Playing Dumb Can be Smart: Behavioral Biases and Payoff Improvements in Strategic Environments

DAVID EASLEY, Cornell University
ARPITA GHOSH, Cornell University

Behavioral biases in individual decision-making are generally regarded as mistakes that reduce the individual's welfare. In this paper, we investigate one specific setting with widely documented behavioral biases—that of updating beliefs using data, where empirical and experimental evidence indicates that people do not accurately update beliefs using Bayes rule—and demonstrate, in a simple model capturing this setting, that biases need not be welfare reducing. If the parameters in the decision environment are purely exogenous, biased updating is indeed suboptimal. However, when the parameters of the decision environment are endogenously determined, as they are in several real scenarios, by a strategic agent with decision-contingent utility, a behavioral bias may actually help the decision-maker by inducing equilibrium decision parameters that improve her net payoff—despite her bias—as well as overall welfare.

Additional Key Words and Phrases: Bayes, behavioral biases, games

1. INTRODUCTION

Individual decision-making is a fundamental component in several realms of economic analysis, and has been studied from various perspectives and in various contexts, both theoretical and experimental. A growing component of this literature studies biases in individual decision-making, documenting how people make suboptimal decisions that do not match those that would be made by rational agents.

Biases sound like a bad thing[1], potentially causing harm to the individual displaying the bias, and possibly also to welfare at large. But is that really so—do biases always hurt, and if not, is there a systematic understanding of when or why they might help? In this paper, we investigate this question in one specific setting where behavioral biases have been widely documented—that of updating beliefs using data, where empirical and experimental evidence indicates that people sometimes do not act as if they accurately update beliefs according to Bayes rule.

Posterior beliefs and (non-)Bayesian updating. An extensive literature (see §1.2) documents how individuals systematically make errors by deviating from Bayes rule in updating their beliefs, displaying biases such as the confirmation bias or representativeness bias in their estimates of posterior probabilities. A typical decision-making scenario that requires updating of beliefs is that of making payoff-relevant inferences about an unknown state, or ground truth, from an observed signal: an individual must choose one amongst a set of actions, each of which has a payoff that depends on the realized state of some random variable (the 'ground truth') with some known

[1] Indeed, the Wikipedia article on biases says: "In science and engineering, a bias is a systematic error".

Arpita Ghosh acknowledges support from NSF grant CCF-1512964 and ONR Grant N00014-15-1-2335.
Authors addresses: D. Easley, Economics and Information Science, Cornell Univ.; email: dae3@cornell.edu; A. Ghosh, Information Science, Ccornell Univ.; email: ag865@cornell.edu.

EC'16, July 24–28, 2016, Maastricht, The Netherlands. ACM 978-1-4503-3936-0/16/07.
http://dx.doi.org/10.1145/XXXXXXX.XXXXXXX

prior distribution. The individual does not observe the realization, but only receives data—one amongst a possible set of signals, the probabilities of which depend on the state; she then updates her belief about the (distribution of the) ground truth based on the observed signal and uses this updated belief to guide her decision. This model describes a variety of decision-making environments, ranging from toy problems in experimental setups (such as identifying the type of an urn based on the color of the ball drawn from it, or the taxicab problem [Easley and Kleinberg 2010]), to decisions about trading under information asymmetry in online markets, where buyers without adequate first-hand information about the quality of goods need to decide whether or not to participate in the market.

Inference in strategic environments. In settings such as the urn experiment or the taxicab problem, the prior and signal probabilities are determined exogenously—by nature, or some entity (such as an experimenter) whose utility is not contingent upon the individual's choice. But there are also settings where the individual faces an identical decision problem, but where the parameters of the data-generating process have been chosen by strategic agents whose utility depends on the choice she makes. For example, sellers (whether of goods or services) deciding whether to invest in a good product or not (choosing the prior), or deciding whether to invest in a costly signal such as a warranty, or certification or coursework (choosing the signal probabilities) derive different utilities contingent on the buyer's choice of whether or not to trade in the market. How does the updating rule—Bayesian or biased—of an otherwise rational (*i.e.*, perceived expected utility-maximizing) decision-maker affect her payoff in such environments?

1.1. Outline of results

We consider the simplest version of this decision-making scenario, where the ground truth (say the 'type' of a seller) can take one of two binary values, G (good) or B (bad), the signal can take one of two possible values, H (high) or L (low), and the decision-maker (the 'buyer') chooses between two actions, t (trade) and n (not trade). The type is drawn according to a distribution $(\pi, 1-\pi)$ (the prior), where G occurs with probability π, and the two possible signals, H and L occur with probabilities that depend on the type—H occurs with probability p_G if the ground truth is G, and p_B if the type is B (and L occurs with the remaining probability $1-p_G$ or $1-p_B$). Given the observed signal S $\in$ {H,L}, and knowledge of π, p_G, p_B, the individual updates her beliefs from the prior to a posterior, and makes her decision based on the expected payoff given her posterior beliefs.

When the parameters of the data-generating process—p_G, p_B and π—are exogenous, biases are indeed sub-optimal: the decision-maker maximizes her welfare by accurately updating her beliefs using Bayes rule. However, when the data-generating process is not entirely exogenous, we show that the decision-maker can improve her welfare when she updates her beliefs with a systematic behavioral bias instead of using accurate Bayesian updating. Note that this updating rule (whether Bayesian or biased) is fixed prior to, or outside of, the game and is *not* a part of the buyer's strategy space in the game.

We study two versions of the setting where the data-generating process is endogenous—first, where the signaling probabilities p_G and p_B are exogenous, but prior $(\pi, 1-\pi)$ is chosen by a strategic agent who incurs a higher cost from choosing to be good (G) than bad (B) (§3), and second, where the seller chooses the signaling probabilities p_G and p_B, with a higher cost to the H signal when the type is bad (§4). Each of these induces a sequential game whereby the seller chooses his parameters in order to maximize his expected payoff from the buyer's choice, knowing that (a) the buyer will choose to trade if her *perceived* expected utility, given the (probabilistic) received signal (H or L), is greater than zero, and (b) knowing how the buyer computes her posterior beliefs that determine her perceived utility, given the data and the prior. We emphasize here that the buyer always plays a best response given her beliefs, *i.e.*, she chooses the action that maximizes her (perceived) expected utility—the only bias is in how she updates the prior using the data to compute her posterior.

We analyze the equilibrium outcomes in the two games, both of which[2] contain mixed equilibria (under appropriate conditions on the market) in addition to the trivial equilibrium of no trade, and ask how the non-trivial mixed equilibria vary with the buyer's updating rule. We show that a buyer who updates her prior in a consistently 'pessimistic' fashion—in that she underreacts to good news, *i.e.*, underestimates the probability that the seller is good G upon receiving an H signal, and overreacts to bad news, overestimating the probability that the seller is bad upon receiving an L signal—obtains higher payoffs in the mixed equilibria in both games than a buyer who updates accurately using Bayes rule. Interestingly, this higher payoff to the buyer comes at no cost to the seller, whose payoff in equilibrium is independent of whether the buyer updates her beliefs as an accurate Bayesian or with a bias.

The results in §3 and §4 together illustrate our main conceptual point, namely that a behavioral bias—which is suboptimal in a decision environment with exogenous parameters—can actually *help* when that same decision is instead embedded within a game which determines the parameters of the choice problem: despite the suboptimality of her decision-making given the parameters of the data-generating process, the decision-maker is better off with her behavioral bias than without, since an (appropriate) bias in her updating rule nudges the endogenous parameters of the decision environment in a direction that increase her final payoff.

Remark. While our analysis shows that a behavioral bias can be payoff-improving in the specific model we study, we clarify that we are not using this analysis to suggest that such payoff improvements are the—or even a—reason that behavioral biases are actually observed as widely as they are: to make such a claim would require an experimental design that must eliminate wide-ranging and rather complex confounding factors, which is well beyond the scope of this paper. Our primary point is to offer an alternative perspective on behavioral biases, which are typically perceived as purely irrational in the sense of being utility-reducing, by investigating their effects in a game environment rather than a pure decision environment.

1.2. Related work

Economists typically describe individual decision-making under uncertainty using expected utility theory as developed in [von Neumann and Morgenstern 1947] or [Savage 1954]; or sometimes a modern generalization of expected utility theory as in [Gilboa and Schmeidler 1989] and the many papers that build on it. This theory, in its subjective form due to [Savage 1954], is built on axioms which imply that beliefs are updated upon the arrival of new information using Bayes rule (see [Blume and Easley 2006], [Ghirardato 2002]). An immediate implication of this observation is that in the framework underlying this theory a decision maker cannot improve her payoff by misusing information from the Bayesian point of view; that is, non-Bayesian updating, followed by choosing optimally using these incorrect conditional beliefs, can only reduce payoffs.

[2]The buyer's decision problem is essentially identical in the two games, as also in the situation with purely exogenous parameters. However, the seller's strategy space and payoff structure is somewhat different in the two games, since one involves a two-dimensional strategy space (choosing p_G, p_B) while the other involves only choosing π.

Although criticisms of expected utility as a model of individual behavior have a long standing in economics, much of this work is either silent on updating or has Bayesian updating embedded in it. The basic critiques began with the experiments of [Allais 1953] calling into question the implications of objective expected utility theory and the thought experiments of [Ellsberg 1961] calling into question the implications of subjective expected utility theory. More recently, there has been renewed interest in alternative or generalized models of decision making as in [Gilboa and Schmeidler 1989] and the vast decision theory and applied economics literatures following the multiple priors approach, or the vast literature using prospect theory [Kahneman and Tversky 1979]. This literature typically does not focus on updating or uses Bayesian updating. There is however, recent interest in axiomatic decision theories using non-Bayesian updating, see [Epstein and Sandroni 2008], [Epstein and Sandroni 2010], and [Ortoleva 2012]. And, of course, there is a substantial literature in both economics and finance exploring the implications of non-Bayesian updating, see for example [Daniel et al. 1998], [Rabin and Schrag 1999] and [Rabin 2002].

Our work differs from this literature in that we do not attempt to axiomatize a non-Bayesian decision theory or to discover its implications for observed behavior in an experiment or a market. Instead, we show that non-Bayesian updating is individual payoff improving in a variety of strategic settings.

There is, of course, a large experimental literature on Bayesian updating: psychologists, as well as experimental and behavioral economists, have found that in some settings individuals do not act as if they have beliefs which are updated according to Bayes rule, see [Rabin 1998] or [Camerer 1995] for surveys that address this large literature. In some settings, individuals seem to overweight data relative to Bayes rule (or equivalently underweight their prior—base-rate neglect), see [Kahneman and Tversky 1982], and in others they seem to underweight data (or equivalently overweight their prior—conservatism), see [Kahneman and Tversky 1973]. Both of these non-Bayesian behaviors have been widely documented. For a sample of the vast literature in psychology and in economics on non-Bayesian behavior see [Dominiak and Lefort 2012; Grether 1980; Holt and Smith 2009; Kahneman and Tversky 1972, 1982; Tversky and Kahneman 1974].

2. PRELIMINARIES

We ask about the potential for payoff-improving, non-Bayesian updating in two simple, standard games derived from the decision-making scenario outlined in §1.1. Before analyzing these games, we describe the buyer's decision problem as well as other parts of the model that are common to both games.

Inference with payoff-relevant unknown ground truth. Recall from §1.1 that we have a single buyer deciding whether to trade (t) or not (n) with a seller of unknown type, based on the observed data and knowledge of the data-generating process (π, p_G and p_B). The buyer obtains a payoff of x from trading with a good seller (G), and a payoff of y from trading with a bad seller (B). If the buyer chooses to not trade with the seller, the payoff to the buyer is 0. The seller obtains a revenue of p, irrespective of his type, if the buyer trades with him and no revenue if he does not. This revenue p is an exogenous parameter throughout, and not a part of the strategy space of the seller in either game.

We assume that the buyer would (strictly) prefer to trade than not if she knows that the seller is good, and not trade if she knows the seller is bad, so that

$$x > 0 > y.$$

However, the buyer doesn't have this information about the type of the seller—rather, she only receives a signal conveying information about the seller's type (G or B). The probability of receiving a H signal varies with the type of the seller: the signal H (indicating 'high quality') is received with probability p_G if the seller is a good (G) type and with probability p_B if the seller is bad; the signal L (indicating 'low quality') is received with probabilities $1 - p_G$ and $1 - p_B$, respectively.[3] The signals H and L can be thought of as the reputation of the seller, or samples provided by the seller to demonstrate his type.

Timing and information. First the seller's type is drawn from a distribution $(\pi, 1 - \pi)$ on $\{$G,B$\}$. Next, the buyer receives one of two signals H or L drawn according to either p_G or p_B, depending on the outcome of the draw from the distribution on $\{$G,B$\}$. The buyer uses the observed data, along with the prior $(\pi, 1 - \pi)$ and the probabilities p_G and p_B, to update her beliefs about the probability that the seller is good, and uses this inference to make her decision of whether to trade or not. We assume that the buyer knows the distribution $(\pi, 1 - \pi)$ of the seller's type and the conditional probabilities, p_G and p_B, of signals.

Behavioral biases: Bayesian and non-Bayesian updating. A buyer who accurately updates her beliefs using Bayesian updating computes the posterior probability of G, given the observed data and the parameters of the data-generating process (the prior $(\pi, 1 - \pi)$ and the probabilities p_G and p_B), using Bayes rule. We denote the posterior probabilities corresponding to the two signals H and L for a Bayesian buyer as $p_H = \Pr(G|H)$ and $p_L = \Pr(G|L)$:

$$p_H = \frac{p_G \pi}{p_G \pi + p_B(1 - \pi)} \tag{1}$$

$$p_L = \frac{(1 - p_G)\pi}{(1 - p_G)\pi + (1 - p_B)(1 - \pi)} \tag{2}$$

We will allow for buyers who make inferences that are biased, or non-Bayesian, by assuming that such buyers *perceive* the posterior probabilities of the seller being good upon seeing a signal S $\in \{$H,L$\}$ as

$$q_S = f(p_S).$$

We assume that f is strictly increasing, *i.e.*, $f(p_1) > f(p_2)$ iff $p_1 > p_2$. This assumption says that the buyer, despite making errors in updating beliefs, preserves the 'ordering' of posterior probabilities, *i.e.*, she perceives more likely events as having a higher probability.

Thus, the non-Bayesian buyer knows the parameters of the data generating process (*i.e.*, the values of π, p_G and p_B), but may make errors (relative to Bayes) in how she uses data (high and low signals) to make inferences about the probability that the seller is good or bad (that individuals actually make such errors is widely documented in both the economics and psychology literature; see section §1.2). Note that these errors are *only* in inference—the buyer has accurate knowledge of the prior $(\pi, 1 - \pi)$ and the frequencies of high signals (p_G and p_B), and therefore accurately computes, for instance, the unconditional probability of seeing a high signal. Also, while the buyer may make errors in updating her beliefs, perhaps by under- or over-weighting the data she receives, she otherwise behaves 'as usual'—given her (possibly mistaken)

[3] Of course, whether H is actually a signal that the seller is a G type (or L a signal that the seller is a B type) is determined in equilibrium.

posterior beliefs, she acts to correctly maximize her expected utility as she perceives it.

Decision-making. There are four pure strategies that the buyer can use in her decision problem, corresponding to two actions (t and n) for each of the two possible signals (H, L): (i) always trade irrespective of the signal (tH, tL), (ii) trade when a high signal is received but not trade on a low signal (tH, nL), (iii) never trade (nH, nL), and (iv) trade on a low signal but not on a high signal, (nH, tL).

Suppose a signal S $\in$ {H,L} is observed. Let q_S denote the buyer's updated belief that the seller is good (G), conditional on having received the signal S—*i.e.*, q_H is the buyer's perception of the probability $\Pr(G|H)$ and q_L is the buyer's perception of the probability $\Pr(G|L)$. Define u_H and u_L to be the buyer's *perceived* expected payoffs, based on her perceived (posterior) beliefs q_H, q_L, from trading on a high (H) and low (L) signal respectively:

$$u_H = q_H x + (1 - q_H)y, \tag{3}$$

$$u_L = q_L x + (1 - q_L)y. \tag{4}$$

Her payoffs from each of her pure strategies are therefore

(1) (tH,tL): $\Pr(H)u_H + \Pr(L)u_L$
(2) (tH,nL): $\Pr(H)u_H$
(3) (nH,nL): 0
(4) (nH,tL): $\Pr(L)u_L$,

where $\Pr(H) = p_G\pi + p_B(1 - \pi)$ and $\Pr(L) = (1 - p_G)\pi + (1 - p_B)(1 - \pi)$ denote the unconditional probabilities of the buyer observing a high and low signal, respectively.[4]

Observe that if any of the parameters governing the data-generating process (either π or p_G, p_B) are endogenously determined by a strategic seller, the buyer's payoffs will have the seller's strategy embedded in them since the values of u_H and u_L depend on the values of $q_H = f(p_H)$ and $q_L = f(p_L)$, each of which are functions of π, p_G, p_B. The buyer's payoff for each strategy (and consequently her choice of strategy) therefore depends on the seller's strategy, which will be determined in an equilibrium of the corresponding game; we analyze this next.

3. STRATEGIC HIDDEN EFFORT

We now introduce strategic behavior in the buyer's decision-making environment. We start with strategic choice of π, the type of the seller: the seller can choose to be good (G), incurring a cost of c, or bad (B), which costs nothing. The seller's strategy, which may be a mixture over G and B, determines the distribution $(\pi, 1 - \pi)$.

The probabilities p_G and p_B remain exogenous, and we will assume throughout this section that

$$0 < p_B < p_G < 1 \text{ and } (p_G - p_B)p > c.$$

The first assumption says that signals are not perfectly informative. The second assumption comes from requiring that if a high signal generates trade (*i.e.*, the buyer plays t on observing H), then the return to being a good seller G, namely $p_G p - c$, is greater than the return $p_B p$ to being a bad seller B—otherwise, it would never pay for a seller to invest in being good. (Recall that the revenue p obtained upon trading is exogenous, and not part of the seller's strategy space in this game, so that assuming $(p_G - p_B)p > c$ is indeed a valid assumption on the parameters of the game).

[4]Note that these unconditional probabilities of seeing an H or L signal are accurately estimated by the buyer, in accordance with the discussion above.

This structure induces a standard, simple hidden-action problem as the buyer cannot observe the seller's action (G or B); she cares about the action choice (because the payoff to trading with a good seller x is greater than the payoff from not trading ($x > 0$), and vice versa for a bad seller ($y < 0$)); and, she observes a signal (Low or High) correlated with the seller's action. We use this canonical hidden-action game to illustrate one type of setting in which non-Bayesian updating can be advantageous.

Strategies and equilibria. The pure strategies available to the seller are to be good (G), paying a cost c, or bad (B) at a cost of 0. A seller can also use a mixed strategy; we denote the probability of choosing G in the seller's mixed strategy by π. Recall from §2 that the buyer has four pure strategies (tH, tL), (tH, nL), (nH, nL), and (nH, tL), corresponding to trading or not for each possible observed signal. A mixed strategy for the buyer is a probability distribution over these pure strategies. Both the buyer and the seller choose their strategies before any random draws are made.

A distribution $(\pi, 1-\pi)$ for the seller on (G,B), and a distribution $\alpha = (\alpha_1, \alpha_2, \alpha_3, \alpha_4)$ on the four possible strategies for the buyer is a (mixed-strategy) Nash equilibrium if neither buyer nor seller can improve their expected payoff by changing α or π given the choice of the other player.

3.1. Equilibrium analysis

We will begin by analyzing the equilibria of this game $\mathcal{G}_E$, and then investigate the equilibrium payoffs in §3.2.

LEMMA 3.1. *There are exactly three equilibria in $\mathcal{G}_E$, one pure and two mixed:*

(1) *Equilibrium E_0: The buyer plays the pure strategy (nH,nL) and the seller plays the pure strategy B.*
(2) *Equilibrium E_1: The buyer mixes over (tH,tL) and (tH,nL) with probability α_1^* given by (7) on (tH,tL) and seller mixed over G and B with probability π_1^* given by (6) on G.*
(3) *Equilibrium E_2: The buyer mixes over (tH,nL) and (nh,nL) with probability α_2^* on (nH,nL) given by (10) and the seller mixes over G and B with probability π_2^* given by (9) on G.*

Before proceeding to prove this claim, we note the following details about how non-Bayesian updating affects these equilibria and their structure; the justifications for these follow directly from the equilibrium analysis below.

(1) The no-trade and no-investment-in-quality equilibrium, ((nH,nL), B), is unaffected by non-Bayesian updating: in this equilibrium the seller does not invest, and knowing this, the buyer is best off not buying irrespective of the signal observed.
(2) In the two mixed strategy equilibria E_1 and E_2, the buyer plays strategy (tH, nL)—*i.e.*, trading conditional on a high signal—with the same probability $1-\alpha_1^* = 1-\alpha_2^*$. Note that this probability, which is the buyer's equilibrium mixture and is determined by equating the *seller's* payoffs in equilibrium from G and B, is unaffected by non-Bayesian updating.
(3) Non-Bayesian updating also does not change the set of equilibria. However, it does affect the seller's mixture over G and B in the mixed equilibria.

The remainder of §3.1 provides the argument for why these are the (only) equilibria in $\mathcal{G}_E$ and can be skipped by the reader not interested in the details.

3.1.1. Pure-strategy equilibria. We first observe that there is a single (and rather dull) pure-strategy equilibrium where the seller chooses B and the buyer chooses (nH, nL), where the seller does not try and the buyer does not trade. This is immediate from writing the payoffs to both players in normal form (recall that $x > 0 > y$):

Buyer / seller	G	B
(tH, tL)	$x, p - c$	y, p
(tH, nL)	$p_G x, p_G p - c$	$p_B y, p_B p$
(nH, nL)	$0, -c$	$0, 0$
(nH, tL)	$(1 - p_G)x, (1 - p_G)p - c$	$(1 - p_B)y, (1 - p_B)p$

3.1.2. Mixed-strategy equilibria. We begin by noting that the buyer does not play strategy (nH, tL) with positive probability in any equilibrium of the game. To see this, first rewrite the payoff from a signal S, S ∈ {H,L}, as

$$u_S = q_S x + (1 - q_S)y = q_S(x - y) + y.$$

Since $x - y > 0$, u_S is increasing in q_S. Since $p_G > p_B$, we have that $p_H > p_L$ from (3). Therefore, since $q = f(p)$ is increasing in p, $q_H > q_L$ and so $u_H > u_L$. Note also that since $0 < p_B, p_G < 1$, $\Pr(L)$ and $\Pr(H)$ are both nonzero for all $\pi \in [0, 1]$.

With this, we can now see that (nH, tL) is not a best response for any value of $\pi \in [0, 1]$ that the seller may choose: (i) if $u_H > 0$ then (tH,tL) has a greater payoff than (nH,tL), and (ii) if $u_H = 0$ then u_L must be less than 0 (since $u_H > u_L$), in which (nH,nL) has a greater payoff than (nH,tL). So (nH,tL) will not be played in an equilibrium of the game.

We next show that there are, in fact, exactly two mixed-strategy equilibria in this game. In one equilibrium, the buyer mixes between always buying (tH, tL) and buying only on a high signal (tH, nL)—that is, she mixes between trading and not when she *sees* the *low* signal, while in the other equilibrium she mixes between never buying (nH, nL) and buying only on the high signal (tH, nL) (corresponding to mixing upon seeing the *high* signal).

1. *Equilibrium E_1: Buyer mixing (tH,tL) and (tH,nL).* Suppose the buyer mixes by playing (tH,tL) with probability α_1, and (tH, nL) with probability $1 - \alpha_1$, and the seller mixes playing G w.p. π and B w.p. $1 - \pi$.

For the buyer to mix ($\alpha \in (0, 1)$), she must perceive equal payoffs from the two strategies (tH,tL) and (tH, nL) when the seller uses the strategy π. Recall that the payoff from (tH,tL) is $\Pr(H)u_H + \Pr(L)u_L$ and that from (tH, nL) is $\Pr(H)u_H$: for these to be equal, we must have $\Pr(L)u_L = 0$. Given our assumption that $0 < p_B, p_G < 1$, the probability of seeing a low signal $\Pr(L) = \pi p_G + (1 - \pi)p_B > 0$ for all $\pi \in [0, 1]$. Therefore, u_L must be 0 for the buyer to mix between these two strategies, *i.e.*,

$$q_L x + (1 - q_L)y = 0 \qquad \Rightarrow q_L = \frac{-y}{x - y}. \tag{5}$$

Note that $q_L \in (0, 1)$ since $x > 0 > y$. So in this mixed equilibrium the seller's strategy π_1^* must solve

$$q_L = f(p_L) = f\left(\frac{\pi_1^*(1 - p_G)}{\pi_1^*(1 - p_G) + (1 - \pi_1^*)(1 - p_B)}\right) = \frac{-y}{x - y}.$$

The equilibrium value of π is thus

$$\pi_1^* = \frac{(1 - p_B)f^{-1}(z)}{(1 - p_G)(1 - f^{-1}(z)) + (1 - p_B)f^{-1}(z)}, \tag{6}$$

where $z = \frac{-y}{x-y}$. Since $\pi_1^* \in [0, 1]$, it describes a valid mixed strategy for the seller.

If we can now also find an $\alpha_1 \in (0, 1)$ such that the seller is indifferent between the strategies of G and B when the buyer plays (tH,tL) with probability α_1, and (tH, nL) with probability $1 - \alpha_1$, we have a mixed strategy equilibrium. Note that the value of α_1 characterizing the buyer's mixed strategy does *not* depend on whether or not the buyer uses accurate Bayesian updating or not: α_1 is determined by equating the *seller's* payoffs from choosing G and B when the buyer mixes according to α. Those payoffs depend on the *true* probabilities p_G and p_B of seeing an H signal, rather than the buyer's *perception* of those probabilities.

The payoff to the seller when she chooses G is

$$-c + [\alpha_1 p + (1 - \alpha_1)(p_G p)]$$

while the payoff from choosing B is

$$\alpha_1 p + (1 - \alpha_1)(p_B p).$$

Equating these two allows us to solve for the value of α_1 at which the seller is indifferent between G and B:

$$\alpha_1^* = 1 - \frac{c}{p(p_G - p_B)} \tag{7}$$

which lies in $(0, 1)$ by the assumption that $p_G p - c > p_B p$.

Since both π_1^* and α_1^* are valid probabilities in $[0, 1]$, we have a mixed-strategy equilibrium.

2. *Equilibrium E_2: Buyer mixing (tH,nL) and (nH,nL).* There is another mixed equilibrium in this game where the buyer mixes between (tH,nL) and (nH, nL). Again, for the buyer to mix, the seller's mixed strategy must be such that she receives equal payoffs from the two strategies (tH,nL) and (nH, nL). Equating the payoff from (tH,nL), which is $\Pr(H)u_H$ and that from (nH,nL), which is zero, we obtain:

$$q_H x + (1 - q_H) y = 0 \qquad \Rightarrow q_H = f(p_H(\pi)) = \frac{-y}{x - y}. \tag{8}$$

Once again, note that $q_H \in (0, 1)$ since $x > 0 > y$. So in this mixed equilibrium the seller's strategy π_2^* must solve

$$q_H = f(p_H) = f\left(\frac{\pi_2^* p_G}{\pi_2^* p_G + (1 - \pi_2^*) p_B}\right) = \frac{-y}{x - y}.$$

The equilibrium value of π is thus

$$\pi_2^* = \frac{f^{-1}(z) p_B}{p_G(1 - f^{-1}(z)) + p_B f^{-1}(z)}, \tag{9}$$

where $z = \frac{-y}{x-y}$. Note that $\pi_2^* \in [0, 1]$.

For the seller to be willing to mix, he must obtain equal payoffs from G and B when the buyer mixes with probability $1 - \alpha_2$ on (tH, nL) and α_2 on (nH, nL). The payoff to the seller when she chooses G and the buyer mixes wp α is

$$(1 - \alpha_2)(p_G p - c) + \alpha_2(-c),$$

while the payoff from choosing B is

$$(1 - \alpha_2) p_B p + \alpha_2 \cdot 0.$$

Equating these two allows us to solve for the value of α at which the seller is indifferent between G and B, which is

$$\alpha_2^* = 1 - \frac{c}{p(p_G - p_B)} \in (0, 1). \tag{10}$$

Again, note that α_2^* does not vary with $f(\cdot)$.

3.1.3. No other equilibria. We now argue that the three equilibria above are the only equilibria in the game. We have already seen that there is exactly one pure strategy equilibrium of the game, and that there is exactly one equilibrium where the buyer mixes between (tH, nL) and (tH, tL), and exactly one where the buyer mixes between (tH, nL) and (nH, nL). Now note that:

(a) There is no equilibrium where the buyer mixes between (tH, tL) and (nH, nL), since the seller's unique best response to both strategies is the pure strategy B, to which the buyer's unique best response is (nH,nL).
(b) There is no equilibrium where the buyer plays all three strategies (tH, tL), (tH, nL) and (nH, nL) with non-zero probability, for any mix π chosen by the seller: If so, the buyer must obtain equal payoffs from all three strategies, which requires

$$\Pr(H)u_H + \Pr(L)u_L = \Pr(H)u_H = 0,$$

which cannot be simultaneously satisfied, since we have already argued that (i) both $\Pr(H)$ and $\Pr(L)$ are nonzero for all π, and (ii) that $u_H > u_L$, so that we cannot simultaneously have $u_H = 0$ and $u_L = 0$.

Together, we have shown that there exist no equilibria other than the three equilibria above.

3.2. Equilibrium payoffs

We now investigate equilibrium payoffs, and how they depend on how the buyer updates her beliefs, $f(p)$.

First, note that the seller's payoff in any equilibrium depends only on the *buyer's* mixture over strategies. Since the buyer's mixture over strategies is independent of f, which describes how the buyer mis-updates posterior beliefs, the *seller's* payoff is independent of f as well. Thus, the buyer's inaccurate updating does not harm, or help, the seller.

Next we investigate the buyer's payoffs. The buyer's payoff in the pure strategy equilibrium where she plays (nH,nL) and the seller plays B is always 0, irrespective of f. For the mixed equilibria, the buyer's *true* payoffs (in contrast with her perceived payoffs) can be computed as the weighted sum of the payoffs from each of her component pure strategies. Note that when the buyer updates accurately, the payoffs from all pure strategies that the buyer mixes over are equal, but if the buyer updates inaccurately, it is the *perceived*, rather than the true, payoffs from strategies belonging to a mixture that are equal.

The true payoff from the strategy (tH, tL) is

$$u_{tH,tL} = \pi x + (1 - \pi)y = \pi(x - y) + y,$$

from the strategy (tH,nL) is

$$u_{tH,nL} = \pi p_G x + (1 - \pi)p_B y = \pi(p_G x - p_B y) + p_B y,$$

while the payoff from (nH,nL) is 0.

Note that the payoffs $u_{tH,tL}$ and $u_{tH,nL}$ are both (strictly) increasing in π: since $x > 0$ and $y < 0$, and $p_G > p_B$, the coefficients of π are positive. Therefore, the total payoff in both the equilibria E_1 and E_2 are increasing in π, since they are the sum of terms each of which is (weakly) increasing in π.

Finally, we ask how the seller's probability π of playing G in the two mixed equilibria E_1 and E_2 depends on f, the function describing how the buyer updates her posterior belief. For the first equilibrium where the buyer mixes between (tH,tL) and (tH,nL), recall that π_1^* is such that

$$f(p_L) = f\left(\frac{\pi_1^*(1-p_G)}{\pi_1^*(1-p_G)+(1-\pi_1^*)(1-p_B)}\right) = \frac{-y}{x-y},$$

while for the other equilibrium where the buyer mixes between (tH,nL) and (nH,nL), the seller mixes according to π_2^* which satisfies

$$f(p_H) = f\left(\frac{\pi_2^* p_G}{\pi_2^* p_G + (1-\pi_2^*)p_B}\right) = \frac{-y}{x-y}.$$

Rewriting p_L and p_H, we can see immediately that p_H and p_L are both strictly increasing in π:

$$\begin{aligned} p_L(\pi) &= \frac{\pi(1-p_G)}{\pi(1-p_G)+(1-\pi)(1-p_B)} = \frac{1}{1+\left(\frac{1}{\pi}-1\right)\frac{(1-p_B)}{(1-p_G)}} \\ p_H(\pi) &= \frac{\pi p_G}{\pi p_G + (1-\pi)p_B} = \frac{1}{1+\left(\frac{1}{\pi}-1\right)\frac{p_B}{p_G}}. \end{aligned}$$

Since $f(p)$ is strictly increasing in its argument, and since the arguments p_H and p_L are strictly increasing in π, we have that f is strictly increasing in π.

Now suppose that the buyer updates according to a function f that consistently underestimates posterior probabilities [5], *i.e.*, such that $f(p) \leq p$. For such a buyer, the equilibrium values of π_1^* and π_2^* are *larger* than they are for a Bayesian buyer with $f(p) = p$: a larger π is needed to generate a larger p_H (respectively p_L) so that $f(p_H)$ (respectively $f(p_L)$) takes on the same value of $\frac{-y}{x-y}$. Therefore, the seller chooses higher values of π in both equilibria of the game when the buyer perceives posterior probabilities according to a function $f(p) < p$: that is, if a buyer is 'pessimistic'—relative to the truth—in her estimate of the probability that the seller is good conditional on the received data, the seller will choose a more favorable mix over being good and bad.

In fact, this buyer pessimism is *Pareto-improving*, since this increase in π comes at no cost to the seller—the seller mixes only over strategies that give him equal payoff, and therefore *any* mix over the two pure strategies is payoff-equivalent for the seller. That is, the buyer's pessimism, $f(p) < p$, has made her strictly better off without making the seller worse off. We summarize these arguments as a theorem below.

THEOREM 3.2. *In each of the non-trivial equilibria E_1 and E_2 of game $\mathcal{G}_E$, the payoff to a non-Bayesian 'pessimistic' buyer, who perceives her posterior beliefs according to $f(p) < p$, is strictly larger than that of a Bayesian buyer ($f(p) = p$). Further, the seller's payoff is independent of the buyer's updating rule $f(p)$, so that non-Bayesian updating with $f(p) < p$ Pareto-improves welfare.*

This Pareto improvement generated by non-Bayesian updating may seem surprising, but can be understood simply as follows. Relative to a Bayesian buyer our pessimistic, non-Bayesian buyer appears to be not best-responding to the seller's strategy, and a commitment to not best-respond can be valuable in a variety of games (of course,

[5] Note that this is not at odds with f being strictly increasing in its argument; for example, the function $f(p) = p^2$ is strictly increasing in p and less than p for $p \in [0,1]$).

note that this biased updating is *not* strategically chosen within the game but comes instead from a behavioral rule that is fixed prior to the game). Thus, while our buyer *is* best-responding *given* her incorrect posterior beliefs—she chooses the action that maximizes her *perceived* expected utility—the fact that her beliefs are incorrect effectively functions as a credible threat to not best-respond, a threat the buyer does in fact carry out; this credible threat earns the buyer more in a favorable distribution over types than she loses with her suboptimal updating.

4. STRATEGIC SIGNALING

We next consider a simple, standard signaling game using the structure outlined previously. The difference from the previous section is that now strategic behavior is with respect to signaling: the distribution over types, $(\pi, 1-\pi)$, is exogenous, but the probabilities p_G and p_B of a high signal given a good or bad type are chosen by a strategic seller. The cost to a seller to sending the high signal H is c_G if he is good (G), and c_B is he is bad (B), with $c_B > c_G$; the cost of sending a low signal L is always zero. We assume that the revenue p that a seller receives from trade is greater than c_B, so that both types of sellers would be willing to signal if this would lead to trade.

Assumptions on market parameters π, x, y and updating function f. To make the game non-trivial, we make the following two assumptions: (i) when the seller invests in a perfectly informative signal, the buyer wants to trade when the signal indicates the seller is good and not when it implies he is bad, and (ii) the buyer is not willing to trade unless the seller invests in an informative signal. For a Bayesian buyer, these assumptions translate to $x > 0 > y$ and $\pi x + (1-\pi)y < 0$ respectively. For a non-Bayesian buyer, they generalize[6] as follows:

(1) The assumption corresponding to when the seller invests in a perfectly informative signal (HGLB or LGHB) translates to the inequalities $f(1)\cdot x+(1-f(1))\cdot y > 0$, and $f(0)\cdot x + (1-f(0))\cdot y < 0$: With the strategy HGLB, $p_G = 1, p_B = 0$ so that $p_H = 1$ and $p_L = 0$. Requiring the perceived expected utilities upon seeing H (respectively L) to be bigger (respectively smaller) than 0 yields these inequalities; analyzing LGHB gives the same inequalities again.
(2) The assumption that the buyer does not trade if the seller does not invest in an informative signal (*i.e.*, plays HGHB or LGLB) translates to $f(\pi)\cdot x+(1-f(\pi))\cdot y < 0$. This inequality follows from noting that if $p_G = p_B$ then $p_H = p_L = \pi$. Also, since a buyer who always trades—*i.e.*, plays (tH,tL)—is essentially trading against the prior π, we would like the buyer's perception to be 'accurate enough' that she perceives her expected payoff from this strategy to be negative, irrespective of the seller's strategy. This leads to two inequalities on f corresponding to informative and uninformative seller strategies: the first is the same as the assumption we just made, while the second adds the assumption that $\pi(f(1)\cdot x + (1-f(1))\cdot y) + (1-\pi)(f(0)\cdot x + (1-f(0))\cdot y) < 0$.

We note that the stronger, but simpler, assumption that $f(0) = 0$ and $f(1) = 1$ in addition to $f(\pi)\cdot x + (1-f(\pi))\cdot y < 0$ would be sufficient for our analysis; we state the less simple assumptions above because they are a strictly weaker set of requirements on the f function.

[6](Recall that $f(p) = p$ for a Bayesian buyer.)

4.1. Equilibrium analysis

The seller has four pure strategies in this game: HGHB, HGLB, LGLB, and LGHB (as before, the buyer has the same four pure strategies as in §2, and both the buyer and the seller choose their strategies before any random draws are made). Let $\gamma = (\gamma_1, \gamma_2, \gamma_3, \gamma_4)$ denote the distribution placed on these strategies by the seller. The probabilities p_G and p_B relate to the probabilities γ_i of the pure strategies as:

$$p_G = \gamma_1 + \gamma_2; \qquad p_B = \gamma_1 + \gamma_4. \tag{11}$$

Again, we start with the equilibria of this game $\mathcal{G}_S$, and then investigate equilibrium payoffs.

LEMMA 4.1. *There are two equilibria in $\mathcal{G}_S$, one pure and one mixed:*

(*1*) *The buyer plays the pure strategy (nH,nL) and the seller plays the pure strategy (LGLB).*
(*2*) *The buyer mixes over (tH,nL) and (nH,nL) with probability α_2 on (tH,nL) given by (13) and seller mixes over HGHB and HGLB with probability γ_1 on HGHB implicitly given by (12).*

Remark. That there are only two equilibria in this game—as opposed to an odd total number of equilibria—can be understood by noting the following facts: (i) Strategies (tH,tL), LGLB and (nH,tL) are not rationalizable. (ii) In the restricted two-by-two game with buyer strategies (tH,nL) and (nH,nL) and seller strategies HGHB and HGLB, there is a unique (mixed) equilibrium. (iii) There is one additional (pure strategy) equilibrium when the strategy LGLB for the seller is added to the two-by-two game. (iv) The payoffs to both players are 0 if the seller plays LGLB and the buyer's payoffs to the strategy (nH,nL) are 0 regardless of the seller's strategy. This special payoff structure is responsible for the unusual, even number of equilibria in this game.

There are a few details to note about these equilibria:

(1) First, the no-trade and no-investment in signaling equilibrium, ((nH,nL), LGLB), is unaffected by non-Bayesian updating. In this equilibrium, the seller never invests in signaling, and the buyer never buys.
(2) In the mixed strategy equilibrium, the *buyer's* strategy, specifically her probability of playing (tH, nL)—trading conditional on a high signal—is not affected by non-Bayesian updating.
(3) Non-Bayesian updating again does not change the set of equilibria, but it does affect the seller's mixed strategy.

The remainder of §4.1 provides the argument for why these are the (only) equilibria and can be skipped by the reader not interested in the details.

4.1.1. Pure strategy equilibrium. We begin by analyzing pure-strategy equilibria, writing the payoffs in normal form. For brevity, the payoffs below are written for a Bayesian buyer, with $u_b = \pi x + (1-\pi)y$ denoting the expected payoff to the buyer from trading when her belief that the seller is good is π, her prior. For a non-Bayesian buyer, the entries need to be modified appropriately to reflect the buyer's perceived expected utility. However, the *comparisons* between the buyer's payoffs for each combination of buyer-seller strategy remains the same as for a Bayesian buyer because of our assumptions on the function f, and therefore the arguments that follow in the equilibrium analysis below remain valid.[7]

[7]For example, the payoff from ((tH,tL), HGHB), as also ((tH,tL), HGHB) is $f(\pi)x + (1 - f(\pi))y < 0$; the payoff from ((tH,nL), HGLB) is $\pi(f(1)x + (1 - f(1))y) > 0$, and so on.

Buyer / seller	HGHB	HGLB	LGLB	LGHB
(tH, tL)	$u_b, p - (\pi c_G + (1-\pi)c_B)$	$u_b, p - \pi c_G$	u_b, p	$u_b, p - (1-\pi)c_B$
(tH, nL)	$u_b, p - (\pi c_G + (1-\pi)c_B)$	$\pi x, \pi(p - c_G)$	$0, 0$	$(1-\pi)y, (1-\pi)(p - c_B)$
(nH, nL)	$0, -(\pi c_G + (1-\pi)c_B)$	$0, -\pi c_G$	$0, 0$	$0, -(1-\pi)c_B$
(nH, tL)	$0, -(\pi c_G + (1-\pi)c_B)$	$(1-\pi)y, (1-\pi)p - \pi c_G$	u_b, p	$\pi x, \pi p - (1-\pi)c_B$

Note that since the buyer's expected payoff from trading when her belief that the seller is good is the prior, u_b, is negative, (tH, tL) is a strictly dominated strategy for the buyer: it is dominated by the strategy (nH, nL) which always yields 0 payoff.

Next, it is easy to see that ((nH,nL), LGLB) is a pure-strategy equilibrium, and in fact is the unique pure equilibrium in the game: there is no equilibrium where the seller plays (i) HGHB (the buyer's best response to HGHB is (nH,nL) and (nH, tL), for both of which the seller would be better off playing LGLB); (ii) HGLB, since the buyer's unique best response is (tH,nL) to which the seller's unique best response is HGHB; (iii) LGHB, since the buyer's unique best response is (nH,btL) to which the seller's unique best response is LGLB. Finally, note that the only pure strategy equilibrium with the seller playing LGLB has the buyer playing (nH,nL), since the seller's best response to (tH,nL), the other best response of the buyer to LGLB, is to deviate to HGHB.

4.1.2. Mixed strategy equilibrium. The outline of the mixed equilibrium analysis is as follows. With (tH,tL) being dominated for the buyer, we show that (LGHB) is never a best response for the seller. With LGHB eliminated, we show that (nH,tL) is never played by the buyer in any equilibrium. We now have a reduced game with two strategies for the buyer —(tH,nL) and (nH,nL)—and three for the seller, HGHB, HGLB, and LGLB. First we show that LGLB is not played with non-zero probability in any equilibrium other than the pure equilibrium (nH,nL)-LGLB. We then solve for the unique mixed-strategy equilibrium where the buyer mixes between (tH,nL) and (nH,nL) and the seller mixes between HGHB and HGLB.

We first show that the seller never plays LGHB given that (tH, tL) is a strictly dominated strategy for the buyer. To see this, suppose the buyer uses a mixture with probabilities $\alpha_2, \alpha_3, \alpha_4$ respectively on the remaining three strategies (tH,nL), (nH,nL) and (nH,tL). If the seller ever plays LGHB, it must be that LGHB is a best response for *some* values of $\alpha_2, \alpha_3, \alpha_4$. The seller's payoffs from the four strategies are

(1) HGHB: $\alpha_2 p - (\pi c_G + (1-\pi)c_B)$
(2) HGLB: $\alpha_2 \pi p + \alpha_4(1-\pi)p - \pi c_G$
(3) LGLB: $\alpha_4 p$
(4) LGHB: $\alpha_2(1-\pi)p + \alpha_4 \pi p - (1-\pi)c_B$

For LGHB to be a best response its payoff must be at least as large as that from HGHB. This yields the following constraint on the α values:

$$\begin{aligned} \alpha_2(1-\pi)p + \alpha_4\pi p - (1-\pi)c_B \ &\geq \ \alpha_2 p - (\pi c_G + (1-\pi)c_B) \\ \Rightarrow (\alpha_2 - \alpha_4) \ &\leq \ \frac{c_G}{p}. \end{aligned}$$

Also, the payoff from LGHB must be at least as large as that from LGLB:

$$\begin{aligned} \alpha_2(1-\pi)p + \alpha_4\pi p - (1-\pi)c_B \ &\geq \ \alpha_4 p \\ \Rightarrow (\alpha_2 - \alpha_4) \ &\geq \ \frac{c_B}{p}. \end{aligned}$$

But $c_B > c_G$, so these two inequalities can never be simultaneously satisfied. Therefore LGHB is never a best response.

Given this, we will next show that there does not exist an equilibrium in which (nH, tL) is played with positive probability. First note that for any $\gamma_1, \gamma_2, \gamma_3$ with $\gamma_1+\gamma_2+\gamma_3 = 1$ the strategy (nH, nL) is a strictly better response that (nH, tL) unless $\gamma_1 = 1$. So (nH, tL) can only be played if the seller plays that pure strategy HGHB. Second, note that at $\gamma_1 = 1$ the buyer is indifferent between (nH, nL) and (nH, tL). But if the buyer plays any mixture of (nH, nL) and (nH, tL) the seller's unique best response is LGLB, not HGHB. So there is no equilibrium in which (nH, tL) is played.

Together, these arguments leave us with the following truncated game:

Buyer / seller	HGHB	HGLB	LGLB
(tH, nL)	$u_b, p-(\pi c_G+(1-\pi)c_B)$	$\pi x, \pi(p-c_G)$	0, 0
(nH, nL)	$0, -(\pi c_G+(1-\pi)c_B)$	$0, -\pi c_G$	0, 0

Recall that we denote the probabilites used by the buyer on strategies (tH,nL) and (nH,nL) by α_2 and α_3 (with $\alpha_2 + \alpha_3 = 1$), and the probabilities used by the seller on strategies HGHB, HGLB, LGLB by $\gamma_1, \gamma_2, \gamma_3$ respectively (again, $\gamma_1 + \gamma_2 + \gamma_3 = 1$).

Note that $p_G = \Pr(H|G) = \gamma_1 + \gamma_2$ (the seller sends a H conditional on being good in the strategies HGHB and HGLB) and $p_B = \gamma_1 + \gamma_4 = \gamma_1$ (the seller sends H if he is a bad type in strategies HGHB and LGHB respectively; recall that we argued that $\gamma_4 = 0$ always).

First we argue that there is no mixed equilibrium with $\gamma_3 > 0$: 1. If all three probabilities $\gamma_1, \gamma_2, \gamma_3$ are non-zero, the buyer's mixture over the two strategies (α_2 on (tH,nL) and $\alpha_3 = 1 - \alpha_2$ on (nH,nL) must be such that the seller derives equal payoffs from HGHB, HGLB, and LGLB, which requires:

$$\alpha_2 p - (\pi c_G + (1-\pi)c_B) = \alpha_2 \pi p - \pi c_G = 0,$$

which requires α_2 to simultaneously equal $\frac{c_B}{p}$ and $\frac{c_G}{p}$, which is impossible (we assumed $c_G < c_B$). 2. There is also no equilibrium with $\gamma_3 > 0$ and only one of the remaining two probabilities non-zero: for the buyer to derive the same payoff from the strategies of (tH,nL) and (nH,nL), we must have $\Pr(H)u_H = 0$. Suppose $\gamma_1 = 0$, then $p_B = 0$. With $p_B = 0$, we have that $p_H = 1$ (unless $p_G = 0$ as well in which case p_H is undefined, but for $p_G = 0$ we need $\gamma_2 = 0$ also, which implies $\gamma_3 = 1$ corresponding to the pure strategy LGLB for the seller, which we have already analyzed. With $p_H = 1$, and our assumption that

$$f(1) \cdot x + (1 - f(1)) \cdot y > 0,$$

we have $u_H > 0$ so we cannot have $\Pr(H)u_H = 0$ since $\Pr(H) > 0$ (recall that $\pi \in (0, 1)$.) Similarly, if $\gamma_2 = 0$, $p_G = p_B$ which implies that $p_H = p_L = \pi$. With the assumption that

$$f(\pi) \cdot x + (1 - f(\pi)) \cdot y < 0,$$

$u_H(\pi) \neq 0$ and $\Pr(H) > 0$ as before, so we cannot have $\Pr(H)u_H = 0$. So we have ruled out equilibria that mix LGLB with one of HGHB or HGLB for the seller. Therefore, if there is a mixed equilibrium, it must satisfy $\gamma_3 = 0$.

We now solve for this equilibrium. With $\gamma_3 = 0$, we know that $p_G = \gamma_1 + \gamma_2 = 1$. This means that $\Pr(H) \neq 0$: for the buyer to mix between (tH,nL) and (nH,nL), therefore, γ_1 must be such that $u_H = 0$, or

$$q_H x + (1 - q_H)y = 0 \Rightarrow f(p_H) = q_H = \frac{-y}{x - y}. \tag{12}$$

For the seller to mix, we must have

$$\alpha_2 p - (\pi c_G + (1-\pi)c_B) = \alpha_2 \pi p - \pi c_G \qquad \Rightarrow \qquad \alpha_2 = \frac{c_B}{p}. \tag{13}$$

4.2. Equilibrium payoffs

The buyer's *true* payoff in the mixed equilibrium is

$$\begin{aligned}\Pr(H)(p_H(x-y)+y) &= \Pr(G,H)(x-y)+\Pr(H)y \\ &= \pi(x-y)+(\pi+(1-\pi)\gamma_1)y \\ &= \pi x+(1-\pi)\gamma_1 y.\end{aligned}$$

Since $y < 0$, the buyer's payoff is decreasing in γ_1 $(= p_B)$.

The equilibrium value of γ_1 is given by the solution to

$$f(p_H) = f\left(\frac{\pi}{\pi+(1-\pi)\gamma_1}\right) = \frac{-y}{x-y}.$$

When $f(p) < p$, the value that γ_1 needs to take to satisfy this equation is smaller than that when $f(p) = p$, since f is strictly increasing in its argument and the argument is strictly decreasing in γ_1 (recall $\pi \in (0,1)$. Therefore, the buyer's welfare increases when she miscomputes posterior probabilities by $f(p) < p$ and again, buyer pessimism, $f(p) < p$, is Pareto improving since seller payoff is unaffected by f. We summarize this analysis in a result very similar to Theorem 3.2:

THEOREM 4.2. *The payoff to a non-Bayesian buyer who perceives her posterior beliefs according to $f(p) < p$ is strictly larger than that of a Bayesian buyer ($f(p) = p$) in the non-trivial equilibrium of game $\mathcal{G}_S$. Further, since the seller's payoff does not vary with the buyer's updating rule $f(p)$, non-Bayesian updating with $f(p) < p$ Pareto-improves welfare.*

5. CONCLUSION AND FURTHER DIRECTIONS

In this paper, we asked whether behavioral biases, which are typically perceived as welfare-reducing errors in individual choice, are indeed always 'bad'—and if not, whether there is a systematic reason they might be valuable. We investigated this question in the specific context of biases in Bayesian updating of beliefs, providing an affirmative answer: Our results in Theorems 3.2 and 4.2 show that an 'unusually suspicious' buyer, who underreacts to good news and overreacts to bad news, obtains higher payoffs in equilibrium than a buyer who accurately uses Bayes rule to compute her posterior beliefs from data. Thus, while biases are known to be suboptimal when the input parameters to a decision are exogenous, they may actually improve payoffs when the same decision must be made with input parameters that are endogenously determined by a strategic agent with decision-contingent payoffs: when the input parameters to the decision are determined in a game, the decision-maker's overall payoff may improve—despite her bias, which remains suboptimal given the parameters—because the parameters of the environment are (adequately) modified in her favor.

We note here that we do not intend to conclude from our results that a pessimism bias is to be expected in a population that routinely makes decisions in environments similar to our model: whether a bias at all improves payoffs, and what specific direction it must take to improve payoffs, is very possibly dependent on the specific game in question. Our primary intent is to illustrate our central point—namely that behavioral biases needn't be bad—in the simplest possible model that captures the essence of a range of situations where an individual makes payoff-relevant inferences from data.[8]

Future work. This observation—that a decision-maker's behavioral bias may actually improve her final payoff in a decision environment with endogenous parameters—

[8]This is also the reason we do not analyze a natural third game where both π and p_G, p_B are controlled by a strategic agent.

suggests some intriguing questions, in addition to the immediate and natural one of the welfare effects of other behavioral biases in decision-making in game environments:

— *Design of online marketplaces: Reputations and social norms.* The fact that the buyer's welfare improves without cost to the seller suggests potential implications for the design of online markets with information asymmetries, if such a behavioral bias can be implemented via the design of the platform itself. In online marketplaces for goods or services, a buyer cannot directly observe the quality (type) of a seller or provider, but only sees signals in the form of reviews from previous buyers. The essence of the analysis of equilibrium buyer welfare—namely that a buyer compares her expected payoff conditional on seeing a signal (or a set of signals) to decide whether or not to trade—provides some insight here as well, despite the larger signal set: as long the probability of the seller being good conditional upon observing a signal remains increasing in the prior, a pessimism bias in the update will induce the seller to choose a higher π in an equilibrium (provided that equilibrium continues to exist). *If* posterior beliefs about sellers are controlled to any degree by the design of the marketplace—for instance, the aggregated rating of a seller, or a reputation score—this suggests that a scheme that essentially punishes sellers more for bad reviews than it rewards them for good ones might improve quality and welfare in the market.
  Another question relates to observed social norms in online marketplaces, with under-reporting of negative experiences in some online communities, or over-reporting of extreme experiences (positive or negative) in other marketplaces. Suppose good and bad types generate positive (and negative, mediocre, and so on) experiences at certain 'true rates but are selectively reported, yielding $\tilde{p}_G, \tilde{p}_B$ that differ from the true rates. This induces a bias in the corresponding conditional beliefs, leading to the following question: are the biases induced by the prevailing social norms payoff-improving for the participants in the market, and can these social norms emerge as ones that survive in an evolutionary model of such a marketplace?
— *Endogenous selection of behavioral biases in the population.* We do not model or analyze how (decision-makers with) biases might be 'selected for' in a game, which could be analyzed either within an evolutionary model[9], or a model with endogenous entry from a diverse population. If equilibrium payoffs depend on biases, which bias occurs in the population might depend on which game is being played, and it is conceivable that payoff-improving behavioral biases might persist in the population playing a game. This suggests intriguing questions regarding mechanism design—how does the design of a game influence the biases in the population playing the game? If the mechanism does determine which population plays the game, this has a potentially significant conceptual implication for design: rather than assume a certain (homogeneous) model of behavior for her agent population, a mechanism designer should instead account for the fact that (the game induced by) her design might select for certain behavioral biases—so that the model of behavior (whether biased or otherwise) now becomes an endogenous component of the design problem.

REFERENCES

M. Allais. 1953. Le Comportement de l'Homme Rationnel devant le Risque: Critique des Postulats et Axiomes de l'Ecole Americaine. *Econometrica* 21, 4 (1953), 503–546.

[9]A very simple evolutionary model where this is indeed the outcome can be constructed; we choose omit it because it does not add much insight beyond the existing analysis.

L. Blume and D. Easley. 2006. Rationality and Selection in Asset Markets. In *Santa Fe Institute Volume in honor of Kenneth Arrow: The Economy as an Evolving Complex System*.

Colin Camerer. 1995. Individual Decision Making. In *The Handbook of Experimental Economics*.

Kent Daniel, David Hirshleifer, and Avanidhar Subrahmanyam. 1998. Investor Psychology and Security Market Under- and Overreactions. *The Journal of Finance* 53, 6 (1998), 1839–1885.

Duersch P. Dominiak, A. and J-P. Lefort. 2012. A Dynamic Ellsberg Urn Experiment. *Games and Economic Behavior* 75, 2 (2012), 625–638.

David Easley and Jon Kleinberg. 2010. *Networks, Crowds and Markets: Reasoning About a Highly Connected World*. Cambridge University Press.

Daniel Ellsberg. 1961. Risk, Ambiguity, and the Savage Axioms. *The Quarterly Journal of Economics* 75, 4 (1961), 643–669.

Jawwad Noor Epstein, Larry G. and Alvaro Sandroni. 2008. Non-Bayesian Updating: A Theoretical Framework. *Theoretical Economics* 3, 2 (2008), 193229.

Jawwad Noor Epstein, Larry G. and Alvaro Sandroni. 2010. Non-Bayesian Learning. *B. E. Journal of Theoretical Economics: Advances in Theoretical Economics* 10, 1 (2010).

Paolo Ghirardato. 2002. Revisiting Savage in a conditional world. *Economic Theory* 20, 1 (2002), 83–92.

Itzhak Gilboa and David Schmeidler. 1989. Maxmin expected utility with non-unique prior. *Journal of Mathematical Economics* 18, 2 (1989), 141 – 153.

David Grether. 1980. Bayes Rule as a Descriptive Model: The Representativeness Heuristic. *The Quarterly Journal of Economics* 95, 3 (1980), 537–557.

Charles Holt and Angela Smith. 2009. An Update on Bayesian Updating. *Journal of Economic Behavior and Organization* 69, 2 (2009), 125–134.

Daniel Kahneman and Amos Tversky. 1972. Subjective Probability: A Judgment of Representativeness. *Cognitive Psychology* 3, 3 (1972), 430–454.

Daniel Kahneman and Amos Tversky. 1973. On the Psychology of Prediction. *Psychological Review* 80, 4 (1973).

Daniel Kahneman and Amos Tversky. 1979. Prospect Theory: An Analysis of Decision under Risk. *Econometrica* 47, 2 (1979).

Daniel Kahneman and Amos Tversky. 1982. Evidential Impact of Base Rates. In *Judgment Under Uncertainty: Heuristics and Biases*.

Pietro Ortoleva. 2012. Modeling the Change of Paradigm: Non-Bayesian Reactions to Unexpected News. *American Economic Review* 102, 6 (2012), 2410–36.

Matthew Rabin. 1998. Psychology and Economics. *Journal of Economic Perspectives* (1998).

Matthew Rabin. 2002. Inference by Believers in the Law of Small Numbers. *The Quarterly Journal of Economics* 117, 3 (2002), 775–816.

Matthew Rabin and Joel L. Schrag. 1999. First Impressions Matter: A Model of Confirmatory Bias. *The Quarterly Journal of Economics* 114, 1 (1999), 37–82.

Leonard Savage. 1954. *The Foundations of Statistics*. Wiley.

Amos Tversky and Daniel Kahneman. 1974. Judgment Under Uncertainty: Heuristics and Biases. *Science* 185, 4157 (1974), 1124–1131.

John von Neumann and Oscar Morgenstern. 1947. *Theory of Games and Economic Behavior*. Princeton University Press.