 1 and the subjective model $\Psi(\mu_1, \mu_2; \gamma)$ for any $\gamma > 0$, there exists a cutoff $C(\mu_1, \mu_2; \gamma)$, such that: (i) the agent strictly prefers stopping after any $x_1 > C(\mu_1, \mu_2; \gamma)$; (ii) the agent is indifferent between continuing and stopping after $x_1 = C(\mu_1, \mu_2; \gamma)$; (iii) the agent strictly prefers continuing after any $x_1 < C(\mu_1, \mu_2; \gamma)$.

Suppose $X_1 = x_1$. Consider the payoff difference between accepting it and continuing under the subjective model $\Psi(\mu_1, \mu_2; \gamma)$ for $\gamma \geq 0$:

$$D(x_1; \mu_1, \mu_2, \gamma) := u_1(x_1) - \mathbb{E}_\Psi[u_2(x_1, X_2)|X_1 = x_1],$$

where $\mathbb{E}_\Psi$ means expectation with respect to $(X_1, X_2) \sim \Psi(\mu_1, \mu_2; \gamma)$. I abbreviate this as $D(x_1)$ when Ψ is fixed. Lemma A.1 follows from three properties of D.

Lemma A.2. *D is strictly increasing and continuous. If $\gamma > 0$, then there are $x_1' < x_1''$ so that $D(x_1') < 0 < D(x_1'')$.*

Proof. **Step 1**: D is strictly increasing.

Suppose $x_1^{'} > \bar{x}_1$. Then,

$$\mathbb{E}_\Psi[u_2(\bar{x}_1, X_2)|X_1 = \bar{x}_1] = \mathbb{E}_{\tilde{X}_2 \sim \mathcal{N}(\mu_2 - \gamma(\bar{x}_1 - \mu_1), \sigma^2)}[u_2(\bar{x}_1, \tilde{X}_2)],$$

while

$$\begin{aligned}\mathbb{E}_\Psi[u_2(\bar{x}_1, X_2)|X_1 = x_1^{'}] &= \mathbb{E}_{\tilde{\tilde{X}}_2 \sim \mathcal{N}(\mu_2 - \gamma(x_1^{'} - \mu_1), \sigma^2)}[u_2(\bar{x}_1, \tilde{\tilde{X}}_2)] \\ &= \mathbb{E}_{\tilde{X}_2 \sim \mathcal{N}(\mu_2 - \gamma(\bar{x}_1 - \mu_1), \sigma^2)}[u_2(\bar{x}_1, \tilde{X}_2 - \gamma(x_1^{'} - \bar{x}_1))].\end{aligned}$$

Since u_2 is strictly increasing in its second argument by Assumption 1(a), we get

$$\mathbb{E}_{\tilde{X}_2 \sim \mathcal{N}(\mu_2 - \gamma(\bar{x}_1 - \mu_1), \sigma^2)}[u_2(\bar{x}_1, \tilde{X}_2 - \gamma(x_1^{'} - \bar{x}_1))] \leq \mathbb{E}_{\tilde{X}_2 \sim \mathcal{N}(\mu_2 - \gamma(\bar{x}_1 - \mu_1), \sigma^2)}[u_2(\bar{x}_1, \tilde{X}_2)]$$

seeing that $\gamma(x_1^{'} - \bar{x}_1) \geq 0$. Also, at any $x_2 \in \mathbb{R}$, by Assumption 1(b) we know that

$$\begin{aligned} u_1(x_1^{'}) - u_1(\bar{x}_1) > u_2(x_1^{'}, x_2) - u_2(\bar{x}_1, x_2). \\ \Rightarrow u_1(x_1^{'}) - u_2(x_1^{'}, x_2) > u_1(\bar{x}_1) - u_2(\bar{x}_1, x_2). \end{aligned}$$

This then shows

$$\begin{aligned} & u_1(x_1^{'}) - \mathbb{E}_{\tilde{X}_2 \sim \mathcal{N}(\mu_2 - \gamma(\bar{x}_1 - \mu_1), \sigma^2)}[u_2(x_1^{'}, \tilde{X}_2 - \gamma(x_1^{'} - \bar{x}_1))] \\ & > u_1(\bar{x}_1) - \mathbb{E}_{\tilde{X}_2 \sim \mathcal{N}(\mu_2 - \gamma(\bar{x}_1 - \mu_1), \sigma^2)}[u_2(\bar{x}_1, \tilde{X}_2 - \gamma(x_1^{'} - \bar{x}_1))] \\ & \geq u_1(\bar{x}_1) - \mathbb{E}_{\tilde{X}_2 \sim \mathcal{N}(\mu_2 - \gamma(\bar{x}_1 - \mu_1), \sigma^2)}[u_2(\bar{x}_1, \tilde{X}_2)] \end{aligned}$$

that is $D(x_1^{'}) > D(\bar{x}_1)$.

Step 2: D is continuous.

Fixing some $\bar{x}_1 \in \mathbb{R}$, I show D is continuous at $\bar{x}_1$. Since u_1 is continuous, find $\delta > 0$ so that whenever $|x_1 - \bar{x}_1| < 1$, $|u(x_1) - u(\bar{x}_1)| < \delta$. Consider the function $f : \mathbb{R} \to \mathbb{R}_{\geq 0}$ defined by $f(x_2) := |u_2(\bar{x}_1, x_2 + \gamma)| + |u_2(\bar{x}_1, x_2 - \gamma)| + \delta$.

Claim A.1. Whenever $|x_1 - \bar{x}_1| < 1$, $|u_2(x_1, x_2 + \gamma(\bar{x}_1 - x_1))| \leq f(x_2)$ for every $x_2 \in \mathbb{R}$.

Proof. Since u_2 is increasing its second argument by Assumption 1(a), if $u_2(x_1, x_2 + \gamma(\bar{x}_1 - x_1)) \geq 0$, then $|u_2(x_1, x_2 + \gamma(\bar{x}_1 - x_1))| \leq |u_2(x_1, x_2 + \gamma)|$ since $|x_1 - \bar{x}_1| < 1$. Otherwise, if $u_2(x_1, x_2 + \gamma(\bar{x}_1 - x_1)) < 0$, then $|u_2(x_1, x_2 + \gamma(\bar{x}_1 - x_1))| \leq |u_2(x_1, x_2 - \gamma)|$. But we have

$$|u_2(x_1, x_2 + \gamma)| \leq |u_2(\bar{x}_1, x_2 + \gamma)| + |u_2(x_1, x_2 + \gamma) - u_2(\bar{x}_1, x_2 + \gamma)|$$

for every x_2. By Assumption 1(b), $|u_2(x_1, x_2 + \gamma) - u_2(\bar{x}_1, x_2 + \gamma)| \leq |u_1(x_1) - u_1(\bar{x}_1) < \delta$

whenever $|x_1 - \bar{x}_1| < 1$. Similarly,

$$|u_2(x_1, x_2 - \gamma)| \leq |u_2(\bar{x}_1, x_2 - \gamma)| + |u_2(x_1, x_2 - \gamma) - u_2(\bar{x}_1, x_2 - \gamma)| \leq |u_2(\bar{x}_1, x_2 - \gamma)| + \delta.$$

□

Claim A.2. The function f is absolutely integrable with respect to the Gaussian distribution $\mathcal{N}(\mu_2 - \gamma(\bar{x}_1 - \mu_1), \sigma^2)$.

Proof. This is because $x_2 \mapsto u_2(\bar{x}_1, x_2)$ is absolutely integrable with respect to each of the two Gaussian distributions $\mathcal{N}(\mu_2 - \gamma(\bar{x}_1 - \mu_1) + \gamma, \sigma^2)$ and $\mathcal{N}(\mu_2 - \gamma(\bar{x}_1 - \mu_1) - \gamma, \sigma^2)$, by Assumption 1(d). □

Together, these two claims show that for the family of functions $x_2 \mapsto u_2(x_1, x_2 + \gamma(\bar{x}_1 - x_1))$ for $|x_1 - \bar{x}_1| < 1$, f is an integrable dominating function with respect to the Gaussian distribution $\mathcal{N}(\mu_2 - \gamma(\bar{x}_1 - \mu_1), \sigma^2)$. Consider a sequence $(x_1^{(n)})_{n \in \mathbb{N}}$ with $x_1^{(n)} \to \bar{x}_1$. By continuity, $u_1(x_1^{(n)}) \to u_1(\bar{x}_1)$. For all large enough n, the functions

$$x_2 \mapsto u_2(x_1^{(n)}, x_2 + \gamma(\bar{x}_1 - x_1^{(n)}))$$

falls within the family mentioned before. Since these functions converge pointwise in x_2 to $x_2 \mapsto u_2(\bar{x}_1, x_2)$, the existence of the dominating function f implies the convergence of the integrals by dominated convergence theorem,

$$\mathbb{E}_{\tilde{X}_2 \sim \mathcal{N}(\mu_2 - \gamma(\bar{x}_1 - \mu_1), \sigma^2)}[u_2(x_1^{(n)}, \tilde{X}_2 + \gamma(\bar{x}_1 - x_1^{(n)}))] \to \mathbb{E}_{\tilde{X}_2 \sim \mathcal{N}(\mu_2 - \gamma(\bar{x}_1 - \mu_1), \sigma^2)}[u_2(\bar{x}_1, \tilde{X}_2)].$$

But

$$\begin{aligned}\mathbb{E}_\Psi[u_2(x_1^{(n)}, X_2)|X_1 = x_1^{(n)}] &= \mathbb{E}_{\tilde{\tilde{X}}_2 \sim \mathcal{N}(\mu_2 - \gamma(x_1^{(n)} - \mu_1), \sigma^2)}[u_2(x_1^{(n)}, \tilde{\tilde{X}}_2] \\ &= \mathbb{E}_{\tilde{X}_2 \sim \mathcal{N}(\mu_2 - \gamma(\bar{x}_1 - \mu_1), \sigma^2)}[u_2(x_1^{(n)}, \tilde{X}_2 + \gamma(\bar{x}_1 - x_1^{(n)})],\end{aligned}$$

which shows

$$\begin{aligned}\lim_{n \to \infty} \mathbb{E}_\Psi[u_2(x_1^{(n)}, X_2)|X_1 = x_1^{(n)}] &= \mathbb{E}_{\tilde{X}_2 \sim \mathcal{N}(\mu_2 - \gamma(\bar{x}_1 - \mu_1), \sigma^2)}[u_2(\bar{x}_1, \tilde{X}_2)] \\ &= \mathbb{E}_\Psi[u_2(\bar{x}_1, X_2)|X_1 = \bar{x}_1].\end{aligned}$$

This establishes that $D(x_1^{(n)}) \to D(\bar{x}_1)$, so D is continuous at $\bar{x}_1$.

Step 3: If $\gamma > 0$, then there are $x_1^{'} < x_1^{''}$ so that $D(x_1^{'}) < 0 < D(x_1^{''})$.

I show D is not always negative; the other statement is symmetric.

From $u_1(x_1^g) - u_2(x_1^g, x_2^b) > \kappa > 0$, we get that for any $x_1^{'} \geq x_1^g$, $x_2^{'} \leq x_2^b$,

$$\begin{aligned} u_1(x_1^{'}) - u_2(x_1^{'}, x_2^{'}) &\geq u_1(x_1^g) - u_2(x_1^g, x_2^{'}) \\ &\geq u_1(x_1^g) - u_2(x_1^g, x_2^b) > \kappa \end{aligned}$$

where the first inequality comes from Assumption 1(b) and the second one comes from Assumption 1(a). We have for any x_1,

$$\begin{aligned} D(x_1) =& u_1(x_1) - \mathbb{E}_\Psi[u_2(x_1, X_2)|X_1 = x_1] \\ =& \mathbb{P}_\Psi[X_2 \leq x_2^b | X_1 = x_1] \cdot (u_1(x_1) - \mathbb{E}_\Psi[u_2(x_1, X_2)|X_1 = x_1, X_2 \leq x_2^b]) \\ &+ \mathbb{P}_\Psi[X_2 > x_2^b | X_1 = x_1] \cdot (u_1(x_1) - \mathbb{E}_\Psi[u_2(x_1, X_2)|X_1 = x_1, X_2 > x_2^b]). \end{aligned}$$

When $x_1 \geq x_1^g$, $u_1(x_1) - \mathbb{E}_\Psi[u_2(x_1, X_2)|X_1 = x_1, X_2 \leq x_2^b] > \kappa$. Also, for $x_1 \geq x_1^g$,

$$u_1(x_1) - \mathbb{E}_\Psi[u_2(x_1, X_2)|X_1 = x_1, X_2 > x_2^b] \leq u_1(x_1^g) - \mathbb{E}_\Psi[u_2(x_1^g, X_2)|X_1 = x_1, X_2 > x_2^b].$$

But

$$\begin{aligned} & \mathbb{P}_\Psi[X_2 > x_2^b | X_1 = x_1] \cdot \mathbb{E}_\Psi[u_2(x_1^g, X_2)|X_1 = x_1, X_2 > x_2^b] \\ =& \mathbb{E}_\Psi[\mathbf{1}\{X_2 > x_2^b\} \cdot u_2(x_1^g, X_2)|X_1 = x_1] \\ =& \mathbb{E}_{\tilde{\tilde{X}}_2 \sim \mathcal{N}(\mu_2 - \gamma(x_1 - \mu_1); \sigma^2)}[\mathbf{1}\{\tilde{\tilde{X}}_2 > x_2^b\} \cdot u_2(x_1^g, \tilde{\tilde{X}}_2)] \\ =& \mathbb{E}_{\tilde{X}_2 \sim \mathcal{N}(\mu_2; \sigma^2)}[\mathbf{1}\{\tilde{X}_2 - \gamma(x_1 - \mu_1) > x_2^b\} \cdot u_2(x_1^g, \tilde{X}_2 - \gamma(x_1 - \mu_1))] \\ \leq& \mathbb{E}_{\tilde{X}_2 \sim \mathcal{N}(\mu_2; \sigma^2)}[\mathbf{1}\{\tilde{X}_2 - \gamma(x_1 - \mu_1) > x_2^b\} \cdot |u_2(x_1^g, \tilde{X}_2)|] \end{aligned}$$

when $x_1 > \mu_1$. Since $\mathbb{E}_{\tilde{X}_2 \sim \mathcal{N}(\mu_2; \sigma^2)}[|u_2(x_1^g, \tilde{X}_2)|]$ exists and is finite by Assumption 1(d), as $x_1 \to \infty$ we must have

$$\mathbb{E}_{\tilde{X}_2 \sim \mathcal{N}(\mu_2; \sigma^2)}[\mathbf{1}\{\tilde{X}_2 - \gamma(x_1 - \mu_1) > x_2^b\} \cdot |u_2(x_1^g, \tilde{X}_2)|] \to 0$$

as the indicator converges to 0 everywhere, given that $\gamma > 0$. So this shows for all large enough x_1, $D(x_1) \geq \kappa/2 > 0$. □

The desired Lemma A.1 follows readily from Lemma A.2.

Proof. Applying Lemma A.2 and using the fact that $\gamma > 0$, D changes sign and is strictly increasing and continuous. So, there exists a unique $c^* \in \mathbb{R}$ satisfying $D(c^*) = 0$. It is clear that the best stopping strategy under Ψ is the cutoff strategy that stops after every $x_1 > c^*$ and continues after every $x_1 < c^*$. This establishes property (ii) of the optimal strategy. Properties (i) and (iii) follow from the fact that D is strictly increasing. □

A.2.2 Cutoff Threshold Increasing in μ_2

In the second part, I prove the lemma:

Lemma A.3. Under Assumption 1, for any $\mu_1 \in \mathbb{R}$ and $\gamma > 0$, the indifference threshold $C(\mu_1, \mu_2; \gamma)$ is strictly increasing in μ_2.

Proof. Let $\hat{\mu}_1, \hat{\mu}_2, \hat{\hat{\mu}}_2 \in \mathbb{R}$ with $\hat{\hat{\mu}}_2 > \hat{\mu}_2$. I show that $C(\hat{\mu}_1\hat{\mu}_2; \gamma) < C(\hat{\mu}_1\hat{\hat{\mu}}_2; \gamma)$.

By Lemma A.1, the threshold $C(\hat{\mu}_1, \hat{\mu}_2; \gamma)$ is characterized by the indifference condition,

$$u_1(C(\hat{\mu}_1, \hat{\mu}_2; \gamma)) = \mathbb{E}_{\tilde{X}_2 \sim \mathcal{N}(\hat{\mu}_2 - \gamma(C(\hat{\mu}_1, \hat{\mu}_2; \gamma) - \hat{\mu}_1), \sigma^2)}[u_2(C(\hat{\mu}_1, \hat{\mu}_2; \gamma), \tilde{X}_2)]$$

But if agent were to instead believe $(\hat{\mu}_1\hat{\hat{\mu}}_2)$ where $\hat{\hat{\mu}}_2 > \hat{\mu}_2$, then the conditional distribution of X_2 given $X_1 = C(\hat{\mu}_1, \hat{\mu}_2; \gamma)$ would be $\mathcal{N}(\hat{\hat{\mu}}_2 - \gamma(C(\hat{\mu}_1, \hat{\mu}_2; \gamma) - \hat{\mu}_1), \sigma^2)$. We have

$$u_1(C(\hat{\mu}_1, \hat{\mu}_2; \gamma)) < \mathbb{E}_{\tilde{X}_2 \sim \mathcal{N}(\hat{\hat{\mu}}_2 - \gamma(C(\hat{\mu}_1, \hat{\mu}_2; \gamma) - \hat{\mu}_1), \sigma^2)}[u_2(C(\hat{\mu}_1, \hat{\mu}_2; \gamma), \tilde{X}_2)]$$

by Assumption 1(a). This means $C(\hat{\mu}_1, \hat{\mu}_2; \gamma) < C(\hat{\mu}_1, \hat{\hat{\mu}}_2; \gamma)$ by Lemma A.1, as only values of X_1 below $C(\hat{\mu}_1, \hat{\hat{\mu}}_2; \gamma)$ lead to strict preference for continuing. □

A.2.3 Proof of Lemma 1

Proof. In fact, this lemma holds for any $\mu_1 \in \mathbb{R}$.

For $\mu_2'' > \mu_2'$, write the corresponding optimal cutoffs as $c'' := C(\mu_1, \mu_2''; \gamma)$ and $c' := C(\mu_1, \mu_2'; \gamma)$. I show that $|c'' - c'| < \ell|\mu_2'' - \mu_2'|$.

Under the model $\Psi(\mu_1, \mu_2''; \gamma)$, the expected continuation payoff after $X_1 = c' + \ell(\mu_2'' - \mu_1')$ is

$$\begin{aligned}
&\mathbb{E}_{\tilde{\tilde{X}}_2 \sim \mathcal{N}(\mu_2'' - \gamma(c' + \ell(\mu_2'' - \mu_1') - \mu_1)\sigma^2)}[u_2(c' + \ell(\mu_2'' - \mu_1'), \tilde{\tilde{X}}_2] \\
=&\mathbb{E}_{\tilde{X}_2 \sim \mathcal{N}(\mu_2' - \gamma(c' - \mu_1)\sigma^2)}[u_2(c' + \ell(\mu_2'' - \mu_1'), \tilde{X}_2 + (\mu_2'' - \mu_1') - \gamma\ell(\mu_2'' - \mu_1')] \\
=&\mathbb{E}_{\tilde{X}_2 \sim \mathcal{N}(\mu_2' - \gamma(c' - \mu_1)\sigma^2)}[u_2(c' + \ell d, \tilde{X}_2 + (1 - \gamma\ell)d)]
\end{aligned}$$

where we put $d = |\mu_2'' - \mu_2'| > 0$. From Assumption 2, for every $x_2 \in \mathbb{R}$, $u_2(c' + \ell d, x_2 + (1 - \gamma\ell)d) - u_2(c', x_2) < u_1(c' + \ell d) - u_1(c')$. This means

$$\mathbb{E}_{\tilde{X}_2 \sim \mathcal{N}(\mu_2' - \gamma(c' - \mu_1)\sigma^2)}[u_2(c' + \ell d, \tilde{X}_2 + (1 - \gamma\ell)d) - u_2(c', \tilde{X}_2)] < u_1(c' + \ell d) - u_1(c')$$
$$\mathbb{E}_{\tilde{X}_2 \sim \mathcal{N}(\mu_2' - \gamma(c' - \mu_1)\sigma^2)}[u_2(c' + \ell d, \tilde{X}_2 + (1 - \gamma\ell)d)] - u_1(c' + \ell d) < \mathbb{E}_{\tilde{X}_2 \sim \mathcal{N}(\mu_2' - \gamma(c' - \mu_1)\sigma^2)}[u_2(c', \tilde{X}_2)] - u_1(c').$$

The cutoff c' satisfies the indifference condition,

$$u_1(c') = \mathbb{E}_{\tilde{X}_2 \sim \mathcal{N}(\mu_2' - \gamma(c' - \mu_1), \sigma^2)}[u_2(c', \tilde{X}_2)],$$

so RHS is 0. But LHS is the difference between expected continuation payoff and stopping payoff at $X_1 = c' + \ell(\mu_2'' - \mu_1')$ under the model $\Psi(\mu_1, \mu_2''; \gamma)$, which shows the agent strictly prefers stopping. This means $c'' < c' + \ell(\mu_2'' - \mu_1')$. But $\mu_2 \mapsto C(\mu_1, \mu_2; \gamma)$ is increasing by Lemma A.3, which means $c'' > c'$. Together, these two inequalities imply $|c'' - c'| < \ell(\mu_2'' - \mu_1')$. $\square$

A.2.4 Lipschitz Continuity with Constant $1/\gamma$

Now I prove the lemma:

Lemma A.4. Under Assumption 1, for every $\gamma > 0$ and $\mu_1 \in \mathbb{R}$, $\mu_2 \mapsto C(\mu_1, \mu_2; \gamma)$ is Lipschitz continuous with Lipschitz constant $1/\gamma$.

Proof. The proof of Lemma 1 also applies when $\ell = \frac{1}{\gamma}$, which implies that when the inequality in Assumption 2 is satisfied with $\ell = \frac{1}{\gamma}$, $\mu_2 \mapsto C(\mu_1, \mu_2; \gamma)$ is Lipschitz continuous with Lipschitz constant $1/\gamma$. But this reduces the inequality to $u_1(x_1 + \frac{1}{\gamma}d) - u_1(x_1) \geq u_2(x_1 + \frac{1}{\gamma}d, x_2) - u_2(x_1, x_2)$, which is true by Assumption 1(b). $\square$

A.3 Proof of Claim 2

Proof. For $d > 0$,

$$u_1(x_1 + \frac{1}{1+\gamma}d) - u_1(x_1) = \frac{1}{1+\gamma}d$$

while

$$\begin{aligned}
&u_2(x_1 + \frac{1}{1+\gamma}d, x_2 + (1 - \frac{\gamma}{1+\gamma})d) - u_2(x_1, x_2) \\
=&u_2(x_1 + \frac{1}{1+\gamma}d, x_2 + \frac{1}{1+\gamma}d) - u_2(x_1, x_2) \\
=&q \max(x_1 + \frac{1}{1+\gamma}d, x_2 + \frac{1}{1+\gamma}d) + (1-q)(x_2 + \frac{1}{1+\gamma}d) \\
&- q \max(x_1, x_2) - (1-q)x_2 \\
=&q\frac{1}{1+\gamma}d + (1-q)\frac{1}{1+\gamma}d = \frac{1}{1+\gamma}d.
\end{aligned}$$

This shows that when $\ell = \frac{1}{1+\gamma}$, we have $u_1(x_1 + \ell d) - u_1(x_1) = u_2(x_1 + \ell d, x_2 + (1 - \gamma\ell)d) - u_2(x_1, x_2)$ for every $x_1, x_2 \in \mathbb{R}$, $d > 0$. $\square$

A.4 Proof of Proposition 2

Proof. Rewrite Definition 5 as

$$\begin{aligned}
&\int_c^\infty \phi(x_1;\mu_1^\bullet,\sigma^2)\cdot \ln\left(\frac{\phi(x_1;\mu_1^\bullet,\sigma^2)}{\phi(x_1;\mu_1,\sigma^2)}\right)dx_1\\
&+\int_{-\infty}^c \phi(x_1;\mu_1^\bullet,\sigma^2)\cdot\int_{-\infty}^\infty \phi(x_2;\mu_2^\bullet,\sigma^2)\cdot\ln\left[\frac{\phi(x_1;\mu_1^\bullet,\sigma^2)}{\phi(x_1;\mu_1,\sigma^2)}\right]dx_2dx_1\\
&+\int_{-\infty}^c \phi(x_1;\mu_1^\bullet,\sigma^2)\cdot\int_{-\infty}^\infty \phi(x_2;\mu_2^\bullet,\sigma^2)\cdot\ln\left[\frac{\phi(x_2;\mu_2^\bullet,\sigma^2)}{\phi(x_2;\mu_2-\gamma(x-\mu_1),\sigma^2)}\right]dx_2dx_1
\end{aligned}$$

which is:

$$\begin{aligned}
&\int_{-\infty}^\infty \phi(x_1;\mu_1^\bullet,\sigma^2)\cdot \ln\left(\frac{\phi(x_1;\mu_1^\bullet,\sigma^2)}{\phi(x_1;\mu_1,\sigma^2)}\right)dx_1\\
&+\int_{-\infty}^c \phi(x_1;\mu_1^\bullet,\sigma^2)\cdot\int_{-\infty}^\infty \phi(x_2;\mu_2^\bullet,\sigma^2)\ln\left[\frac{\phi(x_2;\mu_2^\bullet,\sigma^2)}{\phi(x_2;\mu_2-\gamma(x_1-\mu_1),\sigma^2)}\right]dx_2dx_1
\end{aligned}$$

The KL divergence between $\mathcal{N}(\mu_{\text{true}},\sigma^2_{\text{true}})$ and $\mathcal{N}(\mu_{\text{model}},\sigma^2_{\text{model}})$ is

$$\ln\frac{\sigma_{\text{model}}}{\sigma_{\text{true}}}+\frac{\sigma^2_{\text{true}}+(\mu_{\text{true}}-\mu_{\text{model}})^2}{2\sigma^2_{\text{model}}}-\frac{1}{2},$$

so we may simplify the first term and the inner integral of the second term:

$$\frac{(\mu_1-\mu_1^\bullet)^2}{2\sigma^2}+\int_{-\infty}^c \phi(x_1;\mu_1^\bullet,\sigma^2)\cdot\left[\frac{\sigma^2+(\mu_2-\gamma(x_1-\mu_1)-\mu_2^\bullet)^2}{2\sigma^2}-\frac{1}{2}\right]dx_1.$$

Dropping constant terms not depending on μ_1 and μ_2 and multiplying by σ^2, we get a simplified expression of the objective,

$$\xi(\mu_1,\mu_2):=\frac{(\mu_1-\mu_1^\bullet)^2}{2}+\int_{-\infty}^c \phi(x_1;\mu_1^\bullet,\sigma^2)\cdot\left[\frac{(\mu_2-\gamma(x_1-\mu_1)-\mu_2^\bullet)^2}{2}\right]dx_1$$

We have the partial derivatives by differentiating under the integral sign,

$$\frac{\partial\xi}{\partial\mu_2}=\int_{-\infty}^c \phi(x_1;\mu_1^\bullet,\sigma^2)\cdot(\mu_2-\gamma(x_1-\mu_1)-\mu_2^\bullet)dx_1$$

$$\begin{aligned}
\frac{\partial\xi}{\partial\mu_1}&=(\mu_1-\mu_1^\bullet)+\gamma\int_{-\infty}^c \phi(x_1;\mu_1^\bullet,\sigma^2)\cdot(\mu_2-\gamma(x_1-\mu_1)-\mu_2^\bullet)dx_1\\
&=(\mu_1-\mu_1^\bullet)+\gamma\frac{\partial\xi}{\partial\mu_2}
\end{aligned}$$

By the first order conditions, at the minimum (μ_1^*, μ_2^*), we must have:

$$\frac{\partial \xi}{\partial \mu_2}(\mu_1^*, \mu_2^*) = \frac{\partial \xi}{\partial \mu_1}(\mu_1^*, \mu_2^*) = 0 \Rightarrow \mu_1^* = \mu_1^\bullet$$

So μ_2^* satisfies $\frac{\partial \xi}{\partial \mu_2}(\mu_1^\bullet, \mu_2^*) = 0$, which by straightforward algebra shows

$$\mu_2^*(c) = \mu_2^\bullet - \gamma\left(\mu_1^\bullet - \mathbb{E}\left[X_1 \mid X_1 \leq c\right]\right).$$

□

A.5 Proof of Proposition 3

Proof. Consider the map $\mathcal{I}$ as discussed in the text,

$$\mathcal{I}(\mu_2) := \mu_2^\bullet - \gamma\left(\mu_1^\bullet - \mathbb{E}\left[X_1 | X_1 \leq C(\mu_1^\bullet, \mu_2; \gamma)\right]\right).$$

If $\hat{\mu}_2$ is a fixed point of $\mathcal{I}$, then there is a steady state with $\mu_1^\infty = \mu_1^\bullet$, $\mu_2^\infty = \hat{\mu}_2$, $c^\infty = C(\mu_1^\bullet, \hat{\mu}_2; \gamma)$. This follows readily from definition of $\mathcal{I}$ and the closed-form expression of pseudo-true fundamentals from Proposition 2. So, existence of steady states follows from existence of fixed points of $\mathcal{I}$.

Conversely, suppose $(\mu_1^\infty, \mu_2^\infty, c^\infty)$ is a steady state. From Proposition 2, $\mu_1^\infty = \mu_1^*(c^\infty) = \mu_1^\bullet$, $\mu_2^\infty = \mu_2^*(c^\infty) = \mu_2^\bullet - \gamma\left(\mu_1^\bullet - \mathbb{E}\left[X_1 | X_1 \leq c^\infty\right]\right)$, and $c^\infty = C(\mu_1^\infty, \mu_2^\infty; \gamma) = C(\mu_1^\bullet, \mu_2^\infty; \gamma)$. That is to say, $\mu_2^\infty = \mu_2^\bullet - \gamma\left(\mu_1^\bullet - \mathbb{E}\left[X_1 | X_1 \leq C(\mu_1^\bullet, \mu_2^\infty; \gamma)\right]\right) = \mathcal{I}(\mu_2^\infty)$, so μ_2^∞ is a fixed point of $\mathcal{I}$. So, uniqueness of steady states follows from uniqueness of fixed points of $\mathcal{I}$.

I show $\mathcal{I}$ is a contraction mapping with modulus $\ell\gamma$. We have

$$\mathcal{I}(\mu_2^{'}) - \mathcal{I}(\mu_2^{''}) = \gamma \cdot \left(\mathbb{E}\left[X_1 | X_1 \leq C(\mu_1^\bullet, \mu_2^{'}; \gamma)\right] - \mathbb{E}\left[X_1 | X_1 \leq C(\mu_1^\bullet, \mu_2^{''}; \gamma)\right]\right).$$

By formula of the mean of a truncated Gaussian random variable, when $X_1 \sim \mathcal{N}(\mu_1^\bullet, \sigma^2)$ and $c \in \mathbb{R}$, we get $\mathbb{E}[X_1 | X_1 \leq c] = \mu_1^\bullet - \left(\frac{\phi((c-\mu_1^\bullet)/\sigma)}{\Phi((c-\mu_1^\bullet)/\sigma)}\right)\sigma$. Therefore,

$$\begin{aligned}\mathcal{I}(\mu_2^{'}) - \mathcal{I}(\mu_2^{''}) &= \gamma \cdot \left(\left(\mu_1^\bullet - \frac{\phi((C(\mu_1^\bullet, \mu_2^{'}; \gamma)/\sigma))}{\Phi((C(\mu_1^\bullet, \mu_2^{'}; \gamma)/\sigma))} \cdot \sigma\right) - \left(\mu_1^\bullet - \frac{\phi((C(\mu_1^\bullet, \mu_2^{''}; \gamma)/\sigma))}{\Phi((C(\mu_1^\bullet, \mu_2^{''}; \gamma)/\sigma))} \cdot \sigma\right)\right) \\ &= -\gamma\sigma \cdot \left(\frac{\phi(C(\mu_1^\bullet, \mu_2^{'}; \gamma)/\sigma)}{\Phi(C(\mu_1^\bullet, \mu_2^{'}; \gamma)/\sigma)} - \frac{\phi(C(\mu_1^\bullet, \mu_2^{''}; \gamma)/\sigma)}{\Phi(C(\mu_1^\bullet, \mu_2^{''}; \gamma)/\sigma)}\right).\end{aligned}$$

The function $z \mapsto \frac{\phi(z)}{1-\Phi(z)}$ is the Gaussian inverse Mills ratio and its derivative is bounded by 1 in magnitude[21]. By symmetry this also applies to the function $z \mapsto \frac{\phi(z)}{\Phi(z)}$. This means

[21]See for example Corollary 1.6 in Pinelis (2018)

$\left|\frac{\phi(z')}{\Phi(z')} - \frac{\phi(z'')}{\Phi(z'')}\right| \le |z' - z''|$. So we have

$$\left|\frac{\phi(C(\mu_1^\bullet, \mu_2'; \gamma)/\sigma)}{\Phi(C(\mu_1^\bullet, \mu_2'; \gamma)/\sigma)} - \frac{\phi(C(\mu_1^\bullet, \mu_2''; \gamma)/\sigma)}{\Phi(C(\mu_1^\bullet, \mu_2''; \gamma)/\sigma)}\right| \le \frac{1}{\sigma} \cdot \left|C(\mu_1^\bullet, \mu_2'; \gamma) - C(\mu_1^\bullet, \mu_2''; \gamma)\right| \le \frac{1}{\sigma} \cdot \ell \cdot |\mu_2' - \mu_2''|$$

by Lemma 1. This then showing $|\mathcal{I}(\mu_2') - \mathcal{I}(\mu_2'')| \le \gamma\ell \cdot |\mu_2' - \mu_2''|$ for all $\mu_2', \mu_2'' \in \mathbb{R}$. So Υ is a contraction mapping with contraction constant $\gamma\ell \in (0, 1)$ and the proposition readily follows from properties of contraction mappings. $\square$

A.6 Proof of Proposition 5

Proof. Suppose $(\mu_1^\bullet, \mu_2^\infty, c^\infty)$ is a steady state. If $c^\bullet = \infty$, then $c^\infty < c^\bullet$ trivially as $c^\infty \in \mathbb{R}$.

Now suppose $c^\bullet \neq \infty$. By Lemma 1, agent is indifferent between stopping and continuing after $X_1 = c^\infty$ under the subjective model $\Psi(\mu_1^\bullet, \mu_2^\infty; \gamma)$. This implies

$$\begin{aligned} u_1(c^\infty) &= \mathbb{E}_{\Psi(\mu_1^\bullet, \mu_2^\infty; \gamma)}[u_2(c^\infty, X_2)|X_1 = c^\infty] \\ &= \mathbb{E}_{\tilde{X}_2 \sim \mathcal{N}(\mu_2^\infty - \gamma(c^\infty - \mu_1^\bullet), \sigma^2)}[u_2(c^\infty, \tilde{X}_2)] \end{aligned}$$

By the definition of steady state, $\mu_2^\infty = \mu_2^*(c^\infty) = \mu_2^\bullet - \gamma\left(\mu_1^\bullet - \mathbb{E}\left[X_1 | X_1 \le c^\infty\right]\right)$. But we have

$$\mu_2^\infty - \gamma(c^\infty - \mu_1^\bullet) < \mu_2^\infty - \gamma(\mathbb{E}\left[X_1|X_1 \le c^\infty\right] - \mu_1^\bullet) = \mu_2^\bullet$$

since $c^\infty > \mathbb{E}\left[X_1|X_1 \le c^\infty\right]$.

Therefore, $\mathcal{N}(\mu_2^\infty - \gamma(c^\infty - \mu_1^\bullet), \sigma^2)$ is first-order stochastically dominated by $\mathcal{N}(\mu_2^\bullet, \sigma^2)$. Since u_2 is strictly increasing in its second argument by Assumption 1(a), we therefore have

$$u_1(c^\infty) < \mathbb{E}_{\tilde{X}_2 \sim \mathcal{N}(\mu_2^\bullet, \sigma^2)}[u_2(c^\infty, \tilde{X}_2)].$$

The LHS is the objective payoff of stopping at c^∞ while the RHS is the objective expected payoff of continuing at c^∞. Since the best stopping strategy under the objective model $\Psi^\bullet$ has the cutoff form, we must have $c^\infty < c^\bullet$.[22] $\square$

A.7 Proof of Theorem 1′

The hypotheses of Theorem 1′ will be maintained throughout this section.

[22]In particular this implies if there exists at least one steady state, then $c^\bullet \neq -\infty$.

A.7.1 Optimality of Cutoff Strategies

I first develop an extension of Lemma A.1. I show that for an agent who knows $\mu_1^\bullet$ and has some belief over μ_2 with supported bounded by $[\underline{\mu}_2, \bar{\mu}_2]$, there exists a cutoff strategy that uniquely maximizes payoff across all cutoff strategies, so the "myopically optimal" cutoff strategy is well defined. Furthermore, this myopically optimal cutoff strategy also achieves weakly larger expected payoff compared to any arbitrary stopping strategy[23]. So, restriction to cutoff strategies is without loss.

Lemma A.5. *For an agent who knows $\mu_1^\bullet$ and who holds some belief $\nu \in \Delta([\underline{\mu}_2, \bar{\mu}_2])$ about second-period fundamental, there exists $c^* \in \mathbb{R}$ such that: (i) the cutoff strategy S_{c^*} achieves weakly higher expected payoff than any other (not necessarily cutoff-based) stopping strategy $S : \mathbb{R} \to \{Stop,\ Continue\}$; (ii) for any other $c' \neq c^*$, S_{c^*} achieves strictly higher expected payoff than $S_{c'}$.*

A.7.2 The Log Likelihood Process

Next, I define the processes of data log likelihood (for a given fundamental). For each $\mu_2 \in [\underline{\mu}_2, \bar{\mu}_2]$, let $\ell_t(\mu_2)(\omega)$ be the log likelihood that the true second-period fundamental is μ_2 and histories $(\tilde{H}_s)_{s \le t}(\omega)$ are generated by the end of round t. It is given by

$$\ell_t(\mu_2)(\omega) := \ln(g(\mu_2)) + \sum_{s=1}^{t} \ln(\text{lik}(\tilde{H}_s(\omega); \mu_2))$$

where $\text{lik}(x_1, \varnothing; \mu_2) := \phi(x_1; \mu_1^\bullet, \sigma^2)$ and $\text{lik}(x_1, x_2; \mu_2) := \phi(x_1; \mu_1^\bullet, \sigma^2) \cdot \phi(x_2; \mu_2 - \gamma(x_1 - \mu_1^\bullet); \sigma^2)$.

I record a useful decomposition of $\ell_t'(\mu_2)$, the derivative of the log-likelihood process.

Define two stochastic processes:

$$\varphi_s(\mu_2) := \sigma^{-2} \cdot (X_{2,s} - \mu_2 + \gamma(X_{1,s} - \mu_1^\bullet)) \cdot \mathbf{1}\{X_{1,s} \le \tilde{C}_s\}$$

$$\bar{\varphi}_s(\mu_2) := \sigma^{-2}\mathbb{P}[X_1 \le \tilde{C}_s] \cdot (\mu_2^\bullet - \mu_2 - \gamma(\mu_1^\bullet - \mathbb{E}[X_1 | X_1 \le \tilde{C}_s])),$$

with a slight abuse of notation, $\mathbb{P}[X_1 \le x]$ means the probability that each first-period draw falls below x, and $\mathbb{E}[X_1|X_1 \le x]$ the conditional expectation of the first draw given that it falls below x. Note that $\bar{\varphi}_s(\mu_2)$ is measurable with respect to $\mathcal{F}_{s-1}$, since (C_t) is a predictable process. Write $\xi_s(\mu_2) := \varphi_s(\mu_2) - \bar{\varphi}_s(\mu_2)$ and $y_t(\mu_2) := \sum_{s=1}^{t} \xi_s(\mu_2)$. Write $z_t(\mu_2) := \sum_{s=1}^{t} \bar{\varphi}_s(\mu_2)$.

Lemma A.6. $\ell_t'(\mu_2) = \frac{g'(\mu_2)}{g(\mu_2)} + y_t(\mu_2) + z_t(\mu_2)$

[23] One can construct other stopping strategies with the same expected payoff by, for example, modifying the stopping decision of the optimal cutoff strategy at finitely many x_1.

Proof. We may expand $\ell_t(\mu_2)$ as

$$\ln(g(\mu_2)) + \sum_{s=1}^{t} \ln(\phi(X_{1,s}; \mu_1^\bullet, \sigma^2)) + \sum_{s=1}^{t} \ln(\phi(X_{2,s}; \mu_2 - \gamma(X_{1,s} - \mu_1^\bullet); \sigma^2)) \cdot \mathbf{1}\{X_{1,s} \le \tilde{C}_s\}.$$

The derivative of the first term is $\frac{g'(\mu_2)}{g(\mu_2)}$. The second term does not depend on μ_2. In the third term, we substitute the Gaussian density to get:

$$\sum_{s=1}^{t} \ln((2\pi\sigma^2)^{-1/2}) \cdot \mathbf{1}\{X_{1,s} \le \tilde{C}_s\} + \sum_{s=1}^{t} -\frac{(X_{2,s} - \mu_2 + \gamma(X_{1,s} - \mu_1^\bullet))^2}{2\sigma^2} \cdot \mathbf{1}\{X_{1,s} \le \tilde{C}_s\}.$$

Its derivative with respect to μ_2 is $\varphi_s(\mu_2)$. So in sum, $\ell_t'(\mu_2) = \frac{g'(\mu_2)}{g(\mu_2)} + \sum_{s=1}^{t} \varphi_s(\mu_2)$. The lemma then follows from simple rearrangements. $\square$

Now I derive two results about the $\xi_t(\mu_2)$ processes for different values of μ_2.

Lemma A.7. *There exists $\kappa_\xi < \infty$ so that for every $\mu_2 \in [\underline{\mu}_2, \bar{\mu}_2]$ and for every $t \ge 1$, $\omega \in \Omega$, $\mathbb{E}[\xi_t^2(\mu_2)|\mathcal{F}_{t-1}](\omega) \le \kappa_\xi$.*

Proof. Note that $\bar{\varphi}_t(\mu_2)$ is measurable with respect to $\mathcal{F}_{t-1}$. Also, $\varphi_t(\mu_2)|\mathcal{F}_{t-1} = \varphi_t(\mu_2)|\tilde{C}_t$, because by independence of X_t from $(X_s)_{s=1}^{t-1}$, the only information that $\mathcal{F}_{t-1}$ contains about $\varphi_t(\mu_2)$ is in determining the cutoff threshold C_t.

At a sample path ω so that $\tilde{C}_t(\omega) = c \in \mathbb{R}$,

$$\begin{aligned}
\mathbb{E}[\varphi_t(\mu_2)|\mathcal{F}_{t-1}](\omega) &= \mathbb{E}[\sigma^{-2} \cdot (X_2 - \mu_2 + \gamma(X_1 - \mu_1^\bullet))\mathbf{1}\{X_1 \le c\}] \\
&= \sigma^{-2}(\mathbb{E}[(X_2 - \mu_2) \cdot \mathbf{1}\{X_1 \le c\}] + \mathbb{E}[\gamma(X_1 - \mu_1^\bullet) \cdot \mathbf{1}\{X_1 \le c\}]) \\
&= \sigma^{-2}\mathbb{P}[X_1 \le c] \cdot (\mu_2^\bullet - \mu_2 - \gamma(\mu_1^\bullet - \mathbb{E}[X_1|X_1 \le c])),
\end{aligned}$$

where we used the fact that $X_{1,t}$ and $X_{2,t}$ are independent. This shows that $\mathbb{E}[\varphi_t(\mu_2)|\mathcal{F}_{t-1}](\omega) = \bar{\varphi}_t(\mu_2)(\omega)$. Since this holds regardless of c, we get that $\mathbb{E}[\varphi_t(\mu_2)|\mathcal{F}_{t-1}] = \bar{\varphi}_t(\mu_2)$ for all ω, that is to say

$$\begin{aligned}
\mathbb{E}[\xi_t^2(\mu_2)|\mathcal{F}_{t-1}] &= \mathrm{Var}[\varphi_t(\mu_2)|\mathcal{F}_{t-1}] \\
&\le \mathbb{E}[\varphi_t^2(\mu_2)|\mathcal{F}_{t-1}] \\
&= \sigma^{-4} \cdot \mathbb{E}[(X_{2,s} - \mu_2 + \gamma(X_{1,s} - \mu_1^\bullet))^2 \cdot \mathbf{1}\{X_{1,s} \le \tilde{C}_s\}] \\
&\le \sigma^{-4} \cdot \mathbb{E}[(X_2 - \mu_2 + \gamma(X_1 - \mu_1^\bullet))^2]
\end{aligned}$$

The RHS of the final line is independent of ω and t, while $\mu_2 \mapsto \mathbb{E}[(X_2 - \mu_2 + \gamma(X_1 - \mu_1^\bullet))^2]$ is a continuous function on $[\underline{\mu}_2, \bar{\mu}_2]$. Therefore it is bounded uniformly by some $\kappa_\xi < \infty$, which also provides a bound for $\mathbb{E}[\xi_t^2(\mu_2)|\mathcal{F}_{t-1}](\omega)$ for every t, ω and $\mu_2 \in [\underline{\mu}_2, \bar{\mu}_2]$. $\square$

Lemma A.8. *For every $t \geq 1$, $\mu_2 \in [\underline{\mu}_2, \bar{\mu}_2]$ and $\omega \in \Omega$, $|\xi_t^{'}(\mu_2)(\omega)| \leq 2/\sigma^2$.*

Proof. We have $|\xi_t^{'}(\mu_2)| = |\sigma^{-2} \cdot \mathbf{1}\{X_{1,t} \leq \tilde{C}_t\} + \sigma^{-2}\mathbb{P}[X_1 \leq \tilde{C}_t]| \leq 2/\sigma^{-2}$ since the terms $\mathbf{1}\{X_{1,t} \leq \tilde{C}_t\}$ and $\mathbb{P}[X_1 \leq \tilde{C}_t]$ are both bounded by 1 at every ω. $\square$

A.7.3 Heidhues, Koszegi, and Strack (2018)'s Law of Large Numbers

I use a statistical result from Heidhues, Koszegi, and Strack (2018) to show that the y_t term in the decomposition of $\ell_t^{'}$ almost surely converges to 0 in the long run, and furthermore this convergence is uniform on $[\underline{\mu}_2, \bar{\mu}_2]$. This lets me focus on summands of the form $\bar{\varphi}_s(\mu_2)$, which can be interpreted as the *expected* contribution to the log likelihood derivative from round s data. This lends tractability to the problem as $\bar{\varphi}_s(\mu_2)$ only depends on $\tilde{C}_s$, but not on $X_{1,s}$ or $X_{2,s}$.

Lemma A.9. *For every $\mu_2 \in [\underline{\mu}_2, \bar{\mu}_2]$, $\lim_{t\to\infty} |\frac{y_t(\mu_2)}{t}| = 0$ almost surely.*

Proof. Heidhues, Koszegi, and Strack (2018)'s Proposition 10 shows that if (y_t) is a martingale such that there exists some constant $v \geq 0$ satisfying $[y]_t \leq vt$ almost surely, where $[y]_t$ is the quadratic variation of (y_t), then almost surely $\lim_{t\to\infty} \frac{y_t}{t} = 0$.

Consider the process $y_t(\mu_2)$ for a fixed $\mu_2 \in [\underline{\mu}_2, \bar{\mu}_2]$. By definition $y_t = \sum_{s=1}^t \varphi_s(\mu_2) - \bar{\varphi}_s(\mu_2)$. As established in the proof of Lemma A.7, for every s, $\bar{\varphi}_s(\mu_2) = \mathbb{E}[\varphi_s(\mu_2)|\mathcal{F}_{s-1}]$. So for $t^{'} < t$,

$$\begin{aligned}
\mathbb{E}[y_t(\mu_2)|\mathcal{F}_{t'}] &= \sum_{s=1}^{t'} \varphi_s(\mu_2) - \bar{\varphi}_s(\mu_2) + \mathbb{E}[\sum_{s=t'+1}^{t} \varphi_s(\mu_2) - \bar{\varphi}_s(\mu_2)|\mathcal{F}_{t'}] \\
&= \sum_{s=1}^{t'} \varphi_s(\mu_2) - \bar{\varphi}_s(\mu_2) + \sum_{s=t'+1}^{t} \mathbb{E}[\mathbb{E}[\varphi_s(\mu_2) - \bar{\varphi}_s(\mu_2)|\mathcal{F}_{s-1}] \mid \mathcal{F}_{t'}] \\
&= \sum_{s=1}^{t'} \varphi_s(\mu_2) - \bar{\varphi}_s(\mu_2) + 0 \\
&= y_{t'}(\mu_2).
\end{aligned}$$

This shows $(y_t(\mu_2))$ is a martingale. Also,

$$\begin{aligned}
[y(\mu_2)]_t &= \sum_{s=1}^{t-1} \mathbb{E}[(y_s(\mu_2) - y_{s-1}(\mu_2))^2|\mathcal{F}_{s-1}] \\
&= \sum_{s=1}^{t-1} \mathbb{E}[\xi_s^2(\mu_2)|\mathcal{F}_{s-1}] \\
&\leq \kappa_\xi \cdot t
\end{aligned}$$

by Lemma A.7. Therefore Heidhues, Koszegi, and Strack (2018) Proposition 10 applies. $\square$

Lemma A.10. $\lim_{t\to\infty} \sup_{\mu_2 \in [\underline{\mu}_2, \bar{\mu}_2]} |\frac{y_t(\mu_2)}{t}| = 0$ *almost surely.*

Proof. From the proof of Lemma 11 in Heidhues, Koszegi, and Strack (2018), it suffices to find a sequence of random variables B_t such that $\sup_{\mu_2 \in [\underline{\mu}_2, \bar{\mu}_2]} |\xi_t^{'}(\mu_2)| \leq B_t$ almost surely, $\sup_{t\geq 1} \frac{1}{t} \sum_{s=1}^{t} \mathbb{E}[B_s] < \infty$, and $\lim_{t\to\infty} \frac{1}{t} \sum_{s=1}^{t} (B_s - \mathbb{E}[B_s]) = 0$. But Lemma A.8 establishes the constant random variable $B_t = 2/\sigma^2$ as a bound on $\xi_t^{'}(\mu_2)$ for every t, μ_2, ω, which satisfies these requirements. $\square$

A.7.4 Bounds on Asymptotic Beliefs and Asymptotic Cutoffs

For each t, let $\tilde{G}_t$ be the (random) posterior belief induced by the (random) posterior density $\tilde{g}_t$ after updating prior g using t rounds of histories.

Lemma A.11. For $c^l \geq C(\mu_1^\bullet, \underline{\mu}_2; \gamma)$, if almost surely $\liminf_{t\to\infty} \tilde{C}_t \geq c^l$, then almost surely

$$\lim_{t\to\infty} \tilde{G}_t(\ [\underline{\mu}_2, \mu_2^*(c^l))\ ) = 0.$$

Also, for $c^h \leq C(\mu_1^\bullet, \bar{\mu}_2; \gamma)$, if almost surely $\limsup_{t\to\infty} \tilde{C}_t \leq c^h$, then almost surely

$$\lim_{t\to\infty} \tilde{G}_t(\ (\mu_2^*(c^h), \bar{\mu}_2]) = 0.$$

Proof. I first show that for all $\epsilon > 0$, there exists $\delta > 0$ such that almost surely,

$$\liminf_{t\to\infty} \inf_{\mu_2 \in [\underline{\mu}_2, \mu_2^*(c^l)-\epsilon]} \frac{\ell_t^{'}(\mu_2)}{t} \geq \delta.$$

From Lemma A.6, we may rewrite LHS as

$$\liminf_{t\to\infty} \inf_{\mu_2 \in [\underline{\mu}_2, \mu_2^*(c^l)-\epsilon]} \left[\frac{1}{t}\frac{g^{'}(\mu_2)}{g(\mu_2)} + \frac{y_t(\mu_2)}{t} + \frac{z_t(\mu_2)}{t}\right],$$

which is no smaller than taking the inf separately across the three terms in the bracket,

$$\liminf_{t\to\infty} \inf_{\mu_2 \in [\underline{\mu}_2, \mu_2^*(c^l)-\epsilon]} \frac{1}{t}\frac{g^{'}(\mu_2)}{g(\mu_2)} + \liminf_{t\to\infty} \inf_{\mu_2 \in [\underline{\mu}_2, \mu_2^*(c^l)-\epsilon]} \frac{y_t(\mu_2)}{t} + \liminf_{t\to\infty} \inf_{\mu_2 \in [\underline{\mu}_2, \mu_2^*(c^l)-\epsilon]} \frac{z_t(\mu_2)}{t}.$$

Since $g^{'}$ is continuous and g is strictly positive (and continuous) on $[\underline{\mu}_2, \bar{\mu}_2]$ by the hypotheses of Theorem 1$'$, $g^{'}/g$ is bounded on $[\underline{\mu}_2, \bar{\mu}_2]$, so we in fact have $\lim_{t\to\infty} \inf_{\mu_2 \in [\underline{\mu}_2, \mu_2^*(c^l)-\epsilon]} \frac{1}{t}\frac{g^{'}(\mu_2)}{g(\mu_2)} =$

0. To deal with the second term,

$$\liminf_{t\to\infty} \inf_{\mu_2\in[\underline{\mu}_2,\mu_2^*(c^l)-\epsilon]} \frac{y_t(\mu_2)}{t} \geq \liminf_{t\to\infty} \inf_{\mu_2\in[\underline{\mu}_2\bar{\mu}_2]} \frac{y_t(\mu_2)}{t} = -\liminf_{t\to\infty} \sup_{\mu_2\in[\underline{\mu}_2\bar{\mu}_2]} \frac{-y_t(\mu_2)}{t}.$$

Lemma A.10 gives $\lim_{t\to\infty}\sup_{\mu_2\in[\underline{\mu}_2\bar{\mu}_2]} \frac{-y_t(\mu_2)}{t} = 0$ almost surely, so this second term is non-negative almost surely.

It suffices then to find $\delta > 0$ and show $\liminf_{t\to\infty}\inf_{\mu_2\in[\underline{\mu}_2,\mu_2^*(c^l)-\epsilon]} \frac{z_t(\mu_2)}{t} \geq \delta$ almost surely. Since z_t is the sum of $\bar{\varphi}_s$ terms that are decreasing functions of μ_2, the inner inf is always achieved at $\mu_2 = \mu_2^*(c^l) - \epsilon$. So we get

$$\begin{aligned}\liminf_{t\to\infty} \inf_{\mu_2\in[\underline{\mu}_2,\mu_2^*(c^l)-\epsilon]} \frac{z_t(\mu_2)}{t} &= \liminf_{t\to\infty} \frac{z_t(\mu_2^*(c^l)-\epsilon)}{t} \\ &= \liminf_{t\to\infty} \frac{1}{t}\left[\sum_{s=1}^{t}\bar{\varphi}_s(\mu_2^*(c^l)-\epsilon)\right].\end{aligned}$$

The definition of $\mu_2^*(c^l)$ is such that , $\mu_2^\bullet - \mu_2^*(c^l) - \gamma(\mu_1^\bullet - \mathbb{E}[X_1|X_1\leq c^l]) = 0$. So for any $\tilde{c} \geq c^l$, since $\gamma > 0$,

$$\begin{aligned}\mu_2^\bullet - \mu_2^*(c^l) - \gamma(\mu_1^\bullet - \mathbb{E}[X_1|X_1\leq \tilde{c}]) &\geq 0 \\ \mu_2^\bullet - (\mu_2^*(c^l)-\epsilon) - \gamma(\mu_1^\bullet - \mathbb{E}[X_1|X_1\leq \tilde{c}]) &\geq \epsilon.\end{aligned}$$

So at any $\tilde{c} \geq c^l$,

$$\sigma^{-2}\mathbb{P}[X_1\leq\tilde{c}]\cdot(\mu_2^\bullet - (\mu_2^*(c^l)-\epsilon) - \gamma(\mu_1^\bullet - \mathbb{E}[X_1|X_1\leq\tilde{c}])) \geq \sigma^{-2}\mathbb{P}[X_1\leq c^l]\cdot\epsilon$$

Along any ω where $\liminf_{t\to\infty}\tilde{C}_t \geq c^l$, we therefore have $\liminf_{t\to\infty}\bar{\varphi}_s(\mu_2^*(c^l)-\epsilon) \geq \sigma^{-2}\mathbb{P}[X_1\leq c^l]\cdot\epsilon$. Put $\delta = \sigma^{-2}\mathbb{P}[X_1\leq c^l]\cdot\epsilon$. This shows almost surely,

$$\liminf_{t\to\infty}\frac{1}{t}\left[\sum_{s=1}^{t}\bar{\varphi}_s(\mu_2^*(c^l)-\epsilon)\right] \geq \delta.$$

From here, it is a standard exercise to establish that $\lim_{t\to\infty}\tilde{G}_t(\ [\underline{\mu}_2,\mu_2^*(c^l)-\epsilon)\ ) = 0$ almost surely. Since the choice of $\epsilon > 0$ is arbitrary, this establishes the first part of the lemma.

The proof of the second part of the statement is exactly symmetric. To sketch the argument, we need to show that for all $\epsilon > 0$, there exists $\delta > 0$ such that almost surely,

$$\limsup_{t\to\infty} \sup_{\mu_2 \in [\mu_2^*(c^h)+\epsilon, \bar{\mu}_2]} \frac{\ell_t'(\mu_2)}{t} \leq -\delta.$$

This essentially reduces to analyzing

$$\limsup_{t\to\infty} \frac{1}{t}\left[\sum_{s=1}^{t} \bar{\varphi}_s(\mu_2^*(c^h)+\epsilon)\right].$$

For any $\tilde{c} \leq c^h$, since $\gamma > 0$,

$$\mu_2^\bullet - \mu_2^*(c^h) - \gamma(\mu_1^\bullet - \mathbb{E}[X_1 | X_1 \leq \tilde{c}]) \leq 0$$
$$\mu_2^\bullet - (\mu_2^*(c^h)+\epsilon) - \gamma(\mu_1^\bullet - \mathbb{E}[X_1 | X_1 \leq \tilde{c}]) \leq -\epsilon.$$

For every t and along every ω, $\tilde{C}_t(\omega) \geq C(\mu_1^\bullet, \underline{\mu}_2; \gamma)$, as cutoffs below this value cannot be myopically optimal given any belief about second-period fundamental supported on $[\underline{\mu}_2, \bar{\mu}_2]$. So along any ω such that $\limsup_{t\to\infty} \tilde{C}_t \leq c^h$, we have $\limsup_{t\to\infty} \bar{\varphi}_s(\mu_2^*(c^h)+\epsilon) \leq \sigma^{-2}\mathbb{P}[X_1 \leq C(\mu_1^\bullet, \underline{\mu}_2; \gamma)] \cdot (-\epsilon)$. Setting $\delta := \sigma^{-2}\mathbb{P}[X_1 \leq C(\mu_1^\bullet, \underline{\mu}_2; \gamma)] \cdot (\epsilon)$, we get $\limsup_{t\to\infty} \frac{1}{t}\left[\sum_{s=1}^{t} \bar{\varphi}_s(\mu_2^*(c^h)+\epsilon)\right] \leq -\delta$ almost surely. $\square$

Lemma A.12. For $\underline{\mu}_2 \leq \mu_2^l < \mu_2^h \leq \bar{\mu}_2$, if $\lim_{t\to\infty} \tilde{G}_t([\mu_2^l, \mu_2^h]) = 1$ almost surely, then $\liminf_{t\to\infty} \tilde{C}_t \geq C(\mu_1^\bullet, \mu_2^l; \gamma)$ and $\limsup_{t\to\infty} \tilde{C}_t \leq C(\mu_1^\bullet, \mu_2^h; \gamma)$ almost surely.

Proof. I show $\liminf_{t\to\infty} \tilde{C}_t \geq C(\mu_1^\bullet, \mu_2^l; \gamma)$ almost surely . The argument establishing $\limsup_{t\to\infty} \tilde{C}_t \leq C(\mu_1^\bullet, \mu_2^h; \gamma)$ is symmetric.

Let $c^l = C(\mu_1^\bullet, \mu_2^l; \gamma)$, $\underline{c} = C(\mu_1^\bullet, \underline{\mu}_2; \gamma)$, $\bar{c} = C(\mu_1^\bullet, \bar{\mu}_2; \gamma)$. Fix some $\epsilon > 0$. Since $c \mapsto U(c; \mu_1^\bullet, \mu_2)$ is single peaked for every μ_2, and since $c^l \leq C(\mu_1^\bullet, \mu_2; \gamma)$ for all $\mu_2 \in [\mu_2^l, \mu_2^h]$, we get $U(c^l; \mu_1^\bullet, \mu_2) - U(c^l - \epsilon; \mu_1^\bullet, \mu_2) > 0$ for every $\mu_2 \in [\mu_2^l, \mu_2^h]$. As

$$\mu_2 \mapsto \left(U(c^l; \mu_1^\bullet, \mu_2) - U(c^l - \epsilon; \mu_1^\bullet, \mu_2)\right)$$

is continuous, there exists some $\kappa^* > 0$ so that

$$U(c^l; \mu_1^\bullet, \mu_2) - U(c^l - \epsilon; \mu_1^\bullet, \mu_2) > \kappa^*$$

for all $\mu_2 \in [\mu_2^l, \mu_2^h]$. In particular, if $\nu \in \Delta([\mu_2^l, \mu_2^h])$ is a belief over second-period fundamental supported on $[\mu_2^l, \mu_2^h]$, then

$$\int U(c^l; \mu_1^\bullet, \mu_2) - U(c^l - \epsilon; \mu_1^\bullet, \mu_2) d\nu(\mu_2) > \kappa^*.$$

Now , let $\bar{\kappa} := \sup_{c\in[\underline{c},\bar{c}]} \sup_{\mu_2\in[\underline{\mu}_2,\bar{\mu}_2]} U(c;\mu_1^\bullet,\mu_2)$, $\underline{\kappa} := \inf_{c\in[\underline{c},\bar{c}]} \inf_{\mu_2\in[\underline{\mu}_2,\bar{\mu}_2]} U(c;\mu_1^\bullet,\mu_2)$. Find $p \in (0,1)$ so that $p\kappa^* - (1-p)(\bar{\kappa}-\underline{\kappa}) = 0$. At any belief $\hat{\nu} \in \Delta([\underline{\mu}_2,\bar{\mu}_2])$ that assigns more than probability p to the subinterval $[\mu_2^l,\mu_2^h]$, the optimal cutoff is larger than $c^l - \epsilon$. To see this, take any $\hat{c} \le c^l - \epsilon$ and I will show $\hat{c}$ is suboptimal. If $\hat{c} < \underline{c}$, then it is suboptimal after any belief on $[\underline{\mu}_2,\bar{\mu}_2]$. If $\underline{c} \le \hat{c} \le c^l - \epsilon$, I show that

$$\int U(c^l;\mu_1^\bullet,\mu_2) - U(\hat{c};\mu_1^\bullet,\mu_2)d\hat{\nu}(\mu_2) > 0.$$

To see this, we may decompose $\hat{\nu}$ as the mixture of a probability measure ν on $[\mu_2^l,\mu_2^h]$ and another probability measure ν^c on $[\underline{\mu}_2,\bar{\mu}_2]\backslash[\mu_2^l,\mu_2^h]$. Let $\hat{p} > p$ be the probability that ν assigns to $[\mu_2^l,\mu_2^h]$. The above integral is equal to:

$$\hat{p}\int_{\mu_2\in[\mu_2^l,\mu_2^h]} U(c^l;\mu_1^\bullet,\mu_2) - U(\hat{c};\mu_1^\bullet,\mu_2)d\nu(\mu_2) + (1-\hat{p})\int_{\mu_2\in[\underline{\mu}_2,\bar{\mu}_2]\backslash[\mu_2^l,\mu_2^h].} U(c^l;\mu_1^\bullet,\mu_2) - U(\hat{c};\mu_1^\bullet,\mu_2)d\nu^c(\mu_2)$$

Since c^l is to the left of the optimal cutoff for all $\mu_2 \in [\mu_2^l,\mu_2^h]$ and $\hat{c} \le c^l - \epsilon$, then $U(\hat{c};\mu_1^\bullet,\mu_2) \le U(c^l-\epsilon;\mu_1^\bullet,\mu_2)$ for all $\mu_2 \in [\mu_2^l,\mu_2^h]$. The first summand is no less than

$$\hat{p}\int_{\mu_2\in[\mu_2^l,\mu_2^h]} U(c^l;\mu_1^\bullet,\mu_2) - U(c^l-\epsilon;\mu_1^\bullet,\mu_2)d\nu(\mu_2) \ge \hat{p}\kappa^*.$$

Also, the integrand in the second summand is no smaller than $-(\bar{\kappa}-\underline{\kappa})$, therefore $\int U(c^l;\mu_1^\bullet,\mu_2) - U(\hat{c};\mu_1^\bullet,\mu_2)d\hat{\nu}(\mu_2) \ge \hat{p}\kappa^* - (1-\hat{p})(\bar{\kappa}-\underline{\kappa})$. Since $\hat{p} > p$, we get $\hat{p}\kappa^* - (1-\hat{p})(\bar{\kappa}-\underline{\kappa}) > 0$ as desired.

Along any sample path ω where $\lim_{t\to\infty} \tilde{G}_t([\mu_2^l,\mu_2^h])(\omega) = 1$, eventually $\tilde{G}_t([\mu_2^l,\mu_2^h])(\omega) > p$ for all large enough t, meaning $\liminf_{t\to\infty} \tilde{C}_t(\omega) \ge c^l - \epsilon$. This shows $\liminf_{t\to\infty} \tilde{C}_t \ge C(\mu_1^\bullet,\mu_2^l;\gamma) - \epsilon$ almost surely. Since the choice of $\epsilon > 0$ was arbitrary, we in fact conclude $\liminf_{t\to\infty} \tilde{C}_t \ge C(\mu_1^\bullet,\mu_2^l;\gamma)$ almost surely. □

A.7.5 The Contraction Map

I now combine the results established so far to prove Theorem 1′.

Proof. Let $\mu_{2,[1]}^l := \underline{\mu}_2$, $\mu_{2,[1]}^h := \bar{\mu}_2$. For $k = 2, 3, ...$, iteratively define $\mu_{2,[k]}^l := \mathcal{I}(\mu_{2,[k-1]}^l;\gamma)$ and $\mu_{2,[k]}^h := \mathcal{I}(\mu_{2,[k-1]}^h;\gamma)$.

From Lemma A.12, if $\lim_{t\to\infty} \tilde{G}_t([\mu_{2,[k]}^l,\mu_{2,[k]}^h]) = 1$ almost surely, then $\liminf_{t\to\infty} \tilde{C}_t \ge C(\mu_1^\bullet,\mu_{2,[k]}^l;\gamma)$ and $\limsup_{t\to\infty} \tilde{C}_t \le C(\mu_1^\bullet,\mu_{2,[k]}^h;\gamma)$ almost surely. But using these conclusions in Lemma A.11, we further deduce that $\lim_{t\to\infty} \tilde{G}_t([\mu_2^*(C(\mu_1^\bullet,\mu_{2,[k]}^l;\gamma)),\mu_2^*(C(\mu_1^\bullet,\mu_{2,[k]}^h;\gamma))]) = 1$ almost surely, that is to say $\lim_{t\to\infty} \tilde{G}_t([\mu_{2,[k+1]}^l,\mu_{2,[k+1]}^h]) = 1$ almost surely.

Under Assumptions 1 and 2, $\mu_2 \mapsto \mathcal{I}(\mu_2;\gamma)$ is a contraction mapping. Since $\underline{\mu}_2 < \mu_2^\infty$ and $\bar{\mu}_2 > \mu_2^\infty$, $(\mu_{2,[k]}^l)_{k\geq 1}$ is a sequence whose limit is μ_2^∞, and $(\mu_{2,[k]}^h)_{k\geq 1}$ is a sequence whose limit is μ_2^∞. Thus, agent's posterior converges in L^1 to μ_2^∞ almost surely (since the support of the prior is bounded).

In addition, $\mu_2 \mapsto C(\mu_1^\bullet, \mu_2; \gamma)$ is continuous, so the sequences of bounds on asymptotic cutoffs also converge, $\lim_{k\to\infty} C(\mu_1^\bullet, \mu_{2,[k]}^l; \gamma) = c^\infty$ and $\lim_{k\to\infty} C(\mu_1^\bullet, \mu_{2,[k]}^h; \gamma) = c^\infty$. This means $\lim_{t\to\infty} \tilde{C}_t = c^\infty$ almost surely. □

A.8 Proof of Lemma 2

Proof. By the same algebraic manipulations as in the proof of Proposition 2, we may rewrite the objective in Equation (1) as:

$$\frac{(\mu_1 - \mu_1^\bullet)^2}{2\sigma^2} + \sum_{\tau=0}^{t-1} \left\{ \int_{-\infty}^{c_\tau} \phi(x_1; \mu_1^\bullet, \sigma^2) \cdot \left[\frac{\sigma^2 + (\mu_2 - \gamma(x_1 - \mu_1) - \mu_2^\bullet)^2}{2\sigma^2} - \frac{1}{2} \right] dx_1 \right\}.$$

Dropping terms not dependent on μ_1, μ_2 and multiplying through by σ^2, we get the simplified objective

$$\xi(\mu_1, \mu_2) := \frac{(\mu_1 - \mu_1^\bullet)^2}{2} + \sum_{\tau=0}^{t-1} \left\{ \int_{-\infty}^{c_\tau} \phi(x_1; \mu_1^\bullet, \sigma^2) \cdot \left[\frac{(\mu_2 - \gamma(x_1 - \mu_1) - \mu_2^\bullet)^2}{2\sigma^2} \right] dx_1 \right\}$$

The same argument as in the proof of Proposition 2 gives $\mu_1 = \mu_1^\bullet$ as the only value satisfying the first-order conditions, and following this the minimizing μ_2 must satisfy $\frac{\partial \xi}{\partial \mu_2}(\mu_1^\bullet, \mu_2) = 0$. We now compute:

$$\frac{\partial \xi}{\partial \mu_2}(\mu_1^\bullet, \mu_2) = \sum_{\tau=0}^{t} \mathbb{P}[X_1 \leq c_\tau] \cdot \left(\mu_2 - \mu_2^\bullet - \gamma \left(\mathbb{E}\left[X_1 | X_1 \leq c_\tau\right] - \mu_1^\bullet\right)\right).$$

Since the derivative $\frac{\partial \xi}{\partial \mu_2}$ is a linear function of μ_2, when $\frac{\partial \xi}{\partial \mu_2}(\mu_1^\bullet, \mu_2^*) = 0$ we can rearrange to find

$$\begin{aligned} \mu_2^* &= \frac{1}{\sum_{\tau=0}^{t-1} \mathbb{P}[X_1 \leq c_\tau]} \cdot \sum_{\tau=0}^{t-1} \mathbb{P}[X_1 \leq c_\tau] \left\{\mu_2^\bullet - \gamma\left(\mu_1^\bullet - \mathbb{E}\left[X_1 | X_1 \leq c_\tau\right]\right)\right\} \\ &= \frac{1}{\sum_{\tau=0}^{t-1} \mathbb{P}[X_1 \leq c_\tau]} \sum_{\tau=0}^{t-1} \mathbb{P}[X_1 \leq c_\tau] \cdot \mu_2^*(c_\tau). \end{aligned}$$

This shows $\mu_1^*(c_0, ..., c_{t-1}) = \mu_1^\bullet$ and

$$\mu_2^*(c_0, ..., c_{t-1}) = \frac{1}{\sum_{\tau=0}^{t-1} \mathbb{P}[X_1 \leq c_\tau]} \sum_{\tau=0}^{t-1} \mathbb{P}[X_1 \leq c_\tau] \cdot \mu_2^*(c_\tau).$$

$\square$

A.9 Proof of Proposition 6

Proof. Suppose $\mu^A_{2,[2]} \geq \mu^A_{2,[1]}$. Under Assumption 1, Lemma 1 applies, so C is strictly increasing in its second argument. This shows $c^A_{[2]} = C(\mu^\bullet_1, \mu^A_{2,[2]}; \gamma) \geq C(\mu^\bullet_1, \mu^A_{2,[1]}; \gamma) = c^A_{[1]}$. But by Proposition 2, $\mu^*_2(c)$ increases in c, so $\mu^A_{2,[3]} = \mu^*_2(c^A_{[2]}) \geq \mu^*_2(c^A_{[1]}) = \mu^A_{2,[2]}$. Continuing this argument shows that $(\mu^A_{2,[t]})_{t\geq 1}$ is a monotonically increasing sequence. Since C is strictly increasing in its second argument, $(c^A_{[t]})_{t\geq 1}$ must also form a monotonically increasing sequence.

Conversely if $\mu^A_{2,[2]} < \mu^A_{2,[1]}$, then the analogous arguments show that $(\mu^A_{2,[t]})_{t\geq 1}$ and $(c^A_{[t]})_{t\geq 1}$ are monotonically decreasing sequences.

When Assumption 2 also holds, $\mathcal{I}(\cdot;\gamma)$ is a contraction mapping by Proposition 3. Since $(\mu^A_{2,[t]})_{t\geq 1}$ are the iterates of $\mathcal{I}(\cdot;\gamma)$, this implies $\lim_{t\to\infty} \mu^A_{2,[t]} = \mu^\infty_2$. Also, since $\mu_2 \mapsto C(\mu^\bullet_1, \mu_2; \gamma)$ is a continuous function by Lemma 1, we may exchange the limit:

$$\lim_{t\to\infty} c^A_{[t]} = \lim_{t\to\infty} C(\mu^\bullet_1, \mu^A_{2,[t]}; \gamma) = C(\mu^\bullet_1, \lim_{t\to\infty} \mu^A_{2,[t]}; \gamma) = C(\mu^\bullet_1, \mu^\infty_2; \gamma) = c^\infty.$$

So the monotonic sequence $(c^A_{[t]})_{t\geq 1}$ converges to c^∞. $\square$

A.10 Proof of Theorem 2

Proof. For the first step of the proof, suppose Assumption 1 holds.

Step 1: If $c_{[1]} > c_{[0]}$, then $(\mu_{2,[t]})_{t\geq 1}$ and $(c_{[t]})_{t\geq 0}$ are two increasing sequence, whereas $c_{[1]} \leq c_{[0]}$ implies $(\mu_{2,[t]})_{t\geq 1}$ and $(c_{[t]})_{t\geq 0}$ are two decreasing sequences.

Suppose $c_{[1]} > c_{[0]}$. Note that by Lemma 2, $\mu_{2,[1]} = \mu^*_2(c_{[0]})$, whereas $\mu_{2,[2]}$ is a weighted average between $\mu^*_2(c_{[0]})$ and $\mu^*_2(c_{[1]})$ where the latter is larger because $c_{[1]} > c_{[0]}$ and $\mu^*_2(c)$ is strictly increasing. This shows we have $\mu_{2,[2]} > \mu_{2,[1]}$ and hence $c_{[2]} > c_{[1]}$ as the cutoff is strictly increasing in its second argument by Lemma 1. Now assume the partial sequences $(c_{[\tau]})^T_{\tau=0}$ and $(\mu_{2,[\tau]})^T_{\tau=1}$ are both increasing. We show that $\mu_{2,[T+1]} > \mu_{2,[T]}$, which would also imply $c_{[T+1]} > c_{[T]}$. By comparing expressions for $\mu_{2,[T+1]}$ and $\mu_{2,[t]}$ given by Lemma 2,

$$\mu_{2,[T+1]} = \delta \cdot \mu^*_2(c_{[T]}) + (1-\delta) \cdot \mu_{2,[t]}$$

where $\delta = \frac{\mathbb{P}[X_1 \leq c_{[T]}]}{\sum_{\tau=0}^{T} \mathbb{P}[X_1 \leq c_{[\tau]}]} > 0$ and $\mu_{2,[t]}$ is itself a weighted average of the collection $\{\mu^*_2(c_{[\tau]})\}_{0\leq\tau\leq T-1}$ by Lemma 2. Now by the first part of the inductive hypothesis, $(c_{[\tau]})^T_{\tau=0}$ is strictly increasing, meaning $\mu^*_2(c_{[T]}) > \mu^*_2(c_{[\tau]})$ for any $\tau < T$, which are the components making up $\mu_{2,[t]}$. Since the weight δ on $\mu^*_2(c_{[T]})$ in the expression of $\mu_{2,[T+1]}$ is strictly positive, this shows

$\mu_{2,[T+1]} > \mu_{2,[t]}$. So by induction, we have shown Step 1. (The other case of $c_{[1]} < c_{[0]}$ is symmetric.)

For the rest of this proof, suppose Assumption 2 also holds.

Step 2: $(\mu_{2,[t]})_{t\geq 1}$ is bounded and converges.

In the case that $c_{[1]} \geq c_{[0]}$ (so $\mu_{2,[2]} \geq \mu_{2,[1]}$), **Step 1** implies that $(\mu_{2,[t]})_{t\geq 1}$ forms an increasing sequence. Since $\mu_2^*(\cdot)$ is bounded above by $\mu_2^\bullet$ by Proposition 2 and $\mu_{2,[t]}$ for any $t \geq 1$ is a convex combinations of such terms, we also have $\mu_{2,[t]} \leq \mu_2^\bullet$ for every t. So in this case the sequence $(\mu_{2,[t]})_{t\geq 1}$ is bounded between $\mu_{2,[1]}$ and $\mu_2^\bullet$.

In the case that $c_{[1]} \leq c_{[0]}$ (so $\mu_{2,[2]} \leq \mu_{2,[1]}$), we notice that $c_{[0]} = c_{[0]}^A$, $c_{[1]} = c_{[1]}^A$, so by Proposition 6 the auxiliary environment has the dynamics $\mu_{2,[t]}^A \searrow \mu_2^\infty$, $c_{[t]}^A \searrow c^\infty$, where (μ_2^∞, c^∞) are associated with the unique steady state. So we have $\mu_{2,[1]} = \mu_{2,[1]}^A$ while $\mu_{2,[2]} \geq \mu_{2,[2]}^A$ since $\mu_{2,[2]}$ is a convex combination between $\mu_2^*(c_{[0]})$ and $\mu_2^*(c_{[1]}) = \mu_{2,[2]}^A$, with the latter being lower. This means $c_{[2]} \geq c_{[2]}^A$. In the third generation,

$$\mu_2^*(c_{[0]}, c_{[1]}, c_{[2]}) \geq \mu_2^*(c_{[2]}) \geq \mu_2^*(c_{[2]}^A).$$

The first inequality follows because $\mu_2^*(c_{[0]}, c_{[1]}, c_{[2]})$ is a weighted average between $\mu_2^*(c_{[0]})$, $\mu_2^*(c_{[1]})$, and $\mu_2^*(c_{[2]})$, with the last one being the lowest since $c_{[t]}$ decreases in t for $t = 0, 1, 2$. This shows $\mu_{2,[3]} \geq \mu_{2,[3]}^A$ and $c_{[3]} \geq c_{[3]}^A$. Iterating this argument shows that $\mu_{2,[t]} \geq \mu_{2,[t]}^A$ for every t in this case. Seeing as $(\mu_{2,[t]})_{t\geq 1}$ forms a decreasing sequence by **Step 1**, it is bounded between μ_2^∞ and $\mu_{2,[1]}$.

Since $(\mu_{2,[t]})_{t\geq 1}$ is a bounded, monotonic sequence, it must converge. I denote this limit as $\mu_{2,[t]} \to \tilde{\mu}_2$. Also, since $\mu_2 \mapsto C(\mu_1^\bullet, \mu_2; \gamma)$ is continuous from Proposition 1, the sequence $c_{[t]}$ must also converge. I denote this limit by $c_{[t]} \to \tilde{c}$.

Step 3: $\tilde{\mu}_2$ is a fixed point of $\mathcal{I}(\cdot;\gamma)$, so in particular $\tilde{\mu}_2 = \mu_2^\infty$ and $\tilde{c} = c^\infty$ since $\mathcal{I}(\cdot;\gamma)$ has a unique fixed point.

Proposition 3 shows that under Assumption 2, $\mathcal{I}(\cdot;\gamma)$ is a contraction mapping and hence must be continuous. Now let any $\epsilon > 0$ be given. I show there exists $\bar{t}$ so that $|\mathcal{I}(\mu_{2,[t]};\gamma) - \mu_{2,[t]}| < \epsilon$ for all $t > \bar{t}$. As this holds for all $\epsilon > 0$, continuity of $\mathcal{I}(\cdot;\gamma)$ then implies $\mathcal{I}(\tilde{\mu}_2;\gamma) - \tilde{\mu}_2 = 0$, that is $\tilde{\mu}_2$ is a fixed point of $\mathcal{I}(\cdot;\gamma)$.

We may write by Lemma 2,

$$\mu_{2,[t]} = \frac{1}{\sum_{\tau=0}^{t-1} \mathbb{P}[X_1 \leq c_{[\tau]}]} \sum_{\tau=0}^{t-1} \mathbb{P}[X_1 \leq c_{[\tau]}] \cdot \mu_2^*(c_{[\tau]}).$$

The probabilities $\mathbb{P}[X_1 \leq c_{[\tau]}]$ are bounded below since the beliefs $(\mu_{2,[t]})_{t\geq 1}$ are bounded by **Step 2**. Also, since $\mu_2^*(\cdot)$ is continuous, there exists $\bar{t}_1$ so that $|\mu_2^*(c_{[t]}) - \mu_2^*(\tilde{c})| < \epsilon/2$ for all $t > \bar{t}_1$, that is to say $|\mathcal{I}(\mu_{2,[t]};\gamma) - \mu_2^*(\tilde{c})| < \epsilon/2$. When $T \to \infty$, the sum of weights assigned to terms $\mu_2^*(c_{[t]})$ with $t \geq \bar{t}_1$ in the expression for $\mu_{2,[T]}$ grows to 1, which means $\limsup_{T\to\infty} |\mu_{2,[T]} - \mu_2^*(\tilde{c})| < \epsilon/2$. Combining these facts give $\limsup_{t\to\infty} |\mathcal{I}(\mu_{2,[t]};\gamma) - \mu_{2,[t]}| < \epsilon$ as desired. This establishes that $\tilde{\mu}_2$ is a fixed point of $\mathcal{I}(\cdot;\gamma)$. Since $\mathcal{I}(\cdot;\gamma)$ has a unique fixed point, $\tilde{\mu}_2 = \mu_2^\infty$. By continuity of C in its second argument from Proposition 1,

$$\tilde{c} = \lim_{t\to\infty} C(\mu_1^\bullet, \mu_{2,[t]};\gamma) = C(\mu_1^\bullet, \lim_{t\to\infty} \mu_{2,[t]};\gamma) = C(\mu_1^\bullet, \tilde{\mu}_2;\gamma) = C(\mu_1^\bullet, \mu_2^\infty;\gamma) = c^\infty.$$

□

A.11 Proof of Corollary 1

Proof. Suppose $c_{[1]} \geq c_{[0]}$. Since $\mu_2^*(c)$ is increasing, we have $\mu_{2,[2]} = \mu_2^*(c_{[1]}, c_{[0]}) \geq \mu_2^*(c_{[0]}) = \mu_{2,[1]}$. So we get $c_{[2]} \geq c_{[1]}$. By Theorem 2, we deduce $(c_{[t]})_{t\geq 0}$ is an increasing sequence, so in particular $c^\infty \geq c^\bullet$. But again by 2, c^∞ is the same as the steady-state cutoff in Theorem 1. This is a contradiction because Theorem 1 implies $c^\infty < c^\bullet$.

This shows $c_{[1]} < c_{[0]}$ and similar arguments show $(c_{[t]})_{t\geq 0}$ is a strictly decreasing sequence. Since $c^\bullet$ is the objectively optimal cutoff threshold under the true model $\Psi^\bullet$, and since expected payoff under the true model is a single-peaked function in acceptance threshold, this shows expected payoff is strictly decreasing across generations. □

A.12 Proof of Lemma 3

Proof. Indifference condition $c^{''} = C_{u_1,u_2^{''}}(\mu_1, \mu_2;\gamma)$ implies that

$$u_1(c^{''}) = \mathbb{E}_{\tilde{X}_2 \sim \mathcal{N}(\mu_2 - \gamma(c^{''} - \mu_1), \sigma^2)}[u_2^{''}(c^{''}, \tilde{X}_2)].$$

Since $u_2^{'}(c^{''}, x_2) \geq u_2^{''}(c^{''}, x_2)$ for all $x_2 \in \mathbb{R}$, with strict inequality on a positive-measure set, this shows

$$u_1(c^{''}) < \mathbb{E}_{\tilde{X}_2 \sim \mathcal{N}(\mu_2 - \gamma(c^{''} - \mu_1), \sigma^2)}[u_2^{'}(c^{''}, \tilde{X}_2)].$$

Because $(u_1, u_2^{'})$ satisfy Assumptions 1, the best stopping strategy in the subjective model $\Psi(\mu_1, \mu_2;\gamma)$ has a cutoff form by Proposition 1. This shows $C_{u_1,u_2^{'}}(\mu_1, \mu_2;\gamma)$ is strictly above $c^{''}$. □

A.13 Proof of Proposition 7

Proof. Under Assumptions 1 and 2, each of $(u_1, u_2^{'})$ and $(u_1, u_2^{''})$ has a unique steady state, $(\mu_1^\bullet, \mu_{2,A}^\infty, c_A^\infty)$, $(\mu_1^\bullet, \mu_{2,B}^\infty, c_B^\infty)$ respectively. Let $\mathcal{I}_A, \mathcal{I}_B$ be the iteration maps corresponding to these two stage games, that is to say

$$\mathcal{I}_A(\mu_2) := \mu_2^\bullet - \gamma(\mu_1^\bullet - \mathbb{E}[X_1 \mid X_1 \leq C_{u_1,u_2'}(\mu_1^\bullet, \mu_2; \gamma)])$$

$$\mathcal{I}_B(\mu_2) := \mu_2^\bullet - \gamma(\mu_1^\bullet - \mathbb{E}[X_1 \mid X_1 \leq C_{u_1,u_2''}(\mu_1^\bullet, \mu_2; \gamma)]).$$

From Proposition 3, both $\mathcal{I}_A$ and $\mathcal{I}_B$ are contraction mappings. Consider their iterates with a starting value of 0. That is, put $\mu_{2,A}^{[0]} = 0$, $\mu_{2,B}^{[0]} = 0$ and let $\mu_{2,A}^{[t]} = \mathcal{I}_A(\mu_{2,A}^{[t-1]})$, $\mu_{2,B}^{[t]} = \mathcal{I}_B(\mu_{2,B}^{[t-1]})$ for $t \geq 1$. By property of contraction mappings and since the fixed points of the iteration maps are the steady state beliefs, $\mu_{2,A}^{[t]} \to \mu_{2,A}^{\infty}$ and $\mu_{2,B}^{[t]} \to \mu_{2,B}^{\infty}$.

By induction, $\mu_{2,B}^{[t]} \leq \mu_{2,A}^{[t]}$ for every $t \geq 0$. The base case of $t = 0$ is true by definition. If $\mu_{2,B}^{[T]} \leq \mu_{2,A}^{[T]}$, then

$$C_{u_1,u_2''}(\mu_1^\bullet, \mu_{2,B}^{[T]}; \gamma) \leq C_{u_1,u_2''}(\mu_1^\bullet, \mu_{2,A}^{[T]}; \gamma) < C_{u_1,u_2'}(\mu_1^\bullet, \mu_{2,A}^{[T]}; \gamma).$$

The first inequality comes from C being increasing in the second argument and the inductive hypothesis, while the second inequality is due to Lemma 3. Therefore, $\mathcal{I}_B(\mu_{2,B}^{[T]}) \leq \mathcal{I}_A(\mu_{2,A}^{[T]})$, so $\mu_{2,A}^{[T+1]} \leq \mu_{2,A}^{[T+1]}$.

Since weak inequalities are preserved by limits, we have $\mu_{2,A}^{\infty} \geq \mu_{2,B}^{\infty}$. It is impossible to have $\mu_{2,A}^{\infty} = \mu_{2,B}^{\infty}$, because this would lead to $c_A^{\infty} > c_B^{\infty}$ by Lemma 3, which in turn implies $\mu_{2,A}^{\infty} = \mu_2^*(c_A^{\infty}) > \mu_2^*(c_B^{\infty}) = \mu_{2,B}^{\infty}$. This inequality contradicts the equality $\mu_{2,A}^{\infty} = \mu_{2,B}^{\infty}$. Therefore, we in fact have $\mu_{2,A}^{\infty} > \mu_{2,B}^{\infty}$. The conclusion that $c_A^{\infty} > c_B^{\infty}$ follows from Lemma 3 and the fact that C is increases in its second argument. □

A.14 Proof of Proposition 8

Proof. Rewrite Equation (2) as

$$\int_{-\infty}^{\infty} \phi(x_1; \mu_1^\bullet, (\sigma^\bullet)^2) \cdot \ln\left(\frac{\phi(x_1; \mu_1^\bullet, (\sigma^\bullet)^2)}{\phi(x_1; \mu_1, \sigma_1^2)}\right) dx_1$$
$$+ \int_{-\infty}^{c} \phi(x_1; \mu_1^\bullet, (\sigma^\bullet)^2) \cdot \int_{-\infty}^{\infty} \phi(x_2; \mu_2^\bullet, (\sigma^\bullet)^2) \ln\left[\frac{\phi(x_2; \mu_2^\bullet, (\sigma^\bullet)^2)}{\phi(x_2; \mu_2 - \gamma(x_1 - \mu_1), \sigma_2^2)}\right] dx_2 dx_1.$$

The KL divergence between $\mathcal{N}(\mu_{\text{true}}, \sigma_{\text{true}}^2)$ and $\mathcal{N}(\mu_{\text{model}}, \sigma_{\text{model}}^2)$ is $\ln \frac{\sigma_{\text{model}}}{\sigma_{\text{true}}} + \frac{\sigma_{\text{true}}^2 + (\mu_{\text{true}} - \mu_{\text{model}})^2}{2\sigma_{\text{model}}^2} - \frac{1}{2}$, so we may simplify the first term and the inner integral of the second term.

$$\ln \frac{\sigma_1}{\sigma^\bullet} + \frac{(\mu_1 - \mu_1^\bullet)^2}{2\sigma_1^2} + \frac{(\sigma^\bullet)^2}{2\sigma_1^2} - \frac{1}{2}$$
$$+ \int_{-\infty}^{c} \phi(x_1; \mu_1^\bullet, \sigma^\bullet) \cdot \left[\ln \frac{\sigma_2}{\sigma^\bullet} + \frac{(\sigma^\bullet)^2 + (\mu_2 - \gamma(x_1 - \mu_1) - \mu_2^\bullet)^2}{2\sigma_2^2} - \frac{1}{2}\right] dx_1.$$

Dropping terms not dependent on any of the four variables gives a simplified version of the objective,

$$\begin{aligned}\xi(\mu_1,\mu_2,\sigma_1,\sigma_2) := \ln\frac{\sigma_1}{\sigma^\bullet} + \frac{(\mu_1-\mu_1^\bullet)^2}{2\sigma_1^2} + \frac{(\sigma^\bullet)^2}{2\sigma_1^2} \\ + \int_{-\infty}^{c}\phi(x_1;\mu_1^\bullet,(\sigma^\bullet)^2)\cdot\left[\ln\frac{\sigma_2}{\sigma^\bullet} + \frac{(\sigma^\bullet)^2 + (\mu_2-\gamma(x_1-\mu_1)-\mu_2^\bullet)^2}{2\sigma_2^2}\right]dx_1.\end{aligned}$$

Differentiating under the integral sign,

$$\frac{\partial\xi}{\partial\mu_2} = \int_{-\infty}^{c}\phi(x_1;\mu_1^\bullet,(\sigma^\bullet)^2)\cdot\left[\frac{(\mu_2-\gamma(x_1-\mu_1)-\mu_2^\bullet)}{\sigma_2^2}\right]dx_1$$

$$\begin{aligned}\frac{\partial\xi}{\partial\mu_1} &= \frac{(\mu_1-\mu_1^\bullet)}{\sigma_1^2} + \gamma\int_{-\infty}^{c}\phi(x_1;\mu_1^\bullet,(\sigma^\bullet)^2)\cdot\left[\frac{(\mu_2-\gamma(x_1-\mu_1)-\mu_2^\bullet)}{\sigma_2^2}\right]dx_1 \\ &= \frac{(\mu_1-\mu_1^\bullet)}{\sigma_1^2} + \gamma\frac{\partial\xi}{\partial\mu_2}.\end{aligned}$$

At FOC $(\mu_1^*,\mu_2^*,\sigma_1^*,\sigma_2^*)$, we have $\frac{\partial\xi}{\partial\mu_2}(\mu_1^*,\mu_2^*,\sigma_1^*,\sigma_2^*)=0$, hence $\mu_1^*=\mu_1^\bullet$. Similar arguments as before then establish $\mu_2^*=\mu_2^\bullet-\gamma\left(\mu_1^\bullet-\mathbb{E}\left[X_1\mid X_1\le c\right]\right)$, where expectation is taken with respect to the true distribution of X_1 (with the true variance $(\sigma^\bullet)^2$). Then,

$$\frac{\partial\xi}{\partial\sigma_1}(\mu_1^*,\mu_2^*,\sigma_1^*,\sigma_2^*) = \frac{1}{(\sigma_1^*)} - \frac{(\sigma^\bullet)^2}{(\sigma_1^*)^3} = 0,$$

this gives $\sigma_1^*=\sigma^\bullet$ (since $\sigma_1^*\ge 0$).

Finally, from the FOC for σ_2,

$$\int_{-\infty}^{c}\phi(x_1;\mu_1^\bullet,(\sigma^\bullet)^2)\cdot\left[\frac{1}{\sigma_2^*} - \frac{(\sigma^\bullet)^2+(\mu_2^*-\gamma(x_1-\mu_1^*)-\mu_2^\bullet)^2}{(\sigma_2^*)^3}\right]dx_1 = 0.$$

Substituting in values of μ_1^*,μ_2^* already solved for,

$$\begin{aligned}(\sigma_2^*)^2 &= (\sigma^\bullet)^2 + \mathbb{E}[(\mu_2^*-\gamma(X_1-\mu_1^\bullet)-\mu_2^\bullet)^2|X_1\le c] \\ &= (\sigma^\bullet)^2 + \mathbb{E}[(\mu_2^\bullet-\gamma\left(\mu_1^\bullet-\mathbb{E}\left[X_1\mid X_1\le c\right]\right)-\gamma(X_1-\mu_1^\bullet)-\mu_2^\bullet)^2|X_1\le c] \\ &= (\sigma^\bullet)^2 + \gamma^2\mathbb{E}\left[\left[(X_1-\mu_1^\bullet)-\left(\mathbb{E}\left[X_1\mid X_1\le c\right]-\mu_1^\bullet\right)\right]^2|X_1\le c\right] \\ &= (\sigma^\bullet)^2 + \gamma^2\mathrm{Var}[X_1-\mu_1^\bullet|X_1\le c] \\ &= (\sigma^\bullet)^2 + \gamma^2\mathrm{Var}[X_1|X_1\le c]\end{aligned}$$

as desired. □

A.15 Proof of Proposition 9

I start with a lemma that says, depending on the convexity of the decision problem, a stronger belief in fictitious variation either increases or decreases the subjectively optimal cutoff threshold.

Lemma A.13. *Suppose that under the subjective model $\Psi(\mu_1, \mu_2, \sigma_1^2, \sigma_2^2; \gamma)$, the agent is indifferent between stopping at c and continuing. Suppose $\hat{\sigma}_2^2 > \sigma_2^2$. Then: (i) if $x_2 \mapsto u_2(c, x_2)$ is convex with strict convexity for x_2 in a positive-measure set, then under the subjective model $\Psi(\mu_1, \mu_2, \sigma_1^2, \hat{\sigma}_2^2; \gamma)$ the agent strictly prefers continuing at c; (ii) if $x_2 \mapsto u_2(c, x_2)$ is concave with strict concavity for x_2 in a positive-measure set, then under the subjective model $\Psi(\mu_1, \mu_2, \sigma_1^2, \hat{\sigma}_2^2; \gamma)$ the agent strictly prefers stopping at c.*

Proof. Indifference at $x_1 = c$ under the model $\Psi(\mu_1, \mu_2, \sigma_1^2, \sigma_2^2; \gamma)$ implies that

$$u_1(c) = \mathbb{E}_{X_2 \sim \mathcal{N}(\mu_2 - \gamma(x_1 - \mu_1), \sigma_2^2)}[u_2(c, X_2)].$$

When hypothesis in (i) is satisfied,

$$\mathbb{E}_{X_2 \sim \mathcal{N}(\mu_2 - \gamma(x_1 - \mu_1), \sigma_2^2)}[u_2(c, X_2)] < \mathbb{E}_{X_2 \sim \mathcal{N}(\mu_2 - \gamma(x_1 - \mu_1), \hat{\sigma}_2^2)}[u_2(c, X_2)]$$

since $\hat{\sigma}_2^2 > \sigma_2^2$ implies that $\mathcal{N}(\mu_2 - \gamma(x_1 - \mu_1), \hat{\sigma}_2^2)$ is a strict mean-preserving spread of $\mathcal{N}(\mu_2 - \gamma(x_1 - \mu_1), \sigma_2^2)$. The RHS is the expected continuation payoff under model $\Psi(\mu_1, \mu_2, \sigma_1^2, \hat{\sigma}_2^2; \gamma)$, so the agent strictly prefers continuing when $X_1 = c$. The argument establishing (ii) is analogous. □

Now I give the proof of Proposition 9.

Proof. The result that $\mu_{1,[t]} = \mu_1^\bullet$, $(\sigma_{1,[t]})^2 = (\sigma^\bullet)^2$ for all t follows from Proposition 8.

Suppose $c_{[1]} \leq c_{[0]}$. From Proposition 8, $\mu_{2,[2]} \leq \mu_{2,[1]}$ and $(\sigma_{2,[2]})^2 \leq (\sigma_{2,[1]})^2$. Let $c'_{[2]}$ be the indifference threshold under the model $\Psi(\mu_1^\bullet, \mu_{2,[2]}, (\sigma^\bullet)^2, (\sigma_{2,[1]})^2)$. By Lemma 1, $c'_{[2]} \leq c_{[1]}$. Also, from Lemma A.13, $c_{[2]} \leq c'_{[2]}$ as generation 2 actually believes in the subjective model $\Psi(\mu_1^\bullet, \mu_{2,[2]}, (\sigma^\bullet)^2, (\sigma_{2,[2]})^2)$ where $(\sigma_{2,[2]})^2 \leq (\sigma_{2,[1]})^2$. This shows $c_{[2]} \leq c_{[1]}$. Continuing this argument shows that $(c_{[t]})_{t \geq 1}$ forms a monotonically decreasing sequence. Since the pseudo-true parameters μ_2^* and $(\sigma_2^*)^2$ are monotonic functions of the censoring threshold c, we have established the proposition in the case where $c_{[1]} \leq c_{[0]}$.

The argument for the case where $c_{[1]} \geq c_{[0]}$ is exactly analogous and therefore omitted. □

A.16 Proof of Proposition 10

Proof. In the first generation, both societies A and B observe large datasets of histories with distribution $\mathcal{H}^\bullet(c_{[0]})$. So, by Proposition 8, two societies make the same inferences about the fundamentals.

Suppose the optimal-stopping problem is convex. Then due to fictitious variation in generation 1 and the convexity of u_2, it follows from Lemma A.13 that $c_{[B,1]} > c_{[A,1]}$. In the second generation, $\mu_{2,[B,2]} > \mu_{2,[A,2]}$ because the pseudo-true second-period fundamental increases in the censoring cutoff. Together again with the existence of fictitious variation, we conclude $c_{[B,2]} > c_{[A,2]}$. Continuing this argument establishes the proposition for the case where the optimal-stopping problem is convex. The case of concave optimal-stopping problems is analogous. $\square$

A.17 Proof of Proposition 11

Proof. In the true model, $X_2|(X_1 = x_1) \sim \mathcal{N}(\mu_2^\bullet - \gamma^\bullet(x_1 - \mu_1^\bullet), \sigma^2)$, while the agents' subjective model $\Psi(\mu_1, \mu_2; \gamma)$ has $X_2|(X_1 = x_1) \sim \mathcal{N}(\mu_2 - \gamma(x_1 - \mu_1), \sigma^2)$. So, we can write

$$D_{KL}(\mathcal{H}(\Psi(\mu_1^\bullet, \mu_2^\bullet; \gamma^\bullet); c) \parallel \mathcal{H}(\Psi(\mu_1, \mu_2; \gamma); c))$$

as the following:

$$\int_c^\infty \phi(x_1; \mu_1^\bullet, \sigma^2) \cdot \ln\left(\frac{\phi(x_1; \mu_1^\bullet, \sigma^2)}{\phi(x_1; \mu_1, \sigma^2)}\right) dx_1$$
$$+ \int_{-\infty}^c \left\{ \int_{-\infty}^\infty \begin{matrix} \phi(x_1; \mu_1^\bullet, \sigma^2) \cdot \phi(x_2; \mu_2^\bullet - \gamma^\bullet(x_1 - \mu_1^\bullet), \sigma^2) \cdot \\ \ln\left[\frac{\phi(x_1; \mu_1^\bullet, \sigma^2) \cdot \phi(x_2; \mu_2^\bullet - \gamma^\bullet(x_1 - \mu_1^\bullet), \sigma^2)}{\phi(x_1; \mu_1, \sigma^2) \cdot \phi(x_2; \mu_2 - \gamma(x_1 - \mu_1), \sigma^2)}\right] \end{matrix} dx_2 \right\} dx_1.$$

Performing rearrangements similar to those in the proof of Proposition 2 and using the closed-form expression of KL divergence between two Gaussian distributions, the above can be rewritten as

$$\frac{(\mu_1 - \mu_1^\bullet)^2}{2\sigma^2} + \int_{-\infty}^c \phi(x_1; \mu_1^\bullet, \sigma^2) \cdot \frac{(\mu_2 - \gamma(x_1 - \mu_1) - \mu_2^\bullet + \gamma^\bullet(x_1 - \mu_1^\bullet))^2}{2\sigma^2} dx_1.$$

Multiplying through byσ^2 and dropping terms not depending on μ_1, μ_2, γ, we get a simplified objective with the same minimizers:

$$\xi(\mu_1, \mu_2, \gamma) = \frac{(\mu_1 - \mu_1^\bullet)^2}{2} + \int_{-\infty}^c \phi(x_1; \mu_1^\bullet, \sigma^2) \cdot \frac{1}{2} \cdot [\mu_2 - \gamma(x_1 - \mu_1) - \mu_2^\bullet + \gamma^\bullet(x_1 - \mu_1^\bullet)]^2 dx_1.$$

We have the partial derivatives by differentiating under the integral sign,

$$\frac{\partial \xi}{\partial \mu_2} = \int_{-\infty}^c \phi(x_1; \mu_1^\bullet, \sigma^2) \cdot [\mu_2 - \gamma(x_1 - \mu_1) - \mu_2^\bullet + \gamma^\bullet(x_1 - \mu_1^\bullet)] dx_1,$$

$$
\begin{aligned}
\frac{\partial \xi}{\partial \mu_1} &= (\mu_1 - \mu_1^\bullet) + \gamma \int_{-\infty}^{c} \phi(x_1; \mu_1^\bullet, \sigma^2) \cdot [\mu_2 - \gamma(x_1 - \mu_1) - \mu_2^\bullet + \gamma^\bullet (x_1 - \mu_1^\bullet)] dx_1 \\
&= (\mu_1 - \mu_1^\bullet) + \gamma \frac{\partial \xi}{\partial \mu_2},
\end{aligned}
$$

$$
\frac{\partial \xi}{\partial \gamma} = -\int_{-\infty}^{c} \phi(x_1; \mu_1^\bullet, \sigma^2) \cdot [x_1 - \mu_1] \cdot [\mu_2 - \gamma(x_1 - \mu_1) - \mu_2^\bullet + \gamma^\bullet (x_1 - \mu_1^\bullet)] dx_1.
$$

Suppose $(\mu_1^*, \mu_2^*, \gamma^*)$ is the minimum. By the first-order conditions for μ_1 and μ_2, we have:

$$
\frac{\partial \xi}{\partial \mu_1}(\mu_1^*, \mu_2^*, \gamma^*) = \frac{\partial \xi}{\partial \mu_2}(\mu_1^*, \mu_2^*, \gamma^*) = 0 \Rightarrow \mu_1^* = \mu_1^\bullet.
$$

Substituting this into the first-order condition for μ_2,

$$
\frac{\partial \xi}{\partial \mu_2}(\mu_1^\bullet, \mu_2^*, \gamma^*) = 0 \Rightarrow \mu_2^* = \mu_2^\bullet + (\gamma^\bullet - \gamma^*) \cdot (\mu_1^\bullet - \mathbb{E}[X_1 | X_1 \leq c]) \, .
$$

It remains to show $\gamma^* = \tilde{\gamma}$. We have

$$
\frac{\partial \xi}{\partial \gamma}(\mu_1^*, \mu_2^*, \gamma^*) = -\mathbb{P}[X_1 \leq c] \cdot \mathbb{E}[(X_1 - \mu_1^*) \cdot (\mu_2^* - \gamma^*(X_1 - \mu_1^*) - \mu_2^\bullet + \gamma^\bullet (X_1 - \mu_1^\bullet)) | X_1 \leq c].
$$

We rearrange the expectation term as:

$$
\begin{aligned}
&\mathbb{E}[(X_1 - \mu_1^*) \cdot (\mu_2^* - \gamma^*(X_1 - \mu_1^*) - \mu_2^\bullet + \gamma^\bullet (X_1 - \mu_1^\bullet)) | X_1 \leq c] \\
=& \mathbb{E}[(X_1 - \mu_1^*) | X_1 \leq c] \cdot \mathbb{E}[(\mu_2^* - \gamma^*(X_1 - \mu_1^*) - \mu_2^\bullet + \gamma^\bullet (X_1 - \mu_1^\bullet)) | X_1 \leq c] \\
&+ \mathrm{Cov}(X_1 - \mu_1^*, \mu_2^* - \gamma^*(X_1 - \mu_1^*) - \mu_2^\bullet + \gamma^\bullet (X_1 - \mu_1^\bullet) | X_1 \leq c].
\end{aligned}
$$

The first-order condition for μ_2 implies $\mathbb{E}[(\mu_2^* - \gamma^*(X_1 - \mu_1^*) - \mu_2^\bullet + \gamma^\bullet (X_1 - \mu_1^\bullet)) | X_1 \leq c] = 0$ at the optimum $(\mu_1^*, \mu_2^*, \gamma^*)$. Also, we may drop terms without X_1 in the conditional covariance operator, and we get:

$$
\frac{\partial \xi}{\partial \gamma}(\mu_1^*, \mu_2^*, \gamma^*) = \mathbb{P}[X_1 \leq c] \cdot (\gamma^* - \gamma^\bullet) \cdot \mathrm{Cov}(X_1, X_1 | X_1 \leq c).
$$

We have $\mathbb{P}[X_1 \leq c] > 0$ and $\mathrm{Cov}(X_1, X_1 | X_1 \leq c) > 0$, hence we conclude

$$
\frac{\partial \xi}{\partial \gamma}(\mu_1^*, \mu_2^*, \gamma^*) \begin{cases} > 0 & \text{for } \gamma^* > \gamma^\bullet \\ = 0 & \text{for } \gamma^* = \gamma^\bullet \\ < 0 & \text{for } \gamma^* < \gamma^\bullet \end{cases}.
$$

In case that $\underline{\gamma} > \gamma^\bullet$, at the optimum we must have $\frac{\partial \xi}{\partial \gamma}(\mu_1^*, \mu_2^*, \gamma^*) > 0$. By Karush-Kuhn-Tucker condition, this means the minimizer is $\gamma^* = \underline{\gamma}$. Conversely, when $\bar{\gamma} < \gamma^\bullet$, at the optimum we must have $\frac{\partial \xi}{\partial \gamma}(\mu_1^*, \mu_2^*, \gamma^*) < 0$. In that case, the minimizer is $\gamma^* = \bar{\gamma}$. So in both cases, $\gamma^* = \tilde{\gamma}$ as desired. □

A.18 Proof of Proposition 12

Proof. I start with the expression for the KL divergence from $\mathcal{H}^\bullet(c)$ to $\mathcal{H}(\Psi(\mu, \mu; \gamma); c)$. As in the proof of Proposition 2, this expression can be written as

$$\frac{(\mu - \mu^\bullet)^2}{2} + \int_{-\infty}^{c} \phi(x; \mu^\bullet, \sigma^2) \cdot \left[\frac{\sigma^2 + (\mu - \gamma(x_1 - \mu) - \mu^\bullet)^2}{2} - \frac{1}{2} \right] dx_1.$$

Dropping constant terms not depending on μ, we get a simplified expression of the objective,

$$\xi(\mu) := \frac{(\mu - \mu^\bullet)^2}{2} + \int_{-\infty}^{c} \phi(x; \mu^\bullet, \sigma^2) \cdot \left[\frac{(\mu - \gamma(x_1 - \mu) - \mu^\bullet)^2}{2} \right] dx_1.$$

Taking the first-order condition, $\xi'(\mu) = (\mu - \mu^\bullet) + (1+\gamma) \cdot \int_{-\infty}^{c} \phi(x_1; \mu^\bullet, \sigma^2) \cdot ((1+\gamma)\mu - \gamma x_1 - \mu^\bullet) dx_1$.

The term $\int_{-\infty}^{c} \phi(x_1; \mu^\bullet, \sigma^2) \cdot ((1+\gamma)\mu - \gamma x_1 - \mu^\bullet) dx_1$ may be rewritten as $\mathbb{P}[X_1 \le c] \cdot \mathbb{E}\left[(1+\gamma)\mu - \gamma X_1 - \mu^\bullet | X_1 \le c\right]$.

Setting the first-order condition to 0 and using straightforward algebra,

$$\mu_\triangle^*(c) = \frac{1}{1 + \mathbb{P}[X_1 \le c] \cdot (1+\gamma)^2} \mu^\bullet + \frac{\mathbb{P}[X_1 \le c] \cdot (1+\gamma)^2}{1 + \mathbb{P}[X_1 \le c] \cdot (1+\gamma)^2} \mu_2^\circ(c).$$

□

B Foundation for Inference and Behavior in the Large-Generation Environment

In Section 4, I introduced the large-generations social-learning environment with a continuum of agents in each generation. When agents in generations $\tau = 0, 1, ..., t-1$ choose cutoff thresholds $c_{[0]}, c_{[1]}, ..., c_{[t-1]}$, each generation t agent observes an infinite sample of histories $(h_{\tau,n})_{n \in [0,1]}$ drawn from the history distribution $\mathcal{H}^\bullet(c_\tau)$ for each $0 \le \tau \le t-1$. Agents infer the large-generations pseudo-true fundamentals $\mu_1^*(c_{[0]}, ..., c_{[t-1]})$, $\mu_2^*(c_{[0]}, ..., c_{[t-1]})$ and choose the stopping strategy that best responds to the feasible model with these parameters.

In this section, I provide a finite-population foundation for inference and behavior in the large-generations environment. For $K \ge 1$, let $c_\dagger = (c_\dagger^{(k)})_{k=1}^K \in \mathbb{R}^K$ be a list of cutoff thresholds. I show that when an agent starts with a full-support prior on the space of fundamentals $\mathbb{R}^2$ and observes $N < \infty$ histories drawn i.i.d. from each of $\mathcal{H}^\bullet(c_\dagger^{(k)})$ for $1 \leq k \leq K$, her posterior belief almost surely converges to the dogmatic belief on the large-generations pseudo-true fundamentals $\mu_1^*(c_\dagger), \mu_2^*(c_\dagger)$ as $N \to \infty$. Also, if she chooses the cutoff strategy S_c maximizing her posterior expected payoffs, then as $N \to \infty$ and provided the stage-game payoff functions u_1, u_2 are Lipschitz continuous, her cutoff choice almost surely converges to $C(\mu_1^*(c_\dagger), \mu_2^*(c_\dagger); \gamma)$.

B.1 Setting up the Probability Space

Suppose an agent has a full-support prior density $g : \mathbb{R}^2 \to \mathbb{R}_{>0}$ over fundamentals (μ_1, μ_2). To formally define the problem, consider the $\mathbb{R}^{2K}$-valued stochastic process $(X_n)_{n\geq 1} = (X_{1,n}^{(k)}, X_{2,n}^{(k)})_{1\leq k\leq K, n\geq 1}$, where X_s and $X_{s'}$ are independent for $s \neq s'$. Here, X_n are i.i.d. $\mathbb{R}^{2K}$-valued random variables with independent components, distributions as $X_{1,n}^{(k)} \sim \mathcal{N}(\mu_1^\bullet, \sigma^2)$, $X_{2,n}^{(k)} \sim \mathcal{N}(\mu_2^\bullet, \sigma^2)$ for each $1 \leq k \leq K$. The interpretation is that there are K different populations, who play the stage game using different cutoff thresholds. The random variables $(X_{1,n}^{(k)}, X_{2,n}^{(k)})$ are the *potential* draws in the n-th iteration of the stage game in population k, (but $X_{2,n}^{(k)}$ may not be observed if $X_{1,n}^{(k)}$ is sufficiently large). Clearly, there is a probability space $(\Omega, \mathcal{A}, \mathbb{P})$, with sample space $\Omega = (\mathbb{R}^{2K})^\infty$ interpreted as paths of the process just described, $\mathcal{A}$ the Borel σ-algebra on Ω, and $\mathbb{P}$ the measure on sample paths so that the process $X_n(\omega) = \omega_n$ has the desired distribution. The term "almost surely" means "with probability 1 with respect to the realization of infinite sequence of all (potential) draws", i.e. $\mathbb{P}$-almost surely.

For each $n \geq 1$ and $1 \leq k \leq K$, let $H_n^{(k)}$ be the (random) history given by $H_n^{(k)} = (X_{1,n}^{(k)}, \varnothing)$ if $X_{1,t}^{(k)} \geq c_\dagger^{(k)}$, $H_n^{(k)} = (X_{1,n}^{(k)}, X_{2,n}^{(k)})$ if $X_{1,n}^{(k)} < c_\dagger^{(k)}$. Let $H_n = (H_n^{(1)}, ..., H_n^{(K)})$. After each finite N, the agent Bayesian updates prior density g about the fundamentals, based on the finite dataset of histories $(H_n)_{n\leq N}$. She ends up with a random, non-degenerate posterior density $\tilde{g}_N = g(\cdot|(H_n)_{n\leq N})$, whose randomness comes from the randomness of the $2K \cdot N$ potential draws.

B.2 Inference after Observing Large Samples

Proposition A.1 shows that as $N \to \infty$, the random posterior $\tilde{g}_N$ converges to the large-generations pseudo-true fundamentals in L^1.

Proposition A.1. *Suppose $g : \mathbb{R}^2 \to \mathbb{R}_{>0}$ is integrable and has bounded magnitude. Almost surely,*

$$\lim_{N\to\infty} \mathbb{E}_{(\mu_1,\mu_2)\sim\tilde{g}_N} \left(|\mu_1 - \mu_1^\bullet| + |\mu_2 - \mu_2^*(c_\dagger)|\right) = 0.$$

Belief convergence in L^1 is required to later establish convergence of behavior in Proposition A.3. This convergence does not follow from Berk (1966), because his result only establishes convergence in a weaker mode: for any open set containing the pseudo-true fundamentals, the mass that the posterior belief assigns to the open set almost surely converges to 1. Crucially, the prior distribution in this setting has full support on an unbounded domain of feasible fundamentals, $(\mu_1, \mu_2) \in \mathcal{M} = \mathbb{R}^2$. Indeed, one of the implications of my central inference result, Proposition 2, is that the pseudo-true parameter becomes unboundedly pessimistic as censoring threshold decreases. So, the weak mode of convergence in Berk (1966)'s conclusion leaves open the possibility that posterior beliefs for increasing N put decreasing mass on increasingly extreme values of μ_2. If the magnitudes of these extreme values grow more quickly in N than the speed with which probability concentrates on the open set around the pseudo-true fundamentals, then there can be a positive-probability event where the agent's behavior is bounded away from $C(\mu_1^\bullet, \mu_2^*(c_\dagger); \gamma)$ for every N.

Instead, I apply Bunke and Milhaud (1998)'s results to derive the stronger convergence in L^1 that subsequently allows for convergence of payoffs and behavior as the agent's sample grows large. One technical challenge is that the results of Bunke and Milhaud (1998) only apply in environments where observables are valued in some Euclidean space and given by densities, but censored histories are valued in $\mathbb{H}$ and their distributions have a probability mass on the missing data indicator $\varnothing$. So, I first consider a *noise-added* observation structure where each history $H_n^{(k)}$ is replaced by the $\mathbb{R}^2$-valued pair $(X_{n,1}^{(k)}, Y_n^{(k)})$, where $Y_n^{(k)} = X_{n,2}^{(k)}$ if $X_{n,1}^{(k)} \le c_\dagger^{(k)}$. But if $X_{n,1}^{(k)} > c_\dagger^{(k)}$, then $Y_n^{(k)} \sim \mathcal{N}(0,1)$ is a white noise term that is independent of the draws of any decision problem. The idea is that a censored draw is replaced by noise that is uninformative about the fundamentals, so the distribution of each $(X_{n,1}^{(k)}, Y_n^{(k)})$ pair is given by a density function on $\mathbb{R}^2$. After establishing the analogous belief convergence result in the auxiliary environment, I map the result back into the environment of observing censored histories. This translation is possible because in every finite dataset, the realizations of the white noise terms do not change the *relative* likelihoods of data under different parameters (μ_1, μ_2), hence they do not affect the agent's posterior belief over fundamentals.

I now formally define this noise-added observation structure that replaces censored X_2's with white noise. Let $\mathbb{P}_{Z^\infty}$ be the measure on $(\mathbb{R}^\infty)^K$ corresponding to product of K i.i.d. sequence of $\mathcal{N}(0,1)$ random variables. Consider the expanded probability space $(\bar{\Omega}, \bar{\mathcal{A}}, \bar{\mathbb{P}})$ where $\bar{\Omega} = \Omega \times (\mathbb{R}^\infty)^K$, $\bar{\mathcal{A}}$ is the product σ-algebra on $\bar{\Omega}$ where $(\mathbb{R}^\infty)^K$ is endowed with the usual product Borel σ-algebra, and $\bar{\mathbb{P}}$ is the product measure $\mathbb{P} \otimes \mathbb{P}_{Z^\infty}$ on $\bar{\Omega}$. To interpret, each element $\bar{\omega} = (\omega, z) \in \bar{\Omega}$ consists of the sample path of a sequence of potential draws $(X_n)_{n=1}^\infty$ as well as the sample path of a sequence of white noise realizations $(Z_n)_{n=1}^\infty$, where each Z_n is an $\mathbb{R}^K$-valued random variable.

On the expanded probability space, we can define two kinds of observations. The history dataset of size N is $(H_n(\bar{\omega}))_{n\leq N} = (H_n(\omega_n))_{n\leq N}$, as the K round n histories only depends on ω_n (and not on the white noise z_n). The noise-added dataset of size T is $(X_{1,n}(\omega), Y_n(\omega, z))_{n\leq N} = (X_{1,n}(\omega_n), Y_n(\omega_n, z_n))_{n\leq N}$. Write $\tilde{g}_N^H$ as the posterior density from history dataset of size N, $\tilde{g}_N^{XY}$ as the posterior density from noise-added dataset of size N.

The next lemma formalizes the idea that replacing censored observations with white noise does not affect posterior beliefs.

Lemma A.14. *For every $\bar{\omega} \in \bar{\Omega}$ and $N \in \mathbb{N}$, $\tilde{g}_N^H(\bar{\omega}) = \tilde{g}_N^{XY}(\bar{\omega})$.*

Proof. Suppose $\bar{\omega} = ((x_{1,n}, x_{2,n})_{n=1}^\infty, (z_n)_{n=1}^\infty) \in \Omega \times (\mathbb{R}^\infty)^K$. The noise-added dataset of size N is $(x_{1,n}, y_n)_{n=1}^N$ where $y_n^{(k)} = z_n^{(k)}$ for each n, k where $x_{1,n}^{(k)} \geq c_\dagger^{(k)}$, and $y_n^{(k)} = x_{2,n}^{(k)}$ for each n, k where $x_{1,n}^{(k)} < c_\dagger^{(k)}$. The history dataset of size N is $(h_n)_{n=1}^N$, where $h_n^{(k)} = (x_{1,n}^{(k)}, \varnothing)$ for each n, k where $x_{1,n}^{(k)} \geq c_\dagger^{(k)}$, and $h_n^{(k)} = (x_{1,n}^{(k)}, x_{2,n}^{(k)})$ for each n, k where $x_{1,n}^{(k)} < c_\dagger^{(k)}$.

The likelihood of the noise-added dataset under parameters μ_1, μ_2 is:

$$\prod_{k=1}^K \left\{ \left(\prod_{n=1}^N \phi(x_{1,n}^{(k)}; \mu_1, \sigma^2) \right) \cdot \left(\prod_{n: x_n^{(k)} \leq c_\dagger^{(k)}} \phi(y_n^{(k)}; \mu_2 - \gamma(x_{1,n}^{(k)} - \mu_1), \sigma^2) \right) \cdot \left(\prod_{n: x_n^{(k)} \geq c_\dagger^{(k)}} \phi(y_n^{(k)}; 1, 0) \right) \right\}$$

The likelihood of the history dataset under parameters μ_1, μ_2 is:

$$\prod_{k=1}^K \left\{ \left(\prod_{n=1}^N \phi(x_{1,n}^{(k)}; \mu_1, \sigma^2) \right) \cdot \left(\prod_{n: x_n^{(k)} \leq c_\dagger^{(k)}} \phi(y_n^{(k)}; \mu_2 - \gamma(x_{1,n}^{(k)} - \mu_1), \sigma^2) \right) \right\}$$

So, these likelihoods are equal up to a multiple of $\prod_{k=1}^K \left(\prod_{n: x_n^{(k)} \geq c_\dagger^{(k)}} \phi(y_n^{(k)}; 1, 0) \right)$, which is common across all parameters (μ_1, μ_2). So the posterior likelihood of parameters μ_1, μ_2 must be the same under both $\tilde{g}_N^H$ and $\tilde{g}_N^{XY}$, that is $\tilde{g}_N^H(\bar{\omega}) = \tilde{g}_N^{XY}(\bar{\omega})$. □

On the expanded probability space, inference from history dataset and inference from noise-added dataset give the same posterior beliefs everywhere. If $\tilde{g}_N^{XY}$ converges in L^1 to dogmatic belief on $(\mu_1^\bullet, \mu_2^*(c_\dagger))$ $\bar{\mathbb{P}}$-a.s., then $\tilde{g}_N^H$ also converges in L_1 to the same belief $\bar{\mathbb{P}}$-a.s. Further, by relationship between the expanded probability space and the original probability space, this would also show that $\tilde{g}_N$ converges in L_1 to dogmatic belief on $(\mu_1^\bullet, \mu_2^*(c_\dagger))$ $\mathbb{P}$-a.s., which proves Proposition A.1. Therefore, to prove Proposition A.1 one just needs the following on the expanded probability space.

Lemma A.15. *$\tilde{g}_N^{XY}$ converges in L^1 to the dogmatic belief on $(\mu_1^\bullet, \mu_2^*(c_\dagger))$ $\bar{\mathbb{P}}$-a.s.*

Proof. First, I write down the KL divergence objective in the noise-added observation structure and show its minimizers are exactly the large-generations pseudo-true fundamentals.

Each observation $(X_{1,n}^{(k)}, Y_n^{(k)})_{k=1}^K$ is an element of $\mathbb{R}^{(2K)}$, whose distribution is given by a K densities over K copies of $\mathbb{R}^2$. For $1 \le k \le K$, the k-th such density is

$$f^{\bullet,(k)}(x,y) = \begin{cases} \phi(x;\mu_1^\bullet,\sigma^2)\cdot\phi(y;\mu_2^\bullet,\sigma^2) & \text{if } x < c_\dagger^{(k)} \\ \phi(x;\mu_1^\bullet,\sigma^2)\cdot\phi(y;0,1) & \text{if } x \ge c_\dagger^{(k)}. \end{cases}$$

Under the fundamentals $(\mu_1,\mu_2) \in \mathbb{R}^2$, the agent thinks the observations are distributed according to the product of K densities where the k-th density is

$$f_{\hat{\mu}_1,\hat{\mu}_2}^{(k)}(x,y) = \begin{cases} \phi(x;\hat{\mu}_1,\sigma^2)\cdot\phi(y;\hat{\mu}_2-\gamma\cdot(x-\hat{\mu}_1),\sigma^2) & \text{if } x < c_\dagger^{(k)} \\ \phi(x;\hat{\mu}_1,\sigma^2)\cdot\phi(y;0,1) & \text{if } x \ge c_\dagger^{(k)}. \end{cases}$$

The log likelihood ratio of an observation $(x,y) = (x_1^{(k)}, y^{(k)})_{k=1}^K \in \mathbb{R}^{2K}$ is

$$\ln\left[\prod_{k=1}^K \frac{f^{\bullet,(k)}(x_1^{(k)},y^{(k)})}{f_{\hat{\mu}_1,\hat{\mu}_2}^{(k)}(x_1^{(k)},y^{(k)})}\right] = \sum_{k=1}^K \ln\left[\frac{f^{\bullet,(k)}(x_1^{(k)},y^{(k)})}{f_{\hat{\mu}_1,\hat{\mu}_2}^{(k)}(x_1^{(k)},y^{(k)})}\right].$$

So KL divergence is defined as

$$\begin{aligned} &\int_{\mathbb{R}^{2K}} \left(\sum_{k=1}^K \ln\left[\frac{f^{\bullet,(k)}(x_1^{(k)},y^{(k)})}{f_{\hat{\mu}_1,\hat{\mu}_2}^{(k)}(x_1^{(k)},y^{(k)})}\right]\right)\cdot\left(\prod_{k=1}^K f^{\bullet,(k)}(x_1^{(k)},y^{(k)})\right) d(x,y) \\ =&\sum_{k=1}^K \int_{\mathbb{R}^{2K}} \ln\left[\frac{f^{\bullet,(k)}(x_1^{(k)},y^{(k)})}{f_{\hat{\mu}_1,\hat{\mu}_2}^{(k)}(x_1^{(k)},y^{(k)})}\right]\cdot\left(\prod_{j=1}^K f^{\bullet,(j)}(x_1^{(j)},y^{(j)})\right) d(x,y). \end{aligned}$$

But for each k, the integrand $\frac{f^{\bullet,(k)}(x_1^{(k)},y^{(k)})}{f_{\hat{\mu}_1,\hat{\mu}_2}^{(k)}(x_1^{(k)},y^{(k)})}$ only depends on $(x,y) \in \mathbb{R}^{2K}$ through two of its coordinates, $x_1^{(k)}$ and $y^{(k)}$. In addition, the density $\prod_{j=1}^K f^{\bullet,(j)}(x_1^{(j)},y^{(j)})$ is a product density, so in fact the k-th summand is just

$$\int_{\mathbb{R}^2} \ln\left[\frac{f^{\bullet,(k)}(x_1^{(k)},y^{(k)})}{f_{\hat{\mu}_1,\hat{\mu}_2}^{(k)}(x_1^{(k)},y^{(k)})}\right]\cdot f^{\bullet,(k)}(x_1^{(k)},y^{(k)}) d(x_1^{(k)},y^{(k)}).$$

This expression is, up to a constant not depending on $\hat{\mu}_1,\hat{\mu}_2$ (due to the white noise term), equal to the KL divergence between $\mathcal{H}^\bullet(c_\dagger^{(k)})$ and $\mathcal{H}(\Psi(\hat{\mu}_1,\hat{\mu}_2;\gamma);c_\dagger^{(k)})$. Therefore the overall KL divergence is off by a constant from

$$\sum_{k=1}^K D_{KL}(\ \mathcal{H}^\bullet(c_\dagger^{(k)}) \ ||\ \mathcal{H}(\Psi(\hat{\mu}_1,\hat{\mu}_2;\gamma);c_\dagger^{(k)})\ ),$$

the objective defining large-generations pseudo-true fundamentals in Equation (1).

To finish the proof, Bunke and Milhaud (1998) show that provided the true density $f^\bullet$ and the family of subjective densities $\{f_{\hat{\mu}_1,\hat{\mu}_2} : \hat{\mu}_1, \hat{\mu}_2 \in \mathbb{R}\}$ satisfy a number of conditions, then $\tilde{g}_N^{XY}$ $\bar{\mathbb{P}}$-a.s. converges to its KL-divergence minimizers in L^1, which I have shown to be exactly $(\mu_1^\bullet, \mu_2^*(c_\dagger))$. I check the conditions of Bunke and Milhaud (1998) in the Online Appendix for the case of $K = 1$, so both $f^\bullet$ and $f_{\hat{\mu}_1,\hat{\mu}_2}$ are densities on $\mathbb{R}^2$. Checking the conditions for larger K is exactly analogous, because both $f^\bullet$ and $f_{\hat{\mu}_1,\hat{\mu}_2}$ can be separated as the product of K densities on $\mathbb{R}^2$. $\square$

B.3 Behavior after Observing Large Samples

Next, I turn to the convergence of expected payoffs for different cutoff strategies as sample size grows large. For any $c \in \mathbb{R}$ and $N \in \mathbb{N}$, let $U_N(c) := \mathbb{E}_{(\mu_1,\mu_2)\sim\tilde{g}_N}[U(c;\mu_1,\mu_2,\gamma)]$ where $U(c;\mu_1,\mu_2,\gamma)$ is the expected payoff of using the stopping strategy S_c when $(X_1, X_2) \sim \Psi(\mu_1,\mu_2;\gamma)$. Note that $U_N(c)$ is a real-valued random variable representing the agent's subjective expected payoff for the stopping strategy S_c, under the (random) non-degenerate posterior belief after observing a sample of size N. Proposition A.2 shows that $U_N(c)$ converges almost surely to the subjective expected payoff of S_c with a dogmatic belief in the pseudo-true fundamentals, provided the payoff functions u_1, u_2 of the optimal-stopping problem are Lipschitz continuous. Furthermore, this convergence is uniform across all cutoff thresholds.

Proposition A.2. *Suppose there are constants $K_1, K_2 > 0$ so that $|u_1(x_1') - u_1(x_1'')| < K_1 \cdot |x_1' - x_1''|$ and $|u_2(x_1', x_2') - u_2(x_1'', x_2'')| < K_2 \cdot (|x_1' - x_1''| + |x_2' - x_2''|)$ for all $x_1', x_1'', x_2', x_2'' \in \mathbb{R}$. Then $\lim_{N\to\infty} \sup_{c\in\mathbb{R}} |U_N(c) - U(c;\mu_1^\bullet, \mu_2^*(c_\dagger))| = 0$ almost surely.*

The Lipschitz continuity conditions are satisfied in the search problem (Example 1) for any $q \in [0, 1)$. The Lipschitz condition implies the difference between expected payoffs of S_c under the posterior belief $\tilde{g}_N$ and the dogmatic belief on the pseudo-true fundamental is bounded by a constant multiple of the L^1 distance between $\tilde{g}_N$ and the pseudo-true fundamentals, and furthermore this bound is uniform across all c. Proposition A.2 then follows from the L^1 convergence in beliefs given by Proposition A.1.

Towards a proof of Proposition A.2, I start with a lemma that shows when the optimal-stopping problem's payoff functions u_1, u_2 are Lipschitz continuous, then $(\mu_1, \mu_2) \mapsto U(c;\mu_1,\mu_2)$, the expected payoff of the stopping strategy S_c under the subjective model $\Psi(\mu_1,\mu_2;\gamma)$, is locally Lipschitz continuous in (μ_1, μ_2).

Lemma A.16. *Suppose there are constants $K_1, K_2 > 0$ so that $|u_1(x_1') - u_1(x_1'')| < K_1 \cdot |x_1' - x_1''|$ and $|u_2(x_1', x_2') - u_2(x_1'', x_2'')| < K_2 \cdot (|x_1' - x_1''| + |x_2' - x_2''|)$ for all $x_1', x_1'', x_2', x_2'' \in \mathbb{R}$. For each center $(\mu_1^\circ, \mu_2^\circ) \in \mathbb{R}^2$, there corresponds a constant $K > 0$ so that for any $\mu_1, \mu_2 \in \mathbb{R}$ and any $c' \in \mathbb{R}$, $|U(c;\mu_1,\mu_2) - U(c;\mu_1^\circ,\mu_2^\circ)| < K \cdot (|\mu_1 - \mu_1^\circ| + |\mu_2 - \mu_2^\circ|)$.*

Now I prove Proposition A.2.

Proof. Let $\mu_1^\circ = \mu_1^\bullet$, $\mu_2^\circ = \mu_2^*(c_\dagger)$. Lemma A.16 implies there is a constant $K > 0$, independent of c, so that $|U(c;\mu_1,\mu_2) - U(c;\mu_1^\circ,\mu_2^\circ)| \le K \cdot (|\mu_1 - \mu_1^\circ| + |\mu_2 - \mu_2^\circ|)$ for all $\mu_1, \mu_2, c \in \mathbb{R}$. So for ν a joint distribution about the fundamentals (μ_1, μ_2), we get

$$\begin{aligned}|\mathbb{E}_{(\mu_1,\mu_2)\sim\nu}\left[U(c;\mu_1,\mu_2) - U(c;\mu_1^\circ,\mu_2^\circ)\right]| &\le \mathbb{E}_{(\mu_1,\mu_2)\sim\nu}\left[|U(c;\mu_1,\mu_2) - U(c;\mu_1^\circ,\mu_2^\circ)|\right] \\ &\le K \cdot \mathbb{E}_{(\mu_1,\mu_2)\sim\nu}[|\mu_1 - \mu_1^\circ| + |\mu_2 - \mu_2^\circ|]\end{aligned}$$

for every $c \in \mathbb{R}$, therefore we also get the uniform bound,

$$\sup_{c\in\mathbb{R}} |\mathbb{E}_{(\mu_1,\mu_2)\sim\nu}\left[U(c;\mu_1,\mu_2)\right] - U(c;\mu_1^\circ,\mu_2^\circ)| \le K \cdot \mathbb{E}_{(\mu_1,\mu_2)\sim\nu}[|\mu_1 - \mu_1^\circ| + |\mu_2 - \mu_2^\circ|].$$

By Proposition A.1, almost surely

$$\lim_{T\to\infty} \mathbb{E}_{(\mu_1,\mu_2)\sim\tilde{g}_T}[|\mu_1 - \mu_1^\circ| + |\mu_2 - \mu_2^\circ|] = 0.$$

But along any $\omega \in \Omega$ where the above limit holds,

$$\begin{aligned}\lim_{T\to\infty} \sup_{c\in\mathbb{R}} |U_T(c) - U(c;\mu_1^\circ,\mu_2^\circ)| &\le \lim_{T\to\infty} K \cdot \mathbb{E}_{(\mu_1,\mu_2)\sim\tilde{g}_T}[|\mu_1 - \mu_1^\circ| + |\mu_2 - \mu_2^\circ|] \\ &= 0.\end{aligned}$$

This shows that $\mathbb{P}$-a.s., $U_T(c)$ converges to $U(c;\mu_1^*(c_\dagger),\mu_2^*(c_\dagger))$ uniformly across all c as $T \to \infty$. ◻

To reach my main result on convergence of behavior, suppose the agent chooses a cutoff threshold after observing N histories $(h_n)_{n\le N}$. The choices are given by the functions $\tilde{C}_N : \mathbb{H}^N \to \mathbb{R}$, so the cutoff after a sample of size N is a random variable C_N that depends on the first N pairs of potential draws $(X_n)_{n\le N}$.

Definition A.1. Cutoff choice functions $(\tilde{C}_N)$ are **asymptotically myopic** in N if

$$\limsup_{N\to\infty} \left\{ \sup_{c\in\mathbb{R}} U_N(c) - U_N(\tilde{C}_N) \right\} = 0$$

almost surely.

A simple example is that $\tilde{C}_N$ chooses a cutoff whose expected payoff differs from $\sup_{c\in\mathbb{R}} U_N(c)$ by no more than $1/N$ after every sample of size N.

Proposition A.3. *Let $c^* = C(\mu_1^\bullet, \mu_2^*(c_\dagger);\gamma)$. Suppose cutoffs C_N are generated using asymptotically myopic cutoff choice functions. Almost surely, $C_N \to c^*$ as $N \to \infty$.*

The expected payoff of different cutoff strategies under the pseudo-true fundamentals, $c \mapsto U(c; \mu_1^\bullet, \mu_2^*(c_\dagger))$, is single peaked and maximized at c^*. Therefore cutoffs outside an open ball around c^* have expected payoffs bounded away from the subjectively optimal payoff under the model $\Psi(\mu_1^\bullet, \mu_2^*(c_\dagger); \gamma)$.

Lemma A.17. *For each $\mu_1, \mu_2 \in \mathbb{R}$, let c^* be the subjectively optimal cutoff threshold under $\Psi(\mu_1, \mu_2; \gamma)$. For every $\epsilon > 0$, there exists $\delta > 0$ so that whenever $|c - c^*| \geq \epsilon$, we have $U(c; \mu_1, \mu_2) \leq U(c^*; \mu_1, \mu_2) - \delta$.*

Proof. First, I show $c \mapsto U(c; \mu_1, \mu_2)$ is single peaked: it is strictly increasing up to $c = c^*$, then strictly decreasing afterwards. Recall the cutoff form of the best stopping strategy comes from the fact that $u_1(x_1) < \mathbb{E}_{\Psi(\mu_1,\mu_2;\gamma)}[u_2(x_1, X_2)|X_1 = x_1]$ for $x_1 < c^*$, but $u_1(x_1) < \mathbb{E}_{\Psi(\mu_1,\mu_2;\gamma)}[u_2(x_1, X_2)|X_1 = x_1]$ for $x_1 > c^*$. For two cutoffs $c_1 < c_2 < c^*$, the two stopping strategies S_{c_1}, S_{c_2} only differ in how they treat first-period draws in the interval $[c_1, c_2]$, so we can write the difference in their expected payoffs as

$$\int_{c_1}^{c_2} \left(\mathbb{E}_{\Psi(\mu_1,\mu_2;\gamma)}[u_2(x_1, X_2)|X_1 = x_1] - u_1(x_1)\right) \phi(x_1; \mu_1, \sigma^2) dx_1.$$

The integrand is strictly positive on $[c_1, c_2]$, therefore $U(c_1; \mu_1, \mu_2) < U(c_2; \mu_1, \mu_2)$. This shows $U(\cdot; \mu_1, \mu_2)$ is strictly increasing up until c^*; a symmetric argument shows it is strictly decreasing after c^*.

For a given $\epsilon > 0$, let $\delta = U(c^*; \mu_1, \mu_2) - \max(U(c^* - \epsilon; \mu_1, \mu_2), U(c^* + \epsilon; \mu_1, \mu_2))$, where $\delta > 0$ as both cutoffs $c^* - \epsilon$ and $c^* + \epsilon$ must have a strictly positive loss relative to c^*. Since $U(\cdot; \mu_1, \mu_2)$ is single peaked, every c more than ϵ away from c^* must have a loss relative to c^* at least as much as the loss of either $c^* - \epsilon$ or $c^* + \epsilon$, so $U(c^*; \mu_1, \mu_2) - U(c; \mu_1, \mu_2) \geq \delta$. □

This fact, combined with the uniform convergence $U_N(c)$ from Proposition A.2, establishes Proposition A.3.

Proof. Consider any sample path $\omega = (x_n)_{n=1}^{\infty}$ where the conclusion of Proposition A.2 holds and the cutoff choice functions are asymptotically myopic. For every $\epsilon > 0$, find $\delta > 0$ as in Lemma A.17 with $\mu_1 = \mu_1^\bullet$, $\mu_2 = \mu_2^*(c_\dagger)$, and find large enough $\bar{N}_1$ so that $\sup_{c \in \mathbb{R}} |U_N(c)(\omega) - U(c; \mu_1^\bullet, \mu_2^*(c_\dagger))| < \delta/3$ for all $N \geq \bar{N}_1$. This means for $N \geq \bar{N}_1$,

$$\sup_{c \in \mathbb{R}} U_N(c)(\omega) \geq U_N(c^*)(\omega) \geq U(c^*; \mu_1^\bullet, \mu_2^*(c_\dagger)) - \delta/3,$$

while

$$U_N(c^{'})(\omega) \leq U(c^*; \mu_1^\bullet, \mu_2^*(c_\dagger)) - (2\delta)/3$$

for $c^{'} \notin [c^* - \epsilon, c^* + \epsilon]$. Find $\bar{N}_2$ so that for $N \geq \bar{N}_2$, $\sup_{c \in \mathbb{R}} U_N(c)(\omega) - U_N(C_N)(\omega) < \delta/3$. This shows for $N \geq \max(\bar{N}_1, \bar{N}_2)$, $C_N(\omega) \in [c^* - \epsilon, c^* + \epsilon]$. Since $\epsilon > 0$ was arbitrary, this shows $C_N(\omega) \to c^*$.

Therefore, we conclude $C_N \to c^*$ along any sample path ω where the conclusion of Proposition A.2 holds and the cutoff choice functions are asymptotically myopic. Since these two events both happen almost surely, $C_N \to c^*$ almost surely. $\square$

C General Non-Gaussian Feasible Models and Method of Moments Inference

In the analysis so far, both the true distribution of draws and the class of feasible models come from the Gaussian family. The Gaussian distributional assumption makes the agents' inference problem analytically tractable, since the KL divergence between a pair of Gaussian random variables has a simple closed-form expression. But, the intuition behind my main results—over-pessimism about the fundamentals and the positive-feedback loop between beliefs and cutoffs—holds more generally. KL divergence minimization in the Gaussian environment has a method of moments (MOM) interpretation, as mentioned in Footnote 11. In this section, I consider agents using the same MOM procedure as a simpler but natural alternative to Bayesian inference, but generalize the class of feasible models (and the true distribution) of the draws substantially. Proposition A.5 and Corollary A.1 show that over-pessimism and the positive-feedback loop remain robust to this joint relaxation of Bayesian inference and Gaussian distributional assumption.

C.1 Feasible Models for (X_1, X_2)

Each agent starts with a family of feasible models $\{F(\cdot;\theta_1,\theta_2) : \theta_1 \in \Theta_1, \theta_2 \in \Theta_2\}$ for the joint distribution of (X_1, X_2), with parameter spaces $\Theta_1 \subseteq \mathbb{R}$ and $\Theta_2 \subseteq \mathbb{R}$. For each $(\theta_1\theta_2)$, $F(\cdot;\theta_1,\theta_2)$ is a full-support measure on the rectangle $I_1 \times I_2$, where each I_1, I_2 is a possibly infinite interval of $\mathbb{R}$. By "full-support" I mean that for every open ball $B \subseteq I_1 \times I_2$, $F(B;\theta_1,\theta_2) > 0$.

For each joint distribution $F(\cdot;\theta_1,\theta_2)$, let $F_1(\cdot;\theta_1,\theta_2)$ denote its marginal on I_1, and let $F_{2|1}(\cdot|\theta_1,\theta_2;x_1)$ denote its conditional distribution on I_2 given $X_1 = x_1$. I will make the following assumptions on the family of feasible models:

Assumption A.1. *The feasible models $\{F(\cdot;\theta_1,\theta_2) : \theta_1 \in \Theta_1, \theta_2 \in \Theta_2\}$ satisfy :*

(a) $F_1(\cdot;\theta_1,\theta_2)$ is only a function of θ_1 and $\mathbb{E}_{F_1(\cdot;\theta_1,\theta_2)}[X_1]$ is strictly increasing in θ_1.

(b) For each $x_1 \in I_1$ and $\theta_1 \in \Theta_1$, $\mathbb{E}_{F_{2|1}(\cdot;\theta_1,\theta_2|x_1)}[X_2]$ strictly increases in θ_2.

(c) For any $\theta_1 \in \Theta_1$ and $\theta_2 \in \Theta_2$, $\mathbb{E}_{F_{2|1}(\cdot;\theta_1,\theta_2|x_1)}[X_2]$ strictly decreases in x_1.

In light of Assumption A.1(a), the marginal distribution on X_1 can be just written as $F_1(\cdot;\theta_1)$, omitting θ_2. Assumption A.1(c) is the substantive assumption capturing the gambler's fallacy psychology. Every subjective distribution in the family is such that the agent predicts a lower mean for X_2 after a higher realization of X_1. The behavioral economics literature has not settled on a general definition of the gambler's fallacy that works under all distributional assumptions, but Assumption A.1(c) seems like a reasonable first step.

Here are some examples satisfying these assumptions. The first example shows the family of Gaussian distributions I have worked with in the main text satisfies Assumption A.1.

Example A.1. Let $I_1 = I_2 = \mathbb{R}$ and let $\Theta_1 = \Theta_2 = \mathbb{R}$. Fixing some $\sigma^2 > 0$, $\gamma > 0$, let $F(\cdot;\theta_1,\theta_2)$ be $\Psi(\theta_1,\theta_2,\sigma^2,\sigma^2;\gamma)$ for each $\theta_1,\theta_2 \in \mathbb{R}$. The marginal distribution on X_1 is $\mathcal{N}(\theta_1,\sigma^2)$ and does not depend on θ_2. Its mean is θ_1 so it strictly increases in θ_1. The conditional mean of $X_2|X_1 = x_1$ is $\theta_2 - \gamma(x_1 - \theta_1)$, which is strictly increasing in θ_2 and strictly decreasing in x_1 since $\gamma > 0$. So all conditions in Assumption A.1 are satisfied.

The next example features bivariate exponential distributions supported on the half-line $[0,\infty)$.

Example A.2. Gumbel (1960) proposes the following family of bivariate exponential distributions, parametrized by $\alpha \in [-1,1]$: consider a joint distribution with the density function $(\tilde{x}_1,\tilde{x}_2) \mapsto e^{-\tilde{x}_1-\tilde{x}_2} \cdot [1 + \alpha(2e^{-\tilde{x}_1} - 1)\cdot(2e^{-\tilde{x}_2} - 1)]$ for $\tilde{x}_1,\tilde{x}_2 \geq 0$. If $(\tilde{X}_1,\tilde{X}_2)$ are random variables with this density, then they have full support on $[0,\infty)\times[0,\infty)$ and each $\tilde{X}_j$ has the marginal distribution of an exponential random variable with mean 1. The conditional distribution of $\tilde{X}_2$ given a realization of $\tilde{X}_1$ is $\mathbb{E}[\tilde{X}_2|\tilde{X}_1 = \tilde{x}_1] = 1 - \frac{1}{2}\alpha - \alpha e^{-\tilde{x}_1}$. The correlation between $\tilde{X}_1$ and $\tilde{X}_2$ is $\alpha/4$.

Let $I_1 = I_2 = [0,\infty)$ and let $\Theta_1 = \Theta_2 = (0,\infty)$. Fixing some $-1 \leq \alpha < 0$, let $F(\cdot;\theta_1,\theta_2)$ be the joint distribution generated by $X_1 = \theta_1 \cdot \tilde{X}_1$ and $X_2 = \theta_2 \cdot \tilde{X}_2$ where $(\tilde{X}_1,\tilde{X}_2)$ have the Gumbel bivariate distribution with parameter α. Since $(\tilde{X}_1,\tilde{X}_2)$ have full support on $I_1 \times I_2$, the same holds for (X_1,X_2) for any $\theta_1,\theta_2 > 0$. The marginal distribution of X_1 is exponential with a mean of θ_1, so Assumption A.1(a) is satisfied. The conditional mean of $X_2|X_1 = x_1$ is given by $\mathbb{E}[\theta_2\tilde{X}_2|\theta_1\tilde{X}_1 = x_1] = \theta_2 \cdot \mathbb{E}\left[\tilde{X}_2|\tilde{X}_1 = \frac{x_1}{\theta_1}\right] = \theta_2 \cdot \left(1 - \frac{1}{2}\alpha - \alpha e^{-(x_1/\theta_1)}\right)$. As $\alpha < 0$, the term inside the bracket is strictly positive. So this conditional expectation is strictly increasing in θ_2, showing that Assumption A.1(b) is satisfied. Also, since $\theta_1,\theta_2 > 0$, $x_1 \mapsto -\alpha\theta_2 e^{-(x_1/\theta_1)}$ is strictly decreasing and so Assumption A.1(c) is satisfied.

I give a third example where $I_1 = I_2 = [0,1]$ are bounded intervals.

Example A.3. Let $\Theta_1 = \Theta_2 = (0, \infty)$ and consider the family of distribution $F(\cdot;\theta_1,\theta_2)$ such that under parameters (θ_1, θ_2), $X_1 \sim \text{Beta}(\theta_1, 1)$ and $X_2|X_1 = x_1 \sim \text{Beta}((1-x_1)\theta_2, 1)$. For any values of $\theta_1, \theta_2 > 0$, X_1 has full support on $[0, 1]$. Conditional on any $x_1 \in (0, 1)$, X_2 has full support on $[0, 1]$. This shows the distribution $F(\cdot;\theta_1,\theta_2)$ has full-support on $[0, 1]^2$ for every $(\theta_1, \theta_2) \in \Theta_1 \times \Theta_2$. The mean of X_1 is $\frac{\theta_1}{\theta_1+1}$, which only depends on θ_1 and is strictly increasing in it. So Assumption A.1(a) is satisfied. The conditional mean of $X_2|X_1 = x_1$ is $\frac{(1-x_1)\theta_2}{(1-x_1)\theta_2+1}$, which is strictly increasing in θ_2 and strictly decreasing in x_1. So, Assumptions A.1(b) and A.1(c) are satisfied.

Finally, I give a general class of examples that allows any pair of given marginal distributions for X_1 and X_2 to be joined together using a copula as to induce negative dependence for the joint distribution.

Example A.4. Consider two families of distribution functions $Q_1(\cdot;\theta_1) : I_1 \to [0, 1]$, $Q_2(\cdot;\theta_2) : I_2 \to [0, 1]$, such Q_1 and Q_2 are supported on I_1, I_2 respectively for all $\theta_1 \in \Theta_1$ and $\theta_2 \in \Theta_2$. Suppose the mean of Q_1 is increasing in θ_1, and Q_2 is increasing in stochastic dominance order with respect to θ_2. Fix a differentiable copula: that is, a differentiable function $W : [0, 1]^2 \to [0, 1]$ so that $W(u, 0) = W(0, v) = 0$, $W(u, 1) = u$, $W(1, v) = v$ for all $u, v \in [0, 1]$, and so that for $u_1 \le u_2$, $v_2 \le v_2 \in [0, 1]$, we get $W(u_2, v_2) - W(u_2, v_1) - W(u_1, v_2) - W(u_1, v_1) \ge 0$. Consider the family of distribution functions $Q(\cdot;\theta_1,\theta_2)$ on $\mathbb{R}^2$ generated by joining together $Q_1(\cdot;\theta_1)$ with $Q_2(\cdot;\theta_2)$ using the copula W, namely

$$Q((-\infty, x_1] \times (-\infty, x_2]; \theta_1, \theta_2) = W(Q_1^{-1}(x_1|\theta_1), Q_2^{-1}(x_2|\theta_2)).$$

Then $Q(\cdot;\theta_1,\theta_2)$ has marginal distributions on X_1 and X_2 given by distribution functions $Q_1(\cdot;\theta_1)$, $Q_2(\cdot;\theta_2)$, and:

Lemma A.18. *Suppose $\frac{\partial W}{\partial u}(u, v)$ is an increasing function in u and that $\{Q_1(\cdot;\theta_1) : \theta_1 \in \Theta_1\}$, $\{Q_2(\cdot;\theta_2) : \theta_2 \in \Theta_2\}$ satisfy the conditions of this example. Then, the conditions in Assumption A.1 are satisfied for the family of distributions $F(\cdot;\theta_1,\theta_2)$ where $F(\cdot;\theta_1,\theta_2)$ has the distribution function $Q(\cdot;\theta_1,\theta_2)$.*

The condition that $\frac{\partial W}{\partial u}(u, v)$ increases in u is satisfied by, for example, the Gaussian copula with any negative correlation. The derivative of the Gaussian copula is given by $\frac{\partial W}{\partial u}(u, v) = \mathbb{P}[X_2 \le \Phi^{-1}(v)|X_1 = \Phi^{-1}(u)]$ where Φ is the standard Gaussian distribution function and (X_1, X_2) are jointly Gaussian with correlation $-1 < \rho < 0$ and each with an unconditional variance of 1. As it is known that $X_2|X_1 = x_1 \sim \mathcal{N}(\rho x_1, 1 - \rho^2)$, it is clear that $X_2|X_1 = x_1$ decreases in FOSD order as x_1 increases, so for any v we have $\mathbb{P}[X_2 \le \Phi^{-1}(v)|X_1 = \Phi^{-1}(u)]$ increases in u.

C.2 Method of Moments Inference

For a distribution of histories $\mathcal{H} \in \Delta(\mathbb{H})$, let $m_1[\mathcal{H}]$ represent the average first-period draw under this distribution and let $m_2[\mathcal{H}]$ represent the average second-period draw (when uncensored). More precisely, $m_1[\mathcal{H}] := \mathbb{E}_{h\sim\mathcal{H}}[h_1]$, $m_2[\mathcal{H}] := \mathbb{E}_{h\sim\mathcal{H}}[h_2 \mid h_2 \neq \varnothing]$. Suppose that objectively X_1, X_2 are independent with a joint distribution $F^\bullet$, and denote the true distribution of histories under censoring by cutoff stopping rule $c \in \mathbb{R}$ as $\mathcal{H}^\bullet(c)$. Then by independence, $m_1[\mathcal{H}^\bullet(c)]$ and $m_2[\mathcal{H}^\bullet(c)]$ do not in fact depend on c.

Given the family of subjective models $\{F(\cdot;\theta_1,\theta_2) : \theta_1 \in \Theta_1, \theta_2 \in \Theta_2\}$ about the joint distribution of (X_1, X_2), let $\mathcal{H}(\theta_1,\theta_2;c) := \mathcal{H}(F(\cdot;\theta_1,\theta_2);c)$ denote the distribution on histories under the model $F(\cdot;\theta_1,\theta_2)$ and censoring cutoff c. I now define the method of moments estimator.

Definition A.2. The *method-of-moments (MOM) estimator* derived from an infinite dataset of histories with the distribution $\mathcal{H}^\bullet(c)$ is any pair $(\theta_1^M, \theta_2^M) \in \Theta_1 \times \Theta_2$ such that:

1. $m_1[\mathcal{H}(\theta_1^M, \theta_2^M; c)] = m_1[\mathcal{H}^\bullet(c)]$
2. $m_2[\mathcal{H}(\theta_1^M, \theta_2^M; c)] = m_2[\mathcal{H}^\bullet(c)]$

I will sometimes write $\theta_1^M(c)$, $\theta_2^M(c)$ to emphasize the dependence of the MOM estimators on the censoring threshold c. The MOM estimator need not exist — for example, if all values of $\theta_1 \in \Theta_1$ generate a marginal distribution on X_1 that is smaller than $m_1[\mathcal{H}^\bullet(c)]$. However, when it exists, it is unique under the assumptions I made.

Lemma A.19. *When the family of feasible models satisfies Assumption A.1, the MOM estimator is unique when it exists.*

Now I show the MOM estimators have properties similar to the pseudo-true fundamentals. First, the estimators are monotonic in c, that is MOM agents end up with more pessimistic beliefs about the second-period distribution when the dataset is more severely censored. This a key ingredient for the monotonicity learning dynamics in Theorem 2.

Proposition A.4. *Suppose Assumption A.1 holds. Suppose $c^{'} < c^{''}$ are two different interior values in I_1 and that MOM estimators $(\theta_1^M(c^{'}), \theta_2^M(c^{'}))$ and $(\theta_1^M(c^{''}), \theta_2^M(c^{''}))$ exist. Then $\theta_1^M(c^{'}) = \theta_1^M(c^{''})$ and $\theta_2^M(c^{'}) < \theta_2^M(c^{''})$.*

When (θ_1, θ_2) correspond to the unconditional means, the MOM estimators understate the X_2 mean of the objective distribution $F^\bullet$.

Proposition A.5. *Suppose parameters (θ_1,θ_2) index the unconditional X_1, X_2 means in all feasible models, that is $\mathbb{E}_{F(\cdot;\theta_1,\theta_2)}[X_1] = \theta_1$ and $\mathbb{E}_{F(\cdot;\theta_1,\theta_2)}[X_2] = \theta_2$. Suppose $c \in \mathbb{R}$ and the MOM estimators $\theta_1^M(c), \theta_2^M(c)$ exist. Let $\theta_1^\bullet = \mathbb{E}_{F^\bullet}[X_1]$, $\theta_2^\bullet = \mathbb{E}_{F^\bullet}[X_2]$ be the unconditional means of the true distribution of draws. Then, $\theta_1^M(c) = \theta_1^\bullet$, $\theta_2^M(c) < \theta_2^\bullet$.*

Proof. For any $\theta_1 \in \Theta_1$, $\theta_2 \in \Theta_2$, and $c \in \mathbb{R}$, $m_1[\mathcal{H}(\theta_1, \theta_2; c)] = \theta_1$ since (θ_1, θ_2) are assumed to parametrize means. Also, $m_2[\mathcal{H}(\theta_1, \theta_2; c)] = \mathbb{E}_{F(\cdot;\theta_1,\theta_2)}[X_2 \mid X_1 \leq c] > \theta_2 = \mathbb{E}_{F(\cdot;\theta_1,\theta_2)}[X_2]$ due to Assumption A.1(c). Finally, and $m_1[\mathcal{H}^\bullet(c)] = \theta_1^\bullet$, $m_2[\mathcal{H}^\bullet(c)] = \theta_2^\bullet$ due to independence in $F^\bullet$.

This means if θ_1^M, θ_2^M are the MOM estimators under censoring threshold c, then $\theta_1^M = \theta_1^\bullet$. Also, we must have $m_2[\mathcal{H}(\theta_1^M, \theta_2^M; c)] = \theta_2^\bullet$. At the same time we have $m_2[\mathcal{H}(\theta_1^M, \theta_2^M; c)] > \theta_2^M$, so this means $\theta_2^M < \theta_2^\bullet$. □

These conclusions show that the ideas behind my main results do not depend on the Gaussian assumption or on full Bayesianism. Rather, the crucial assumption is the generalized notion of negative dependence between X_1 and X_2, as articulated by Assumption A.1(c) for arbitrary joint distributions.

As a corollary, I characterize the large-generations learning dynamics for MOM agents using a general class of feasible models. The key idea is that the positive feedback between distorted stopping rules and distorted beliefs continue to hold, with the parametric version of gambler's fallacy interpreted as $\gamma > 0$ in a specific Gaussian setup replaced with the general notion of negative dependence as in Assumption A.1(c).

To define the MOM estimators for agents who observe several sub-datasets of histories with different censoring thresholds, we extend the moment functions m_1, m_2 to take as argument multiple history distributions. That is, $m_1(\mathcal{H}^{(1)}, ..., \mathcal{H}^{(K)}) := \mathbb{E}_{h \sim \oplus_{k=1}^K \mathcal{H}^{(k)}}[h_1]$ and $m_2(\mathcal{H}^{(1)}, ..., \mathcal{H}^{(K)}) := \mathbb{E}_{h \sim \oplus_{k=1}^K \mathcal{H}^{(k)}}[h_2 \mid h_2 \neq \varnothing]$, where $\oplus_{k=1}^K \mathcal{H}^{(k)}$ is the mixture distribution assigning $\frac{1}{K}$ weight to each of the k history distributions, $(\mathcal{H}^{(k)})_{k=1}^K$. After K sub-datasets of censored histories with distributions $\mathcal{H}^\bullet(c_{[0]}, ..., c_{[K-1]})$, the MOM estimators $\mu_1^M(c_{[0]}, ..., c_{[K-1]})$, $\mu_2^M(c_{[0]}, ..., c_{[K-1]})$ are such that

$$m_1(\mathcal{H}^\bullet(c_{[0]}), ..., \mathcal{H}^\bullet(c_{[K-1]})) = m_1(\mathcal{H}(\theta_1^M, \theta_2^M; c_{[0]}), ..., \mathcal{H}(\theta_1^M, \theta_2^M; c_{[K-1]}))$$

$$m_2(\mathcal{H}^\bullet(c_{[0]}), ..., \mathcal{H}^\bullet(c_{[K-1]})) = m_2(\mathcal{H}(\theta_1^M, \theta_2^M; c_{[0]}), ..., \mathcal{H}(\theta_1^M, \theta_2^M; c_{[K-1]})).$$

One caveat: we must now ensure the MOM estimator exists in each generation when the previous generation uses any cutoff stopping rule that has a positive probability of continuing into the next period. To guarantee existence, I impose an additional assumption on how the feasible models relates to the true distribution $F^\bullet$.

Assumption A.2. *(a) The supports of X_1 and X_2 under the true distribution $F^\bullet$ are I_1 and I_2, respectively.*

(b) The range of $\theta_1 \mapsto \mathbb{E}_{F_1(\cdot;\theta_1)}[X_1]$ is I_1.

(c) *For every $\theta_1 \in \Theta_1$ and $x_1 \in I_1$, $\theta_2 \mapsto \mathbb{E}_{F_{2|1}(\cdot;\theta_1,\theta_2|x_1)}[X_2]$ is continuous with a range of I_2.*

Assumption A.2(a) is a consistency requirement, saying that the supports for the objective distributions of X_1 and X_2 match their supports under the agents' subjective models. Assumption A.2(b) and Assumption A.2(c) ensures the agents can always match the two moment conditions. It is easily verified that Examples A.1, A.2, and A.3 satisfy Assumption A.2 when the true joint distribution of (X_1, X_2) is supported on $\mathbb{R}^2$, $[0,\infty)^2$, and $[0,1]^2$ respectively.

Corollary A.1. *Fix some objective, independent distribution $F^\bullet$ for (X_1, X_2) and suppose agents' feasible models $\{F(\cdot;\theta_1,\theta_2) : \theta_1 \in \Theta_1, \theta_2 \in \Theta_2\}$ satisfy Assumptions A.1 and A.2. Suppose the payoff function $u_2(x_1, x_2)$ in the optimal-stopping problem is linear in x_2. Initiate the 0th generation at an arbitrary cutoff $c_{[0]}$ in the interior of I_1. Then, beliefs and cutoff thresholds $(\mu^M_{1,[t]})_{t\geq 1}$, $(\mu^M_{2,[t]})_{t\geq 1}$, and $(c^M_{[t]})_{t\geq 1}$ form monotonic sequences.*

This corollary establishes the monotonicity of the beliefs and cutoffs for MOM agents, analogous to the monotonicity result of Theorem 2.

D Optimal-Stopping Problems with L Periods

D.1 An L-Periods Model of the Gambler's Fallacy

In an optimal-stopping problem with L periods, the agent observes a draw $x_\ell \in \mathbb{R}$ in each period $1 \leq \ell \leq L$. At the end of period ℓ, the agent must decide between stopping and receiving a payoff $u_\ell(x_1, ..., x_\ell)$ that depends on the profile of draws $(x_i)_{i=1}^{\ell}$ observed so far, or continuing into the next period. If the agent continues into period L without stopping, then his payoff will be $u_L(x_1, ..., x_L)$.

I first introduce notation for a class of joint distributions of the L possible draws $(X_i)_{i=1}^L$.

Definition A.3. Let $\sigma^2 > 0$ be fixed. For every vector $\boldsymbol{\mu} = (\mu_i)_{i=1}^L$ and triangular array $\boldsymbol{\gamma} = (\gamma_{i,j})_{2\leq i\leq L, 1\leq j\leq i-1}$ with each $\gamma_{i,j} \in \mathbb{R}$, the *subjective model* $\Psi(\boldsymbol{\mu};\boldsymbol{\gamma})$ denotes the joint distribution of $(X_i)_{i=1}^L$ where $X_1 \sim \mathcal{N}(\mu_1, \sigma^2)$ and, for all $i \geq 2$ and $(x_j)_{j=1}^{i-1} \in \mathbb{R}^{i-1}$,

$$X_i|(X_1 = x_1, ..., X_{i-1} = x_{i-1}) \sim \mathcal{N}(\mu_i - \sum_{j=1}^{i-1} \gamma_{i,j} \cdot (x_j - \mu_j), \sigma^2).$$

Under $\Psi(\boldsymbol{\mu};\boldsymbol{\gamma})$, $(X_i)_{i=1}^L$ are jointly Gaussian,[24] such that the conditional mean of X_i given

[24] An equivalent description of the subjective model $\Psi(\boldsymbol{\mu};\boldsymbol{\gamma})$ is to consider a set of L independent Gaussian random variables $Z_i \sim \mathcal{N}(\mu_i, \sigma^2)$ for $1 \leq i \leq L$. Let $X_1 = Z_1$ and iteratively define $X_i = Z_i - \sum_{j=1}^{i-1} \gamma_{i,j}(X_j - \mu_j)$. Using induction, one can show that every X_i is a linear function of the Z_i's, so they are jointly Gaussian.

the previous draws $X_1 = x_1, ..., X_{i-1} = x_{i-1}$ depends linearly on these realizations. I consider agents who entertain a set of subjective models, $\{\Psi(\boldsymbol{\mu};\boldsymbol{\gamma}) : \boldsymbol{\mu} \in \mathbb{R}^L\}$ for a fixed array $\boldsymbol{\gamma}$ where each $\gamma_{i,j} > 0$. The positive $\gamma_{i,j}$ capture the gambler's fallacy, as higher realizations of earlier draws lead agents to predict lower means for future draws. The greater the magnitude of $\gamma_{i,j}$, the more that the agent's prediction of X_i depends on realization of X_j. Agents hold a dogmatic belief in the correlation structure between $(X_i)_{i=1}^L$, but can flexibly estimate $(\mu_i)_{i=1}^L$, the *fundamentals* of the environment. Objectively, $(X_i)_{i=1}^L$ are independent, so the true joint distribution is $\Psi^\bullet = \Psi(\boldsymbol{\mu}^\bullet;\mathbf{0})$ for some $(\mu_i^\bullet)_{i=1}^L$.

A useful functional form to keep in mind is $\gamma_{i,j} = \alpha \cdot \delta^{i-j-1}$ for $\alpha > 0$, $0 \le \delta \le 1$, which corresponds to Rabin and Vayanos (2010)'s specification of gambler's fallacy in multiple periods. Here, α relates to the severity of the bias and δ captures how quickly the influence of past observations decay in predicting future draws.

D.2 Inference from Censored Datasets in L Periods

In general, a stopping strategy in an optimal-stopping problem over L periods is a set of functions $S_i : \mathbb{R}^i \to \{\text{Stop}, \text{Continue}\}$ for $1 \le i \le L-1$, where $S_i(x_1, ..., x_i)$ maps the realizations of the first i draws to a stopping decision. I consider stopping strategies where S_i is a cutoff rule in x_i after each partial history $(x_1, ..., x_{i-1})$, that is there exist $(c_i)_{i=1}^{L-1}$ with $c_1 \in \mathbb{R}$ and for $i \ge 2$, $c_i(x_1, ..., x_{i-1}) \in \mathbb{R}$ for every $(x_1, ..., x_{i-1}) \in \mathbb{R}^{i-1}$, so that the agent stops after $(x_1, ...x_i)$ if and only if $x_i \ge c_i(x_1, ..., x_{i-1})$. A stopping strategy with stopping regions characterized by a profile of cutoff rules $\boldsymbol{c} = (c_i)_{i-1}^{L-1}$ will be abbreviated as $S_{\boldsymbol{c}}$.

For subjective model Ψ and cutoff rule $S_{\boldsymbol{c}}$, let $\mathcal{H}(\Psi; S_{\boldsymbol{c}})$ represent the distribution of histories when applying rule $S_{\boldsymbol{c}}$ to draws $(X_i) \sim \Psi$. More precisely, consider a procedure where $X_1, X_2, ..., X_L$ is drawn according to Ψ and revealed one at a time. At the earliest $1 \le \bar{i} \le L-1$ such that $X_{\bar{i}} \ge c_{\bar{i}}(X_1, ..., X_{\bar{i}-1})$, the process stops and the history records $(X_1, ..., X_{\bar{i}}, \varnothing, ..., \varnothing)$, with $L - \bar{i}$ instances of the censoring indicator $\varnothing$ replacing the unobserved subvector $(X_{\bar{i}+1}, ..., X_L)$. If no such $\bar{i}$ exists, then history records the entire profile of draws, $(X_1, ..., X_L)$. The distribution of histories generated this way is denoted $\mathcal{H}(\Psi; S_{\boldsymbol{c}})$.

Definition A.4. For cutoff strategy $S_{\boldsymbol{c}}$ and fundamentals $\hat{\boldsymbol{\mu}}$, the *KL divergence* between objective distribution of histories and the predicted distribution under censoring is the sum of L integrals,

$$D_{KL}(\ \mathcal{H}(\Psi^\bullet; S_{\boldsymbol{c}}) \ || \ \mathcal{H}(\Psi(\boldsymbol{\mu};\boldsymbol{\gamma}); S_{\boldsymbol{c}})\ ) := \sum_{i=1}^{L} I_i,$$

where

$$I_1 = \int_{c_1}^{\infty} \phi(x_1; \mu_1^\bullet, \sigma^2) \ln\left(\frac{\phi(x_1; \mu_1^\bullet, \sigma^2)}{\phi(x_1; \mu_1, \sigma^2)}\right) dx_1,$$

and for $2 \le i \le L-1$, integral I_i is

$$\int_{-\infty}^{c_1} \cdots \int_{-\infty}^{c_{i-1}(x_1,...,x_{i-2})} \int_{c_i(x_1,...,x_{i-1})}^{\infty} \prod_{k=1}^{i} \phi(x_k; \mu_k^\bullet, \sigma^2) \ln \left(\frac{\prod_{k=1}^{i} \phi(x_k; \mu_k^\bullet, \sigma^2)}{\prod_{k=1}^{i} \phi(x_k; \mu_k - \sum_{j=1}^{k-1} \gamma_{k,j} \cdot (x_j - \mu_j), \sigma^2)} \right) dx_i ... dx_1.$$

Finally, I_L is given by

$$\int_{-\infty}^{c_1} \cdots \int_{-\infty}^{c_{L-1}(x_1,...,x_{L-2})} \int_{-\infty}^{\infty} \cdot \prod_{k=1}^{i} \phi(x_k; \mu_k^\bullet, \sigma^2) \ln \left(\frac{\prod_{k=1}^{i} \phi(x_k; \mu_k^\bullet, \sigma^2)}{\prod_{k=1}^{i} \phi(x_k; \mu_k - \sum_{j=1}^{k-1} \gamma_{k,j} \cdot (x_j - \mu_j), \sigma^2)} \right) dx_i ... dx_1.$$

To interpret, consider a history $h = (x_1, ..., x_i, \varnothing, ..., \varnothing)$ where $x_k < c_k(x_1, ..., x_{k-1})$ for all $k \le i-1$ and $x_i \ge c_i(x_1, ..., x_{i-1})$. This history is possible under the stopping strategy $S_{\boldsymbol{c}}$. It has a likelihood of $\Pi_{k=1}^{i} \phi(x_k; \mu_k^\bullet, \sigma^2)$ under $\Psi^\bullet$ and a likelihood of $\Pi_{k=1}^{i} \phi(x_k; \mu_k - \sum_{j=1}^{k-1} \gamma_{k,j} \cdot (x_j - \mu_j), \sigma^2)$ under $\Psi(\boldsymbol{\mu}; \boldsymbol{\gamma})$. So, the integral I_i calculates the contribution of all possible histories of length i to the KL divergence from $\mathcal{H}(\Psi(\boldsymbol{\mu}; \boldsymbol{\gamma}); S_{\boldsymbol{c}})$ to $\mathcal{H}(\Psi^\bullet; S_{\boldsymbol{c}})$. In the case of $L = 2$, this definition reduces to Definition 5, the KL divergence in the two-periods baseline model, with $\gamma = \gamma_{2,1}$ and $c_1 \in \mathbb{R}$ as the censoring threshold.

The KL-divergence minimizers

$$\min_{\boldsymbol{\mu} \in \mathbb{R}^L} D_{KL}(\ \mathcal{H}(\Psi^\bullet; S_{\boldsymbol{c}}) \ || \ \mathcal{H}(\Psi(\boldsymbol{\mu}; \boldsymbol{\gamma}); S_{\boldsymbol{c}})\ )$$

are the *pseudo-true fundamentals* with respect to stopping strategy $S_{\boldsymbol{c}}$. The next proposition gives an explicit characterization of them.

Proposition A.6. *Let stopping strategy $S_{\boldsymbol{c}}$ be given. For each $i \ge 1$, let R_i represent the region*

$$\{(x_1, ..., x_i) : x_1 < c_1, x_2 < c_2(x_1), ..., x_i < c_i(x_1, .., x_{i-1})\} \subseteq \mathbb{R}^i.$$

The pseudo-true fundamentals with respect to $S_{\boldsymbol{c}}$ are $\mu_1^* = \mu_1^\bullet$ and, iteratively,

$$\hat{\mu}_i^* = \mu_i^\bullet - \sum_{j=1}^{i-1} \gamma_{i,j} \cdot (\mu_j^* - \mathbb{E}_{\Psi^\bullet}[X_j | (X_k)_{k=1}^{i-1} \in R_{i-1}]).$$

The expression for μ_i^* in the general L-periods setting resembles the expression for μ_2^* in the two-period setting. Relative to the truth $\mu_i^\bullet$, the estimate μ_i^* is distorted by the fact that X_i is only observed when previous draws $(X_1, ..., X_{i-1})$ fall into the continuation region $R_{i-1} \subseteq \mathbb{R}^{i-1}$ associated with $S_{\boldsymbol{c}}$. The agent uses this censored empirical distribution of $(X_1, ..., X_{i-1}, X_i)$ to infer the period-i fundamental, under a dogmatic belief about the correlation structure between the draws given by $\boldsymbol{\gamma}$. Importantly, whether a certain realization X_j for $j < i$ should be judged as below-average (and thus predict a higher X_i) or above-average (and thus predict a lower X_i) depends on agent's belief about the period j fundamental, μ_j^*, which gives the iterative structure of the expression for $\hat{\mu}_i^*$.

The proof of this result follows two steps. First, recall that $D_{KL}(\mathcal{H}(\Psi^\bullet; S_{\boldsymbol{c}})||\mathcal{H}(\Psi(\boldsymbol{\mu}; \boldsymbol{\gamma}); S_{\boldsymbol{c}}))$ is defined as the sum $\sum_{i=1}^{L} I_i$, where I_i is the KL-divergence contribution from histories with length i. I rewrite this expression as the sum of L different integrals, $\sum_{i=1}^{L} J_i$, where J_i is the KL-divergence contributions from histories containing X_i. So, J_i is a function of $\mu_1, ..., \mu_i$. The second step is similar to the proof of Proposition 2, where I show $\frac{\partial J_i}{\partial \mu_j}$ is a linear multiple of $\frac{\partial J_i}{\partial \mu_i}$ whenever $j < i$. First-order condition at $\boldsymbol{\mu}^*$ allows for a telescoping rearrangement, yielding $\frac{\partial J_i}{\partial \mu_i}(\boldsymbol{\mu}^*) = 0$ for every i. The proposition readily follows.

Now I turn to a special class of cutoff-based stopping rules where c_k is independent of history. So, a stopping rule of this kind $S_{\boldsymbol{c}}$ can be viewed simply as a list of L constants, $c_1, ..., c_L \in \mathbb{R}$, such that the agent stops after the draw $X_\ell = x_\ell$ if and only if $x_\ell < c_\ell$. I show that the expression for the pseudo-true fundamentals greatly simplifies and admits a path-counting interpretation.

Definition A.5. For $1 \leq j < i \leq L$, a *path* p from i to j is a sequence of pairs $p = ((i_0, i_1), ..., (i_{M-1}, i_M))$ with $M \geq 1$, $i_0 = i$, $i_M = j$, and $i_{m+1} < i_m$ for all $m = 0, 1, ..., M-1$. The length of p is $\#(p) := M$. The weight of p is $W(p) := \Pi_{0 \leq m \leq M-1}(-\gamma_{i_\ell, i_{\ell+1}})$. Denote the set of all paths from i to j as $P[i \to j]$.

That is, we may imagine a network with L nodes, one per period of the optimal-stopping problem. There is a directed edge with weight $-\gamma_{i,j}$ for all pairs $i > j$. A path from i to j is a concatenation of edges, starting with i and ending with j. Its weight is the product of the weights of all the edges used.

The next proposition differs from Proposition A.6 in that the expression for the pseudo-true fundamental μ_i^* does not involve other pseudo-true fundamentals μ_j^*. It shows that the distortion of μ_i^* from the true value $\mu_i^\bullet$ depends on terms $\mu_j^\bullet - \mathbb{E}_{\Psi^\bullet}[X_j | X_j \leq c_j]$ and the total number of paths from i to j in the network that $\boldsymbol{\gamma}$ defines.

Proposition A.7. *For stopping strategy $S_{\boldsymbol{c}} = (c_1, ..., c_L) \in \mathbb{R}^L$, the pseudo-true fundamentals are given by*

$$\mu_i^* = \mu_i^\bullet + \sum_{j=1}^{i-1} \left(\sum_{p \in P[i \to j]} W(p) \right) \cdot \left(\mu_j^\bullet - \mathbb{E}[X_j | X_j \leq c_j] \right).$$

As a corollary, suppose $L \geq 3$ and $\boldsymbol{\gamma}$ have the Rabin and Vayanos (2010) functional form of $\gamma_{i,j} = \alpha \cdot \delta^{i-j-1}$ for $\alpha > 0$, $0 \leq \delta \leq 1$. I show that all pseudo-true fundamentals are too pessimistic in every dataset censored with $S_{\boldsymbol{c}} = (c_1, ..., c_L) \in \mathbb{R}^L$ if and only if $\delta > \alpha$. The idea is the influence of the gambler's fallacy psychology must not decay "too quickly" relative to the influence of the most recent observation. This condition is satisfied in all the calibration exercises in Rabin and Vayanos (2010) and in the structural estimations of Benjamin, Moore, and Rabin (2017). The result shows the over-pessimism from the 2-periods model extends into the L periods model (provided the regularity condition on the parametrization of the L-periods gambler's fallacy holds).

Corollary A.2. *Suppose $L \geq 3$ and $\gamma_{i,j} = \alpha \cdot \delta^{i-j-1}$ for $\alpha > 0$, $0 \leq \delta \leq 1$. If $\delta > \alpha$, then for all stopping strategies $S_{\boldsymbol{c}} = (c_1, ..., c_L) \in \mathbb{R}^L$, the pseudo-true fundamentals satisfy $\mu_i^* < \mu_i^\bullet$ for all i. If $\delta < \alpha$, then there exists a stopping strategy $S_{\boldsymbol{c}} = (c_1, ..., c_L) \in \mathbb{R}^L$ such that $\mu_i^* > \mu_i^\bullet$ for at least one i.*

To understand the intuition, consider an example that violates the condition of the corollary, $\alpha = 0.5$, $\delta = 0$, so that $\gamma_{2,1} = 0.5$, $\gamma_{3,2} = 0.5$, and $\gamma_{3,1} = 0$. The agent expects reversals between the pairs (X_1, X_2) and (X_2, X_3), but his expectation for $X_3|(X_1 = x_1, X_2 = x_2)$ does not vary with x_1. By the same logic as the two-periods censoring effect, inference about the second-period fundamental μ_2^* decreases as c_1 decreases, with $\lim_{c_1 \to -\infty} \mu_2^*(c_1) = -\infty$. This has an important indirect effect on μ_3^*, since a very pessimistic μ_2^* leads the agent to interpret objectively typical draws of X_2 as greatly above average. Expecting low values of X_3 after these surprisingly high draws of X_2, the agent infers the fundamental μ_3^* to be above the sample mean of X_3 in the dataset, hence overestimating it as $c_1 \to -\infty$. When δ is strictly positive, however, there is an opposite effect where lower sample mean of X_1 in observations containing uncensored X_3 lead to more pessimistic inference about the third-period fundamental. When $\delta > 0.5$, overoptimistic inference never happens because this second effect dominates.

E The Censoring Effect in a Finite-Urn Model

Rabin (2002) Section 7 discusses an example with endogenous observations. There is an infinite population of financial analysts, each with quality $\theta \in \{\frac{1}{4}, \frac{1}{2}, \frac{3}{4}\}$. Conditional on quality θ, an analyst generates either a good (signal a) or bad (signal b) return each period, with probabilities θ and $1-\theta$ and independently across periods. The agent, however, believes successive returns from the same analyst are generated through a finite-urn model. Consider an urn with N balls where N is a multiple of 4. For an analyst with quality θ, initialize the urn with θN balls labeled "a" and $(1-\theta)N$ balls labeled "b". Successive returns are successive draws without replacement from the urn. The urn is refreshed every two draws. Rabin (2002) calls an agent with this finite-urn model an "N-Freddy". Since the urn is not refreshed between draws $2k-1$ and $2k$ for $k = 1, 2, 3, ...$, such pairs of draw exhibit negative correlation in agent's subjective model, generating the gambler's fallacy.

Returning to Rabin (2002) Section 7's example, objectively all financial analysts have quality $\theta = \frac{1}{2}$. The agent samples a financial analyst at random and observes his returns over two periods. Depending on the realizations of these two returns, the agent either observes the same analyst for two more periods before sampling a new analyst, or immediately samples a new analyst. This procedure is infinitely repeated. Rabin (2002) investigates a 4-Freddy agent's long-run belief about the proportions of analysts with the three levels of quality in the population.

The endogenous observation in the example is distinct from what I term the "censoring effect" in this paper. The main mechanism behind the censoring effect is that the some rows of the dataset omits signals (X_2) which the biased agent judges to be negatively correlated with signals that are always observed (X_1). This then leads to distorted inference. However, in Rabin (2002)'s finite-urn model, the urn is refreshed every two periods. This means an N-Freddy agent judges the part of the data that is sometimes censored (the analyst's returns in periods 3 and 4) to be independent of the part of the data that is always observed (the analyst's returns in periods 1 and 2). Therefore the driving force behind Rabin (2002) Section 7's example is not the interaction between censoring and the gambler's fallacy, but rather between censoring and the Bayesian aspect of N-Freddy's quasi-Bayesian inference.

In this section, I study a related problem where an N-Freddy agent observes each analyst for either one or two periods, depending on whether the analyst generates a bad first-period return. This setup features the censoring effect, because the finite-urn model generates negative correlation between the first and second draws from each urn. I find that the agent's inference under this censoring structure tends to be too optimistic. This conclusion is in line with predictions about the censoring effect in the baseline model of this paper, for the basic inference result in Proposition 2 shows that when the dataset is censored in the opposite way (i.e. censored when the first draw is good), the resulting inference is too pessimistic[25]. That is, I demonstrate the robustness of my censoring effect to an alternative model of the gambler's fallacy in a binary-signals setting, showing that it is not an artifact of the continuous-signals setup in my baseline model.

Table A.1 displays the likelihood of all signals of length 2 for the 4-Freddy and 8-Freddy agents, for different values of $\theta \in \{\frac{1}{4}, \frac{1}{2}, \frac{3}{4}\}$. The last row of each table also shows the likelihoods of simply observing the signal b in the first period, under the censoring rule that stops observing an analyst if his first return is bad. I first discuss inference without censoring. After aa, Freddy exaggerates the relative likelihood of $\theta = \frac{3}{4}$ to $\theta = \frac{1}{2}$ compared to a Bayesian, whereas after ab Freddy's relative likelihoods of these two qualities are the same as a Bayesian's. Overall, given a sample with an equal number of aa and ab signals,

[25]Proposition OA.12 in the Online Appendix shows that when the dataset is censoring using a strategy that stops when $X_1 \leq c$ for some $c \in \mathbb{R}$, inference about second-period fundamental is always too high.

4-Freddy	$\theta = \frac{1}{4}$	$\theta = \frac{1}{2}$	$\theta = \frac{3}{4}$
aa	0	$\frac{1}{6}$	$\frac{1}{2}$
ab	$\frac{1}{4}$	$\frac{1}{3}$	$\frac{1}{4}$
ba	$\frac{1}{4}$	$\frac{1}{3}$	$\frac{1}{4}$
bb	$\frac{1}{2}$	$\frac{1}{6}$	0
$b\varnothing$	$\frac{3}{4}$	$\frac{1}{2}$	$\frac{1}{4}$

8-Freddy	$\theta = \frac{1}{4}$	$\theta = \frac{1}{2}$	$\theta = \frac{3}{4}$
aa	$\frac{1}{28}$	$\frac{6}{28}$	$\frac{15}{28}$
ab	$\frac{6}{28}$	$\frac{8}{28}$	$\frac{6}{28}$
ba	$\frac{6}{28}$	$\frac{8}{28}$	$\frac{6}{28}$
bb	$\frac{15}{28}$	$\frac{6}{28}$	$\frac{1}{28}$
$b\varnothing$	$\frac{3}{4}$	$\frac{1}{2}$	$\frac{1}{4}$

Table A.1: The likelihoods of observations under different analyst qualities, for 4-Freddy and 8-Freddy agents.

Freddy exaggerates the relative likelihood of $\theta = \frac{3}{4}$ to $\theta = \frac{1}{2}$. This phenomenon is analogous to the continuous version of gambler's fallacy where a biased observer "partially forgives" a mediocre outcome following an outstanding outcome. Here, even though the average outcome in the second period is mediocre, the fact that they follow the best possible outcome a in the first period lead to an overly optimistic estimate about the analyst's ability. By the same logic, observing an equal number of ba and bb signals would lead to exaggeration of the likelihood of $\theta = \frac{1}{4}$ relative to $\theta = \frac{1}{2}$.

However, now suppose the second observation is censored when the first observation is b. The otherwise symmetric situation becomes asymmetric. Following the observation of $b\varnothing$ (where the second draw is censored), Freddy's inference is the same as a Bayesian's. So we have turned off the channel that exaggerates the probability of $\theta = \frac{1}{4}$ but kept the channel that exaggerates the probability of $\theta = \frac{3}{4}$. This is analogous to the censoring effect in my model, where censoring second period draw following unfavorable first period draws would lead to overly optimistic beliefs.

In the long-run, the agent observes a distribution of returns across different analysts: 25% of the time aa is observed, 25% of the time ab is observed, and 50% of the time $b\varnothing$ is observed. To calculate the agent's long-run beliefs, first suppose Freddy's prior specifies either all analysts have $\theta = \frac{1}{4}$ or all analysts have $\theta = \frac{3}{4}$. Then Freddy's long-run inference is given by the parameter maximizing expected log-likelihood of the data. For 4-Freddy, the log-likelihood likelihood under $\theta = \frac{1}{4}$ is $-\infty$. For 8-Freddy, The log-likelihood under $\theta = \frac{1}{4}$ is

$$\frac{1}{4}\ln(1/28) + \frac{1}{4}\ln(6/28) + \frac{1}{2}\ln(3/4) \approx -1.362$$

and the log-likelihood under $\theta = \frac{3}{4}$ is

$$\frac{1}{4}\ln(\frac{15}{28}) + \frac{1}{4}\ln(\frac{6}{28}) + \frac{1}{2}\ln(1/4) \approx -1.234.$$

So in both cases, Freddy will come to believe $\theta = \frac{3}{4}$ over $\theta = \frac{1}{4}$ for all analysts.

Now consider a 4-Freddy who dogmatically believes some $1 - \kappa \in (0, 1)$ fraction of the analysts have $\theta = \frac{1}{2}$, but the remaining analysts either have $\theta = \frac{1}{4}$ or $\theta = \frac{3}{4}$. So, the agent estimates $q_a \in [0, 1 - \kappa]$, the fraction of analysts who have $\theta = \frac{3}{4}$. Straightforward algebra shows that the q_a^* maximizing expected log-likelihood of the data is $q_a^* = \frac{7}{18}\kappa + \frac{1}{9}$ for $\kappa \geq \frac{2}{11}$, $q_a^* = \kappa$ otherwise. Since $\frac{7}{18}\kappa + \frac{1}{9} > \frac{1}{2}\kappa$ for all $\kappa \in (\frac{2}{11}, 1)$, we see that no matter what fraction of analysts 4-Freddy believes to be average, he will end up believing there are more above-average than below-average analysts in the population. That is, his overall belief will be too optimistic.

F The Gambler's Fallacy and Attentional Stability

Many investigations of behavioral learning, including this paper, can be phrased as agents with a prior (or "misspecified theory") over states of the world whose support excludes the true, data-generating state. Agents in this paper start with a prior supported on the class of subjective models $\{\Psi(\mu_1, \mu_2, \sigma^2, \sigma^2; \gamma) : \mu_1, \mu_2 \in \mathbb{R}\}$ for some fixed $\gamma > 0$, with different models viewed as different states of the world. But the true state of the world is the objective distribution $(X_1, X_2) \sim \Psi(\mu_1^\bullet, \mu_2^\bullet, \sigma^2, \sigma^2; 0)$, which does not belong to the feasible set. As an agent's data size grows, her misspecified theory can appear infinitely less likely in the limit than an alternative prior belief (or "light-bulb theory") that includes the true state in its support.

Gagnon-Bartsch, Rabin, and Schwartzstein (2018) offer an explanation for why such misspecified theories persist with learning – attentional stability. Under a misspecified theory, some coarsened information may be sufficient for decision-making. When agents only pay attention to this coarsened information, the aspects of the data that they attend to may be so coarse that their misspecified theory no longer appears infinitely less likely than the light-bulb theory.

In this section, I investigate the attentional stability of the gambler's fallacy bias in my learning setting. The main intuition is that when agents are dogmatic about γ, they are dogmatic about the correlation between X_1 and X_2. Therefore, under their misspecified theory, agents do not find it necessary to separately keep track of the conditional distributions $X_2|(X_1 = x_1)$ for different values of x_1. Agents believe certain "statistics" of the dataset are sufficient for decision-making, and this process of reducing the entire dataset into these sufficient statistics removes features of the dataset that would otherwise have led the agents to question the validity of their theory.

My setting differs in two ways from that of Gagnon-Bartsch, Rabin, and Schwartzstein (2018). Each of my agents acts once (after observing a possibly large or even infinite dataset), while their agents observe one signal each period over an infinite number of periods. Another distinction is that data is endogenous in my setting, whereas Gagnon-Bartsch, Rabin, and Schwartzstein (2018) almost entirely focus on an exogenous-data environment. So, I begin by defining the key concepts surrounding attentional stability in my setting.

F.1 A Definition of Attentional Stability in Large Datasets

In the learning environment where agents act one at a time, Theorem 1 implies that almost surely behavior converges. As a consequence, late enough agents observe finite datasets of histories whose distributions asymptotically resemble $\mathcal{H}^\bullet(c^\infty)$, where c^∞ is the steady-state censoring threshold.

In the learning environment where agents act in large generations, each agent in generation t observes t sub-datasets of infinitely many histories. The overall distribution of histories in the dataset is $\mathcal{H}^\bullet(c_{[0]}, ..., c_{[t-1]}) = \oplus_{k=0}^{t-1}\mathcal{H}^\bullet(c_{[k]})$, where the right-hand side refers to the mixture between the t history distributions that assigns weight $1/t$ to each.

To develop a definition of attentional stability in large datasets, I consider an agent who directly observes a distribution of histories (instead of a dataset with this distribution) $\mathcal{H}^\bullet(c_1, ..., c_L) \in \Delta(\mathbb{H})$. This corresponds to the asymptotic long-run observation in the environment where agents act one at a time, by putting $L = 1$ and $c = c^\infty$. It also represents the medium-run observations of agents in each generation $t \geq 1$ of the large-generations environment.

Definition A.6. Let π, λ be beliefs over the joint distribution of (X_1, X_2). Say π is *inexplicable* relative to λ, conditional on the true model $\Psi^\bullet$ and censoring thresholds $c_1, ..., c_L$, if $\mathcal{H}^\bullet(c_1, ..., c_L) = \mathcal{H}(\Psi; c_1, ..., c_L)$ for some $\Psi \in \text{supp}(\lambda)$, but $\mathcal{H}^\bullet(c_1, ..., c_L) \neq \mathcal{H}(\Psi; c_1, ..., c_L)$ for any $\Psi \in \text{supp}(\lambda)$.

Each subjective model Ψ and list of censoring thresholds $c_1, ..., c_L$ together induce a distribution over histories. If the observed history distribution $\mathcal{H}^\bullet(c_1, ..., c_L)$ can be produced by some subjective model of (X_1, X_2) in the support of the light-bulb theory λ, but not by any distribution in the support of the misspecified theory π, then I call π inexplicable.

I now define a particular kind of limited attention. Given a distribution over histories, the agent maps the entire distribution to finitely many real numbers. This is an extreme form of data coarsening. If there is a strategy optimal under the misspecified theory π that only makes use of these finitely many statistics, then we have a sufficient-statistics strategy.

Definition A.7. A *sufficient-statistics strategy (SSS) for large generations* consists of a statistics map $\Lambda : \Delta(\mathbb{H}) \to \mathbb{R}^K$ for some finite $K < \infty$ and a cutoff map $\sigma : \text{Im}(\Lambda) \to \mathbb{R}$, such that agents in each generation $t \geq 1$ of the large-generations environment find it optimal (under the prior $\Psi \sim \pi$) to use the stopping strategy with cutoff $\sigma(\Lambda(\mathcal{H}))$ whenever $\mathcal{H}$ is a dataset of predecessors' histories $\mathcal{H} = \mathcal{H}^{\bullet}(c_{[0]}, ..., c_{[t-1]})$.

An agent following the strategy (Λ, σ) first extracts K statistics (i.e. K real numbers) from the infinite dataset of predecessors' histories. Then, she applies σ to choose a cutoff threshold that only depends on the dataset through its K extracted statistics, $\Lambda(\mathcal{H})$. The idea is that the agent only pays attention to the finitely many statistics, a perhaps more realistic behavior than paying full attention to the entire infinite dataset. If such a strategy is optimal for an agent believing the true joint distribution of (X_1, X_2) is drawn according to her (misspecified) prior $\Psi \sim \pi$, I call the pair (Λ, σ) an SSS.

A related definition of sufficiency works with finite datasets instead of infinite datasets. This corresponds to limited attention for agents in the environment where they act one at a time.

Definition A.8. A *sufficient-statistics strategy (SSS) in datasets of size* $N < \infty$ consists of a statistics map $\Lambda^{(N)} : \mathbb{H}^N \to \mathbb{R}^K$ for some finite $K < \infty$ and a cutoff map $\sigma^{(N)} : \text{Im}(\Lambda^{(N)}) \times \mathbb{N} \to \mathbb{R}$, such that the subjectively optimal cutoff threshold (under the Bayesian posterior belief about the fundamentals after updating prior density $g(\mu_1, \mu_2)$) is $\sigma^{(N)}(\Lambda^{(N)}((h_n)_{n=1}^N), N_1)$ after observing a dataset $(h_n)_{n=1}^N$ with size N and containing $N_1 \leq N$ instances of second-period draws.

Finally, I combine these concepts to define attentional stability. Roughly speaking, the theory π is attentionally stable if we can find a (Λ, σ) pair that pays "fine" enough attention to be an SSS under π, but "coarse" enough attention so that the resulting statistics can be explained by some model in the support of π.

Definition A.9. Theory π is *attentionally stable*, conditional on the objective model $\Psi^{\bullet}$ and censoring thresholds $c_1, ..., c_L$, if there exists an SSS (Λ, σ) such that $\Lambda(\mathcal{H}^{\bullet}(c_1, ..., c_L)) = \Lambda(\mathcal{H}(\Psi; c_1, ..., c_L))$ for some Ψ in the support of π.

F.2 The Gambler's Fallacy is Inexplicable under Full Attention

Fix $\gamma > 0$. Let π be any full-support belief over $\{\Psi(\mu_1, \mu_2, \sigma^2, \sigma^2; \gamma) : (\mu_1, \mu_2) \in \mathcal{M}\}$, where $\mathcal{M} \subseteq \mathbb{R}^2$ is any specification of feasible fundamentals. Let λ be any belief with $\Psi^{\bullet} = \Psi(\mu_1^{\bullet}, \mu_2^{\bullet}, \sigma^2, \sigma^2; 0)$ in its support. I first show that without channeled attention, agents will come to realize that their misspecified theory π is wrong after seeing a large dataset.

Proposition A.8. *π is inexplicable relative to λ, conditional on $\Psi^{\bullet}$ and any censoring thresholds $c_1, ..., c_L \in \mathbb{R}$.*

Proof. This is because $\Psi^\bullet \in \text{supp}(\lambda)$ but every $\Psi \in \text{supp}(\pi)$ has KL divergence bounded away from 0 relative to $\Psi^\bullet$ in terms of the histories they generate under censoring by $c_1, ..., c_L$, that is to say

$$\begin{aligned}
&\inf_{\Psi \in \text{supp}(\pi)} D_{KL}(\mathcal{H}^\bullet(c_1, ..., c_L) \parallel \mathcal{H}(\Psi; c_1, ..., c_L)) \\
=&\inf_{(\mu_1,\mu_2)\in\mathcal{M}} D_{KL}(\mathcal{H}^\bullet(c_1, ..., c_L) \parallel \mathcal{H}(\Psi(\mu_1, \mu_2, \sigma^2, \sigma^2; \gamma); c_1, ..., c_L)) > 0.
\end{aligned}$$

To see why this inequality holds, recall that the proof of Lemma 2 shows the above KL-divergence minimization problem has a minimum strictly above 0 even over the unrestricted domain $(\mu_1, \mu_2) \in \mathbb{R}^2$. The restriction to some $\mathcal{M} \subseteq \mathbb{R}^2$ can only make the minimum larger. □

F.3 The Gambler's Fallacy is Attentionally Stable

Now I exhibit a family of SSS for finite datasets of size N and another SSS for large generations that naturally corresponds to taking $N \to \infty$. These SSS have the additional property that they lead agents to the same beliefs about the fundamentals as the full-attention Bayesianism assumed in the rest of the paper. So, not only do these SSS justify agents not discarding their misspecified theory after seeing large datasets, they also provide a limited-attention foundation for the learning dynamics that I investigate in the main text of the paper.

In a dataset of size N, consider the statistics map with $K = 2$,

$$\Lambda^{(N)}((h_n)_{n=1}^N) = \left(\frac{1}{N}\sum_{n=1}^N h_{1,n},\ \frac{1}{\#(n : h_{2,n} \neq \varnothing)} \sum_{n:h_{2,n}\neq\varnothing} (h_{2,n} + \gamma h_{1,n}) \right).$$

The first statistic is the sample mean of the first-period draws. The second statistic can be thought of as a "re-centered" observation $v_n := h_{2,n} + \gamma h_{1,n}$ for each history h_n where $h_{2,n} \neq \varnothing$. The agent only pays attention to the sample averages of $x_{1,n} = h_{1,n}$ and v_n. Under the subjective model $\Psi(\mu_1, \mu_2; \gamma)$, we may write the distributions of X_1, X_2 as

$$\begin{aligned}
X_1 &= \mu_1 + \epsilon_1 \\
X_2 &= \mu_2 + \gamma\epsilon_1 + z_2
\end{aligned}$$

where $\epsilon_1, z_2 \sim \mathcal{N}(0, \sigma^2)$, are independent. Defining $V := X_2 + \gamma X_1$, we see that under $\Psi(\mu_1, \mu_2; \gamma)$, $V = \mu_2 + \gamma\mu_1 + z_2$. So, observations of first-period draws are signals about μ_1, while observations of re-centered second-period V are signals about $\mu_2 + \gamma\mu_1$.

Proposition A.9. $\Lambda^{(N)}$ *is part of an SSS in datasets of size N. The cutoff choice in this SSS is the same as for the full-attention agent.*

In the environment where full-attention Bayesian agents move one at a time, their behavior is indistinguishable from agents using this SSS. Roughly speaking, this is because the subjective joint distribution between (X_1, V) is Gaussian and the mean of a sequence of Gaussian random variables is a sufficient statistic for the likelihood of the entire sequence. Even when agents are full-attention Bayesians, their posterior distribution only depends on the histories data through these statistics. Therefore, the statistics are sufficient for any decision problem.

Consider now the large-sample analog of the finite-sample SSS just defined. Again with $K = 2$, consider the statistic map Λ sends each distribution $\mathcal{H}$ to $\mathbb{E}_{h \sim \mathcal{H}}[h_{i,1}]$ and $\mathbb{E}_{h \sim \mathcal{H}}[h_{i,2} + \gamma h_{i,1} \mid h_{i,2} \neq \varnothing]$. I show that Λ makes π attentionally explicable whenever π has full-support over the subjective models indexed by feasible fundamentals $\mathcal{M} = \mathbb{R}^2$.

Proposition A.10. *For any list of censoring thresholds $c_1, ..., c_L \in \mathbb{R}$ and fundamentals $\mu_1, \mu_2 \in \mathbb{R}$,*

$$\Lambda_1(\mathcal{H}(\Psi(\mu_1, \mu_2, \gamma); c_1, ..., c_L)) = \mu_1,$$

$$\Lambda_2(\mathcal{H}(\Psi(\mu_1, \mu_2, \gamma); c_1, ..., c_L)) = \mu_2 + \gamma\mu_1.$$

Also,

$$\Lambda(\mathcal{H}^\bullet(c_1, ..., c_L)) = \Lambda(\mathcal{H}(\ \Psi(\mu_1^\bullet, \mu_2^*(c_1, ..., c_L); \gamma)\ ; c_1, ..., c_L))$$

The first two equations in this claim show that for any $c_1, ..., c_L$, $\Psi \mapsto \Lambda(\mathcal{H}(\Psi; c_1, ..., c_L))$ is a one-to-one function on the support of π, and furthermore any values of the statistics s_1, s_2 can be rationalized through appropriate choices of μ_1, μ_2. We may put $\sigma(s_1, s_2) = C(s_1, s_2 - \gamma s_1; \gamma)$ to make (Λ, σ) an SSS, thus showing the gambler's fallacy is attentionally stable in large datasets. Another implication of this claim is that the limited-attention agent comes to believe the large-generations pseudo-true fundamentals $(\mu_1^\bullet, \mu_2^*(c_1, ..., c_L))$ after seeing the history distribution $\mathcal{H}^\bullet(c_1, ..., c_L)$. Therefore, the large-generations SSS gives the same behavior as the full-attention Bayesianism in the baseline large-generations environment.