Strategic Self-Deception*

Greg Kubitz[1], Claudio Mezzetti[2], and Lionel Page[2]

[1]Queensland University of Technology
[2]University of Queensland

April 2025

Abstract

This paper develops a model of strategic self-deception in a competitive context. We study a setting where an agent can covertly gather information about his own ability before deciding whether to compete with another agent. The agent uses an information gathering-strategy that increases the likelihood of forming confident Bayesian beliefs that induce him to compete. This action credibly signals his self-belief and, as a consequence, deters his competitor and improves his expected payoff. The model provides an explanation for two important behavioural phenomena: self-deception and overconfidence. The agent's information-gathering strategy can be understood as self-deception, as he deliberately limits what he learns in order to form confident beliefs. It generates overconfidence, as the agent competes more often than if he were fully informed.

Keywords: Competitive interaction, covert information gathering, motivated reasoning, self-beliefs, self-deception.

JEL-Classification: D91, D83, D82, C72.

*We thank for feedback the seminar audiences at NASM 2024 at Vanderbilt, Ca' Foscari University of Venice, and City University of Hong Kong, The Paris School of Economics, ESAM 2024, Deakin, FUR 2024, University of Sydney, ENS Paris Saclay, LEG 2023, RMIT, and Monash. Mezzetti gratefully acknowledges financial support through a visiting research scholar grant from the Department of Economics at Ca' Foscari University of Venice and the Australian Research Council Grant DP190102904.

"It is easy to show that man is actually guilty of many *inner lies*, but it seems more difficult to explain how they are possible; for a lie requires a second person whom one intends to deceive, whereas to deceive oneself on purpose seems to contain a contradiction."
Kant (1996, p.183)

1 Introduction

Empirical evidence from psychology and behavioral economics indicates that people's beliefs are often shaped by their preferences. People make use of motivated reasoning, processing information in a way that aligns with their preferences. They suffer from wishful thinking, believing in the truth or likely occurrence of what they want to be true or wish to occur. More generally, they engage in *self-deception*, misleading themselves about their situation. The literature in psychology traditionally attributes self-deception to the desire to maintain a more positive self-image (see, e.g., Kunda, 1990; Baumeister, 1998). Self-deception, however, is likely to induce people to venture into situations that may impose large costs on them; e.g., fights, life threatening challenges, etc. The existence of these costs suggests that a tendency towards making inaccurate inferences about oneself could have been selected by evolution over time only if it comes with some material – and not just psychological – benefits. The biologist Robert Trivers (1976) famously proposed one such possible benefit: self-deception may be useful to deceive others if people are more able to convince others when they truly hold the belief they want to convey.[1] Recent empirical evidence supports the suggestion that self-deception may be beneficial in social interactions (Anderson et al., 2012; Schwardmann and Van der Weele, 2019), by making individuals appear more competent and persuasive.

In this paper, we provide a rational micro-foundation for this strategic role of self-deception by showing that it emerges as a way to gain an advantage in competitive interactions. We postulate that the outcome – win or lose – of a competitive interaction between two parties depends on the strength, or type, of one of the parties, Player 1.

[1] "If [...] deceit is fundamental to animal communication, then there must be strong selection to spot deception and this ought, in turn, to select for a degree of self-deception, rendering some facts and motives unconscious so as not to betray – by the subtle signs of self-knowledge – the deception being practiced."

Player 1 and Player 2 have a common prior belief about Player 1's type, but Player 1 may gather private information *covertly*, before deciding whether to participate in the competitive interaction. That is, he can strategically select his *self-belief*, constrained by Bayes' law and the cost of gathering information. While Player 2 does not observe Player 1's self-belief, the action of competing or quitting by Player 1 can serve as a signal of Player 1's belief about his own strength. If Player 1 competes, then Player 2 may accept or refuse to compete. If both players engage in competition, then Player 1 wins if strong and loses if weak. Both agents prefer to win rather than quit and prefer to quit rather than lose.

This competitive interaction game nests many games of economic importance. For example, in a contest, Player 1, an adversary who may discover upon investigation whether he is strong or weak relative to the opponent, can decide to attack the opponent who then must respond by deciding to retreat or stand their ground. In a market entry game, an entrant who may be either efficient or inefficient, and needs to perform market research to find out, can enter a market where the incumbent firm would respond either by fighting or accommodating entry. Lastly, a prosecutor who may study the evidence to discover whether he is dealing with a strong or a weak case can offer a plea deal to a defendant who responds by accepting or rejecting the deal. In all these cases, Player 1 would need to choose the depth with which to assess his own strength before deciding what action to take.

We model information gathering by Player 1 about himself as a continuous flow of independent signals (drift diffusion process). Given Player 1's initial self-beliefs and the cost of information, he might decide to gather information. If so, he does it until his belief reaches either a lower or an upper *stopping belief*. At the lower stopping belief, Player 1 has become sufficiently pessimistic and quits. At the upper stopping belief, Player 1 decides to compete as he has become sufficiently confident in his ability to profit from competing. By competing, Player 1 sends a costly signal to Player 2 and induces Player 2 to back off with positive probability.[2] Thus, as suggested in Trivers (1976), Player 1 is able to convince Player 2 of his strength from truly holding confident beliefs that he is likely to be strong and acts upon such self-beliefs.

Even if the cost of information vanishes, Player 1's equilibrium upper stopping belief does not change. This is because, when correctly forecast by Player 2, this is the

[2] In the equilibrium with information gathering Player 2 uses a mixed strategy.

only belief that makes Player 2 mix between competing and not competing. As the cost of information decreases, Player 1 can afford to increase the chance of reaching the upper stopping belief by lowering the lower threshold beliefs where he quits. As the lower stopping threshold decreases, Player 1 can spend more time continuing to search for information after receiving initially negative evidence. That way, Player 1 gives himself a greater chance of eventually reaching the upper threshold. Player 1's strategy to form confident self-beliefs is based on an asymmetric stopping rule. As the flow cost of gathering information vanishes, more and more negative evidence (certainty in the limit) is required for Player 1 to quit, while the amount of positive evidence that is sufficient for Player 1 to compete does not change.

We view the fixed upper stopping belief at which Player 1 decides to compete as a manifestation of a Bayesian form of self-deception as defined by Sobel (2020), according to whom deception occurs when an informed sender's message is not as informative to the receiver as it could be. In our setting, Player 1 acts as both the sender and the receiver by both choosing the information-gathering strategy and using its realization to decide whether to compete. Player 1 engages in self-deception because, even when the cost of information is arbitrarily small, he chooses to limit the precision of the signal about his own strength. Moreover, the precision of information is limited in such a way that leads Player 1 to be *overconfident* in his ability to win, in the sense that his posterior belief leads him to compete more often than if he was fully informed.

Related Literature

This paper adds to our understanding of overconfidence as a form of motivated belief. Three types of explanations have been proposed for the tendency to form overconfident beliefs.

The answer most explored in the literature is that overconfidence stems from preferences over beliefs, either in the form of wishful thinking (Brunnermeier and Parker, 2005; Caplin and Leahy, 2019)—the enjoyment of having high expectations—or in the form of ego utility (Köszegi, 2006)—the enjoyment of thinking we are better than we are.

Another possibility is that overconfident beliefs produce objective benefits (not just subjective ones) by helping to attain higher achievements, either by compensating for a lack of willpower, in the form of present bias (Bénabou and Tirole, 2002), or by increasing the level of performance directly (Compte and Postlewaite, 2004).

A third possibility is that overconfidence is generated by self-deception because it provides benefits in social interactions, making it easier to convince others of one's worth or strength (Trivers, 1976; Von Hippel and Trivers, 2011). This idea has recently gained attention in economics, with several empirical studies finding that self-deception and the emergence of overconfident beliefs appear to be influenced by the strategic incentives faced by the agent (Burks et al., 2013; Schwardmann and Van der Weele, 2019; Solda et al., 2020; Schwardmann et al., 2022).

We contribute to this literature in several critical ways. First, we show theoretically that such benefits can indeed exist in strategic interactions between rational players. Second, compared to previous explanations of overconfidence, we show that overconfidence can emerge as a result of an agent's strategic behavior, without imposing exogenous costs or benefits of having certain beliefs, such as ego utility or anticipatory utility. Instead, the costs and benefits of overconfident beliefs emerge endogenously from the strategic incentives that agents face in competitive interactions. Third, we show that this strategic overconfidence can emerge between strategic agents that do not deviate from standard assumptions of rationality. Our result does not require an irremediable pre-existing behavioral bias, such as present bias, that overconfidence would help compensate. In addition, our explanation of overconfidence does not require agents to violate Bayesian rationality, echoing the findings in non-strategic settings of Bénabou and Tirole (2002) and Köszegi (2006). This explanation aligns with the fact that many behavioral patterns labeled as overconfidence, such as the observation that more than 50% of the population believes they are above the median, can emerge from Bayesian updating (Benoît and Dubra, 2011).

Our paper also contributes to the growing literature on optimal dynamic information gathering (Wald, 1992; Arrow et al., 1949). In our model, Player 1 can sequentially sample signals over time at a cost. This setting is widely used in neuroscience to model how the human brain collects information and forms beliefs (Gold and Shadlen, 2002; Bogacz et al., 2006) and is increasingly used in the economics literature to model information gathering (Fudenberg et al., 2018; Webb, 2019; Fudenberg et al., 2020; Henry and Ottaviani, 2019; McClellan, 2022). Recent papers have in particular explored optimal dynamic information gathering in the context of drug trials (Henry and Ottaviani, 2019; McClellan, 2022) to influence the view of other players. A key difference in our study is that the realizations of the information process are not public and can describe internal processes, like evaluating arguments and counterarguments, that need not be verifiable to others. Our paper emphasizes how strategic incentives influence such dynamic private information search, leading to an asymmetric process that favors the development of confident self-beliefs observed in experiments. For example, studies have found that individuals often stop looking for information asymmetrically, being more likely to accept good news, and continue to search for information after bad news (Ditto and Lopez, 1992; Solda et al., 2020), that they are more likely to accept positive news than negative news about their ability (Möbius et al., 2022; Zimmermann, 2020; Hagenbach and Saucet, 2024), and more likely to remember positive news than negative news (Chew et al., 2020).

Finally, our paper is related to the signaling literature (Crawford and Sobel, 1982), and the literature on Bayesian persuasion pioneered by Kamenica and Gentzkow (2011). In a setting of costless information gathering, Kamenica and Gentzkow (2011) demonstrated how a sender can design an optimal information structure to influence the beliefs and actions of a Bayesian receiver. A key assumption in their model, and much of the subsequent literature, is that the sender can commit to a publicly verifiable information disclosure policy before learning the state. The plausibility of this assumption has been questioned in many real-world applications (Bergemann and Morris, 2019; Kamenica, 2019). Recognizing this issue, recent studies have attempted to relax it (Fréchette et al., 2022; Min, 2021). Our paper takes a different approach. In our setting Player 1 does not "send" any information to Player 2 and therefore does not to commit to a verifiable information disclosure policy. Rather, he privately gathers information and, like in the standard signaling model, only "communicates" with Player 2 via the choice of an action. Nevertheless, we show that Player 1 can still persuade Player 2 by first gathering information privately, and then taking an action that credibly signals his posterior beliefs.

The remainder of the paper is organized as follows. In the next section, we introduce the model of the competitive interaction and in Section 2.2 we present the information gathering technology. In Section 3 we derive the equilibrium of our model. We show in Section 4 that our result can be viewed as self-deception—inducing overconfidence when competing—in an intra-personal game. In Section 5, we discuss the extension of our results to the cases of public information gathering and general, costless, information

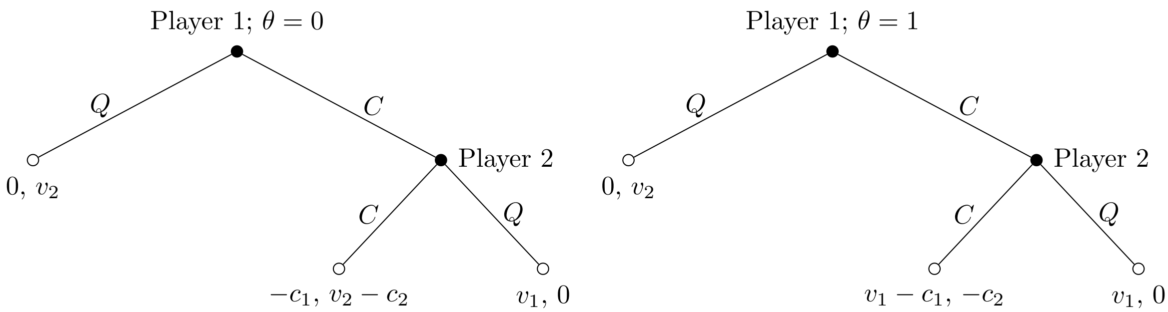

Figure 1: The Competitive Interaction Game

gathering. Section 6 concludes. All proofs are in the Appendix.

2 Model

2.1 The Competitive Interaction

There are two agents, *Player 1* and *Player 2*. Player 1 decides whether to engage in a competitive interaction with Player 2 by choosing between actions C (compete) and Q (quit). If Player 1 chooses C, Player 2 observes this and must then choose between C (compete) and Q (quit).

If Player 1 does not engage in the competitive interaction (chooses Q), then his payoff is zero, while the payoff of Player 2 is the winning prize $v_2 > 0$. If Player 1 competes (chooses C) and upon observing such a choice Player 2 quits (chooses Q), then the payoff of Player 1 is the winning prize $v_1 > 0$, while the payoff of Player 2 is zero. If, after Player 1 has chosen C, Player 2 decides to engage in the competitive interaction (chooses C), then the agents' payoffs depend on a binary state $\theta \in \{0, 1\}$, which can be interpreted as Player 1's ability or competitive strength. If Player 1 is weak, $\theta = 0$, then he loses the competition and his payoff is $-c_1$, while Player 2 wins and obtains the payoff $v_2 - c_2 > 0$; the parameter $c_i > 0$ can be interpreted as Player i's cost of engaging in the competitive interaction with the other agent. If Player 1 is strong, $\theta = 1$, then he wins the competition and his payoff is $v_1 - c_1 > 0$, while Player 2 loses and obtains the payoff $-c_2$. The game tree for each type θ is depicted in Figure 1.

Simply put, each player prefers to win rather than quit and prefers to quit rather than lose.[3] If there is an interaction, then Player 1 wins if strong, and loses if weak. Under complete information, the subgame perfect equilibrium is: when strong ($\theta = 1$), Player 1 competes and wins while Player 2 quits; when weak ($\theta = 0$), Player 1 quits while Player 2 competes and wins.

The agents are expected utility maximizers. Denote Player i's belief that $\theta = 1$ by $\mu_i \in [0, 1]$. When both compete, Player 1's expected payoff $\mu_1 v_1 - c_1$ increases in μ_1 while Player 2's expected payoff $(1 - \mu_2)v_2 - c_2$ decreases in μ_2.

When Player 2 chooses C, Player 1 is indifferent between C and Q if and only if:

$$\mu_1^* v_1 - c_1 = 0 \qquad \text{or} \qquad \mu_1^* = \frac{c_1}{v_1}$$

and prefers C over Q for $\mu_1 > \frac{c_1}{v_1}$. Similarly, when Player 1 chooses C, Player 2 is indifferent between C and Q if and only if:

$$(1 - \mu_2^*)v_2 - c_2 = 0 \qquad \text{or} \qquad \mu_2^* = 1 - \frac{c_2}{v_2}$$

and prefers C over Q for $\mu_2 < 1 - \frac{c_2}{v_2}$.

Before describing the process of information gathering by Player 1, we consider as a benchmark the equilibria of the competitive interaction game when players share a common prior that Player 1 is strong, $\mu^0 = \Pr(\theta = 1)$. The following proposition, whose proof is elementary and hence omitted, presents the equilibrium of this game for different values of the common prior.

Proposition 1. *The Bayesian Nash equilibrium of the competitive interaction is* (C, Q) *if* $\mu^0 > \mu_2^*$, (Q, C) *if* $\mu^0 < \mu_2^*$ *and* $\mu^0 < \mu_1^*$, *and* (C, C) *if* $\mu_1^* < \mu^0 < \mu_2^*$.

When $\mu^0 > \mu_2^*$, Player 1 has no incentive to acquire additional information, as he wins the prize v_1 at no cost. When $\mu^0 < \mu_1^* < \mu_2^*$, Player 1 has an incentive to acquire information about his strength, but not to the level that would deter Player 2 from competing.[4] The case of interest is therefore $\mu^0 < \mu_2^* < \mu_1^*$. In this case, Player 1 may benefit from acquiring private information about his strength, signal this strength by

[3] Our main results hold without requiring a fixed prize for winning and fixed cost for fighting for each player. In general, a sufficient requirement is for each player to strictly gain from winning and strictly lose from losing the competition relative to quitting.

[4] If $\mu^0 < \mu_1^* < \mu_2^*$, then, for any given strategy of Player 2, Player 1 would not acquire costly private information beyond μ_1^*; therefore in any equilibrium of the game described in Section 3, Player 2 would always choose C.

engaging in the competitive interaction and get Player 2 to quit at least some of the time. Therefore, for the remainder of the paper, we make the following assumption to restrict attention to this case of interest:

Assumption 1.

$$\mu^0 < \mu_2^* < \mu_1^* \tag{1}$$

The assumption ranks the threshold beliefs that Player 1 is strong. It says that the threshold belief μ_2^* above which Player 2 would quit if the opponent competed is lower than the threshold belief μ_1^* below which Player 1 would quit if the opponent competed. In the special case of symmetric prizes and costs, $v_1 = v_2 = v$, $c_1 = c_2 = c$, $\mu_2^* < \mu_1^*$ reduces to $c > \frac{v}{2}$; the cost of engaging in the interaction with a competing opponent is high relative to the winning prize.

2.2 Information Gathering

Prior to the competitive interaction, Player 1 can strategically shape his self-belief by gathering information privately through a sequence of costly experiments as in Wald (1945), which we model as taking place in continuous time. Gathering information incurs a flow cost of γ. As long as Player 1 continues paying the flow cost, information arrives according to a Brownian stochastic process B_t with variance σ^2, positive drift ρ in state $\theta = 1$, and negative drift $-\rho$ in state $\theta = 0$. The value of the process at time t is:

$$B_t = \rho(2\theta - 1)t + \sigma W_t, \quad \sigma^2 > 0, \quad \rho > 0, \tag{2}$$

where W_t is a Wiener process (a zero mean, unit variance, standard Brownian motion) that is mutually independent of Player 1's type θ. The variance σ^2 and the drift ρ capture the informativeness of the experiments conducted by Player 1. The smaller σ^2 and the larger ρ, the more informative the experiments' signals are.

We refer to a path of the information process up to time t as a realization (or outcome) at time t and denote it by $\mathcal{B}_t = (B_s)_{s \leq t}$; we refer to $\mathcal{B} = (B_t)_{t \geq 0}$ as an outcome of the information process. Prior to gathering information, both players share common prior about Player 1's strength, $\mu^0 = \Pr(\theta = 1)$. The posterior belief of Player 1 conditional on a realization at time t of the information process, obtained via Bayesian updating, is denoted by $\mu_{1,t}(\mathcal{B}_t) = \Pr(\theta = 1 | \mathcal{B}_t)$.

Player 1 must decide when to stop gathering information and what action to take at that stopping point as a function of the time and realized information flow. That is, Player 1 selects a decision rule $\delta = (\tau, q_1)$ consisting of a stopping time τ which maps outcomes $\mathcal{B}$ of the information process into a time $t \in [0, \infty)$ when collecting information stops, and an action rule q_1 that maps outcomes of the information process to a probability $q_1 \in [0, 1]$ of choosing C (and a complementary probability $1 - q_1$ of choosing Q). The stopping time τ satisfies the property that if $\tau(\mathcal{B}) = t$, then $\tau(\mathcal{B}') = t$ for all $\mathcal{B}'$ such that $\mathcal{B}'_t = \mathcal{B}_t$; that is, the decision to stop collecting information at time t can only depend on information collected up to time t. The action rule q_1 is measurable on the σ-algebra induced by the τ function and can be written as $q_1(\mathcal{B}) = q_1(\mathcal{B}_{\tau(\mathcal{B})}) \in [0, 1]$); that is, the choice of an action only depends on the path of the information process up to the time $\tau(\mathcal{B})$ when collecting information stops. We call a decision rule *adaptable* when it satisfies these conditions.

While we assume that parameters of the information gathering process are common knowledge, the information process is privately observed by Player 1. Player 2 only sees the action taken by Player 1, not observing either the realization of the information process or the stopping time. Therefore, Player 2's strategy simply consists of a potentially mixed strategy between choosing to compete or quit. We denote by $q_2 \in [0, 1]$ the probability that Player 2 chooses C after observing that Player 1 has chosen C. We denote by μ_2 Player 2's belief that $\theta = 1$ after observing that Player 1 has chosen C.[5]

2.3 Perfect Bayesian Equilibrium

We study the perfect Bayesian equilibria (PBE) of the game consisting of the strategic information gathering by Player 1, followed by the competitive interaction. A PBE is an adaptable decision rule $\delta = (\tau, q_1)$ and a posterior belief function $\mu_{1,t}$ of Player 1, and a strategy q_2 and a belief μ_2 of Player 2 that satisfy Bayesian updating – with Player 1 using the information gathered and Player 2 using the equilibrium decision rule of Player 1 – and sequential rationality. We denote the expected utility of player 1 given action $A_1 \in \{C, Q\}$, belief μ_1 and strategy of player 2 q_2 as $u_1(A_1, \mu_1, q_2)$. We denote the expected utility of player 2 given action $A_2 \in \{C, Q\}$ belief μ_2, and $A_1 = C$

[5] We don't need notation for Player 2's belief when Player 1 chooses Q, as it plays no role in the analysis.

as $u_2(A_2, \mu_2)$. Formally, a PBE is $\langle (\tau, q_1), \mu_{1,t};\, q_2, \mu_2 \rangle$ where the following conditions hold.

1. Given a conjecture about player 2's strategy, $\hat{q}_2$, player 1 chooses a decision rule and has belief function where

   (a) for each $\mathcal{B}_t$ the action rule only puts positive weight on actions that maximize expected utility: $\operatorname{supp} q_1(\mathcal{B}_t) \subseteq \arg\max_{A_1 \in \{C,Q\}} u_1(A_1, \mu_{1,t}(\mathcal{B}_t), \hat{q}_2)$,

   (b) the stopping rule satisfies

   $$\tau(\hat{q}_2) = \inf\{t \geq 0 : V(t) = u_1(q_1(\mathcal{B}_t), \mu_{1,t}(\mathcal{B}_t), \hat{q}_2)\}$$

   where $V(t) = \sup_{(\tau', q_1')} \mathbb{E}[u_1(q_1'(\mathcal{B}_t), \mu_{1,\tau'}(\mathcal{B}_{\tau'}), \hat{q}_2)]$, and

   (c) the belief function satisfies Bayesian updating $\mu_{1,t}(\mathcal{B}_t) = \Pr(\theta = 1 | \mathcal{B}_t)$.

2. Given a conjecture about Player 1's decision rule $\hat{\delta} = (\hat{\tau}, \hat{q}_1)$ Player 2 chooses a strategy and has belief where

   (a) the strategy only puts positive weight on actions that maximize expected utility: $\operatorname{supp} q_2(\mu_2) \subseteq \arg\max_{A_2 \in \{C,Q\}} u_2(A_2, \mu_2)$ and

   (b) beliefs satisfy Bayesian updating: $\mu_2 = \mathbb{E}[\Pr(\theta = 1 | \mathcal{B}_{\tau(\mathcal{B})}) | C, \hat{\delta}]$.

3. Conjectures match play: $\hat{\delta} = \delta$ and $\hat{q}_2 = q_2$.

3 Equilibrium Strategies

For any fixed strategy of Player 2, Player 1's payoff depends only on his own action and the realization of his type $\theta \in \{0, 1\}$. The expected payoff of Player 1 from choosing action C is a decreasing function of the probability q_2 that Player 2 competes, and for $\underline{q}_2 = \frac{v_1}{v_1 + c_1} \in (0, 1)$ it is $u_1(C, 0, \underline{q}_2) = u_1(Q, 0, \underline{q}_2)$, or equivalently $(1 - \underline{q}_2)v_1 - \underline{q}_2 c_1 = 0$. Furthermore, $u_1(C, 1, q_2) > u_1(Q, 1, q_2)$ for any $q_2 \in [0, 1]$. Then for any $q_2 \leq \underline{q}_2$, Player 1 prefers to compete regardless of θ. For $q_2 > \underline{q}_2$ Player 1 prefers to compete when $\theta = 1$ and prefers to quit when $\theta = 0$. The following results state that in any PBE the equilibrium strategy of Player 2 must satisfy $q_2 > \underline{q}_2$.

Lemma 1. *In any PBE,* $q_2 > \underline{q}_2 = \frac{v_1}{v_1+c_1}$*; Player 2 plays Compete with sufficiently high probability so that* $u_1(Q,0,q_2) > u_1(C,0,q_2)$ *and* $u_1(C,1,q_2) > u_1(Q,1,q_2)$.

Proof. In the Appendix. □

Lemma 1 implies that in a PBE Player 1's decision rule follows that of Wald's Sequential Probability Ratio Test (SPRT); that is, as shown in the next proposition, for any given conjectured strategy $\hat{q}_2$ of Player 2, we can represent Player 1's optimal choice of a stopping time τ as defined by two stopping beliefs $\mu_1^L(\hat{q}_2)$ and $\mu_1^H(\hat{q}_2)$, with $\mu_1^L(\hat{q}_2) < \mu_1^H(\hat{q}_2)$. When reaching one of these posterior beliefs, Player 1 stops gathering information about his type. Intuitively, as a function of the cost and benefit of gathering additional information, when the information process leads to a belief $\mu_1^L(\hat{q}_2)$, Player 1 is "sufficiently" confident that he is weak, given Player 2's conjectured strategy $\hat{q}_2$, and quits, while when the information process leads to a belief $\mu_1^H(\hat{q}_2)$, Player 1 is "sufficiently" confident that he is strong and competes.

Proposition 2 (Shiryaev (2007), Theorem 4.5, pg 185)**.** *For any conjectured strategy* $\hat{q}_2$ *of Player 2, Player 1's best response decision rule is a solution to the continuous time Wald SPRT, contains two stopping belief thresholds,* $\mu_1^L(\hat{q}_2)$ *and* $\mu_1^H(\hat{q}_2)$*, and has the following properties*

(i) stops gathering information and chooses Q *if* $\mu_{1,t}(\mathcal{B}_t) \leq \mu_1^L(\hat{q}_2)$*,*

(ii) continues gathering information if $\mu_1^L(\hat{q}_2) < \mu_{1,t}(\mathcal{B}_t) < \mu_1^H(\hat{q}_2)$*, and*

(iii) stops gathering information and chooses C *if* $\mu_1^H(\hat{q}_2) \leq \mu_{i,t}(\mathcal{B}_t)$*,*

Proposition 2 highlights that Player 1's best response stopping thresholds do not depend on the shared prior belief μ^0 or the time taken to reach a particular posterior belief. As an immediate consequence, if $\mu^0 \notin (\mu_1^L(\hat{q}_2), \mu_1^H(\hat{q}_2))$, Player 1's optimal decision rule will be to immediately stop gathering information and make the choice described in Proposition 2. We next turn to characterizing the PBE of the game. There are two types of equilibria, those with information gathering and those without.

3.1 Equilibria without Information Gathering

In an equilibrium without information gathering, Player 1's belief about his strength is the prior belief μ^0. Hence, by Assumption 1 there is no equilibrium where Player 1 competes with positive probability. As Player 2 gets to decide whether to compete only off the equilibrium path, there is some degree of freedom in selecting Player 2's action and beliefs that can be part of a PBE. However, it must be the case that Player 2's strategy q_2 induces Player 1 not to gather information and to quit; that is, $\mu_1^L(q_2) \geq \mu^0$. It must also be the case that Player 2's belief is sufficiently optimistic to choose C off the equilibrium path: $\mu_2 \leq \mu_2^*$. Note that, if $\mu_2 < \mu_2^*$, then sequential rationality (off the equilibrium path) requires that $q_2 = 1$. These outcome equivalent equilibria are formalized in the following proposition.

Proposition 3. *In any PBE with no information gathering, i.e., $\tau(\mathcal{B}) = 0$ for all $\mathcal{B}$, Player 1's action choice is $q_1(\mathcal{B}) = 0$ for all $\mathcal{B}$ and Player 2's action choice is such that $\mu_1^L(q_2) \geq \mu^0$. Beliefs of the players satisfy $\mu_{1,0} = \mu^0$ and $\mu_2 \leq \mu_2^*$ (with $q_2 = 1$ if $\mu_2 < \mu_2^*$).*

Proof. In the Appendix. □

Intuitively, as we shall see, these equilibria prevail if and only if the cost of gathering information is high. The cost of acquiring a "unit" of information can be summarized with the term $\omega = \frac{\gamma\sigma^2}{2\rho^2}$.[6] All else equal, $\mu_1^L(q_2)$ decreases as ω reduces for any $q_2 > \underline{q}_2$. It follows that for a given μ^0 and sufficiently small information cost, $\mu_1^L(q_2) < \mu^0$. In the following section we characterize the unique equilibrium with information gathering. It corresponds to the unique equilibrium for low information costs.

3.2 Equilibrium with Information Gathering

We now consider equilibria with information gathering. In these equilibria, Player 2's strategy is such that $\mu_1^L(q_2) < \mu^0 < \mu_1^H(q_2)$ where the thresholds indicate the posteriors where Player 1 stops acquiring information as defined in Proposition 2.

Relative to Wald's classic SPRT, there is an additional choice variable to be determined in equilibrium, namely Player 2's strategy q_2. As outlined in Proposition 2, for a given q_2 unique thresholds for a stopping rule are determined. As q_2 changes, so too does the pair of thresholds. In order to pin down the equilibrium with information gathering we use the fact that Player 2 will be using a mixed strategy ($q_2 \in (0, 1)$).

[6] This cost is proportional to the cost of a trial (the flow cost γ) and the variance σ^2 in the information process. It is inversely proportional to the square of the drift ρ towards the true state θ.

To see this, note first that, by Lemma 1, it must be $q_2 > 0$. Second, by the sequential rationality of Player 1's strategy, $q_2 = 1$ can only be part of an equilibrium with information gathering if $\mu_1^H(1) \geq \mu_1^*$; Player 1 with belief $\mu_1^H(1)$ must find it optimal to compete. On the equilibrium path, Player 2 must have the same beliefs as Player 1; that is, in this case, $\mu_2 = \mu_1^H(1)$. Since by Assumption 1 it is $\mu_2^* < \mu_1^*$, it follows that $q_2 = 1$ is not sequentially rational. Player 2 with belief $\mu_2 = \mu_1^H(1) \geq \mu_1^*$ would find it optimal to choose to quit, i.e., $q_2 = 0$.

Because Player 2 is using a mixed strategy $q_2 \in (0, 1)$ in equilibrium, it must be that $\mu_1^H(q_2)$ is equal to μ_2^*. In other words the mixing probability q_2 of Player 2 must make it optimal for Player 1 to chooses $\mu_2^H(q_2) = \mu_2^*$. Thus, instead of the two belief thresholds being pinned down by Player 1's optimal decision rule, it is the lower threshold and the mixed strategy of Player 2 that are pinned down by this decision rule.

The following proposition characterizes the properties of the equilibrium with information gathering and shows that there is at most one such equilibrium.

Proposition 4. *In a perfect Bayesian equilibrium with information gathering, player 2's strategy q_2^* is such that $\mu_1^L(q_2^*) < \mu^0$ and $\mu_1^H(q_2^*) = \mu_2^*$. The lower threshold and player 2's strategy are the unique solution of the following conditions:*

$$c_1 q_2^* = v_1 - \frac{\gamma\sigma^2}{2\rho^2}\left(\frac{1-\mu_1^L}{\mu_1^L} - \ln\frac{\mu_1^L}{1-\mu_1^L} - \frac{1-\mu_2^*}{\mu_2^*} + \ln\frac{\mu_2^*}{1-\mu_2^*}\right) \tag{3}$$

$$v_1 q_2^* = \frac{\gamma\sigma^2}{2\rho^2}\left(\frac{1-\mu_1^L}{\mu_1^L} - \frac{\mu_1^L}{1-\mu_1^L} - \frac{1-\mu_2^*}{\mu_2^*} + \frac{\mu_2^*}{1-\mu_2^*}\right) + \frac{\gamma\sigma^2}{\rho^2}\left(\ln\frac{\mu_2^*}{1-\mu_2^*} - \ln\frac{\mu_1^L}{1-\mu_1^L}\right) \tag{4}$$

After observing that Player 1 has chosen C, Player 2's belief is $\mu_2 = \mu_2^*$.

Proof. In the Appendix. □

While the equilibria described in Proposition 3 and Proposition 4 represent all possible equilibria of the game, they do not preclude one another. For example, for a single set of parameters, there can be an equilibrium with information gathering where q_2^* satisfies the conditions of Proposition 4 and an equilibrium with no information gathering where $q_2 = 1$. This would happen whenever $\mu_1^L(q_2^*) < \mu^0 < \mu_1^L(1)$. The next result shows that at least one type of equilibrium exists for any set of parameters.

Corollary 1. *For any set of payoffs, information costs and prior that satisfy Assumption 1, there exists a perfect Bayesian equilibrium.*

Proof. In the Appendix. □

As we shall see in the next proposition, for a given prior μ^0, there is a sufficiently low unit cost of information that guarantees a unique equilibrium with information gathering.[7] The proposition also shows that as the unit information cost ω converges to zero, the lower stopping beliefs converges to $\mu_1^L = 0$ and the equilibrium strategy q_2^* of Player 2 converges to $\underline{q}_2 = \frac{v_1}{v_1+c_1}$.

Proposition 5. *For all prior beliefs $\mu^0 < \mu_2^*$, there is a threshold cost of a unit of information $\overline{\omega}$ such that if $\omega < \overline{\omega}$, then Player 1 gathers information in the unique equilibrium.*

As the cost of a unit of information ω decreases and converges to zero:

- *The lower stopping belief $\mu_1^L(q_2^*)$ decreases monotonically and converges to zero;*
- *The strategy q_2^* of Player 2 converges to $\underline{q}_2 = \frac{v_1}{v_1+c_1}$;*
- *The expected payoff of Player 1 converges to $\frac{\mu^0 v_1^2}{v_1+c_1}$.*

Proof. In the Appendix. □

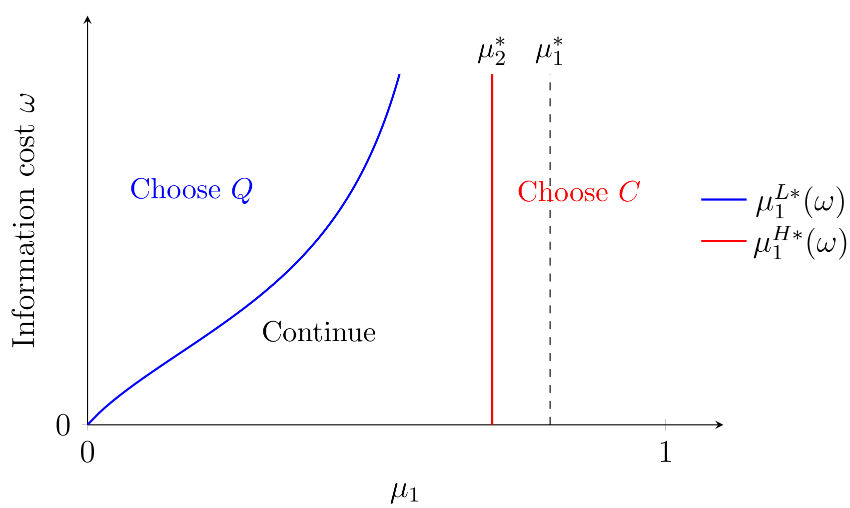

Figure 2: Depiction of Player 1's equilibrium decision rule when gathering information for different levels of information cost. Parameters are $\mu_2^* = 0.7$, $v_1 = 1$ and $c_1 = 0.75$.

Figure 2 shows how the equilibrium belief thresholds change with the cost of a unit of information. As this cost decreases, the lower belief threshold decreases while

[7] At the same time, for a given unit cost of information if the prior μ^0 is sufficiently sufficiently close to μ_2^*, then Player 1 wants to gather information.

the upper belief threshold remains constant and leading to assymmetric belief thresholds. As information cost vanishes, the lower belief threshold approaches zero and the probability Player 1 reaches the upper threshold in an equilibrium with information gathering approaches $\frac{\mu^0}{\mu_2^*}$. Player 1's equilibrium information gathering strategy will lead to him choosing C with probability that exceeds μ^0, the ex-ante likelihood he is strong. In this limit, Player 1 will never choose Q when he is strong, but with positive probability chooses C when he is weak. In the following section we discuss the connections between this equilibrium the ideas of self-deception and overconfidence.

4 Self-Deception and Overconfidence

We now argue that the strategic choice of beliefs by Player 1 can be characterized as "self-deception" when using Sobel (2020)'s definition of deception in an intra-personal game taking place between two selves of Player 1: a planner and a doer. In Sobel's metaphor, the planner-self observes the state at no cost and chooses a messaging strategy to communicate information about the state to the doer-self. The doer-self observes the message of the planner self and chooses an action, which in our setting is either to compete or quit. To connect equilibrium behavior in our setting to this intra-personal game we use Proposition 5 and focus on the limit equilibrium strategy of Player 1 as the unit cost ω converges to zero.[8]

The messaging strategy of the planner-self is determined from our setting by Player 1's choice of stopping beliefs μ_1^L, μ_1^H and the associated probability of reaching each belief for each state, denoted by $P^{\theta}_{\mu^0}(\mu_1^L)$ and $P^{\theta}_{\mu^0}(\mu_1^H)$ for $\theta \in \{0,1\}$. When $\mu_1^L < \mu_1^H$, the planner-self sends the doer-self a high message with probability $P^{\theta}_{\mu^0}(\mu_1^H)$ and a low message with complementary probability $P^{\theta}_{\mu^0}(\mu_1^L)$ when the state is θ. In a full-disclosure strategy, the stopping beliefs of Player 1 are $\mu_1^L = 0$, which is reached with probability one when $\theta = 0$, and $\mu_1^H = 1$, which is reached with probability one when $\theta = 1$. In this strategy, the messaging strategy of the planned self is to send a high message when $\theta = 1$ and a low message when $\theta = 0$. At the other extreme, when $\mu_1^L = \mu_1^H = \mu^0$ as in the no information gathering equilibria, the planner-self always

[8] In Section 5.3, we consider an extension to costless information gathering with no restriction on the information technology and show that the limit equilibrium of Proposition 5 remains an equilibrium in that setting.

sends the same message for all values of θ; that is, he transmits no information to the doer-self.

Following Sobel (2020)'s definition, deception occurs when the sender's message induces beliefs which are further from the truth than those that would have been induced by another available message. A messaging strategy that sends a deceptive message with positive probability is referred to as a deceptive strategy. In our setting, the planner-self's message is deceptive in an equilibria with information gathering whenever the information process hits the high threshold when $\theta = 0$, or whenever it hits the low threshold when $\theta = 1$. The definition of a deceptive strategy in our setting and our use of self-deception is given below.

Definition 1 (Self-Deception).

(a) *A messaging strategy determined from stopping beliefs μ_1^L, μ_1^H is deceptive when $\mu_1^L < \mu_1^H$ and either $P^1_{\mu^0}(\mu_1^L) > 0$ or $P^0_{\mu^0}(\mu_1^H) > 0$.*[9]

(b) *Player 1 engages in self-deception if his planner-self chooses a messaging strategy that is deceptive.*

In the equilibrium of our model, the doer-self of Player 1 decides to quit if the stopping belief reached is μ_1^L and to compete if it is μ_1^H. Hence, it is natural to think that when choosing a deceptive messaging strategy, the planner-self of Player 1 wants to influence the doer-self's beliefs about his ability to win the competition. When $\mu_1^L < \mu_1^H$ and $P^0_{\mu^0}(\mu_1^H) > 0$ the planner-self sends the high message with positive probability when $\theta = 0$. In this case, the message increases the doer-self's beliefs about his ability and induces him to compete despite the fact that he is weak. We refer to this message as one that generates overconfidence in the doer-self of Player 1.

Definition 2 (Overconfidence).

(a) *The doer-self of Player 1 is overconfident about his ability to win when $\mu_1 > 0$ and $\theta = 0$.*

(b) *A messaging strategy determined from stopping beliefs μ_1^L, μ_1^H generates overconfidence when $\mu_1^L < \mu_1^H$ and $P^0_{\mu^0}(\mu_1^H) > 0$.*

[9]By (6) and (7) in the proof of Proposition 4, an equivalent definition of a deceptive messaging is when the stopping rules are such that either $0 < \mu_1^L < \mu_1^H$ or $\mu_1^L < \mu_1^H < 1$.

The next proposition shows that all messaging strategies determined by the stopping beliefs of Player 1's information gathering strategy, apart from disclosing no information and disclosing full information, are deceptive. It follows that the limit equilibrium strategy of Player 1 described in Proposition 5 corresponds to a message strategy that is self-deceptive. More precisely, the planner-self of Player 1 deceives the doer-self by generating overconfidence in his ability to win with probability $(1-\mu^0)P^0_{\mu^0}(\mu_1^H)$. On the other hand, in the limit $\mu_1^L = 0$ and $P^1_{\mu^0}(\mu_1^L) = 0$; that is, the planner-self does not induce underconfidence in the doer-self.[10]

Proposition 6.

(a) Any messaging strategy other than the full-disclosure and the no-disclosure strategy is deceptive.

(b) In the limit PBE obtained from $\omega \to 0$, Player 1 engages in self-deception, with the planner-self generating overconfidence in the doer-self.

Proof. In the Appendix. □

In our competitive setting, an information gathering strategy that uses self-deceptive messages generating overconfidence is valuable to Player 1. With no information gathering the doer-self always quits and Player 1 gets a zero payoff; the self-deceptive messaging strategy, on the contrary, yields a positive payoff to Player 1. Relative to the full information messaging strategy, the self-deceptive messaging strategy increases the likelihood that the doer-self competes. In particular, Player 1 is more likely to compete than he is to be strong.

5 Extensions

5.1 Publicly Observable Information Gathering

When signal realizations are publicly observable, the belief μ_2 of Player 2 directly depends on the information that Player 1 gathers. At each time t, $\mu_{2,t}(\mathcal{B}_t) = \mu_{1,t}(\mathcal{B}_t) = \Pr(\theta = 1 \mid \mathcal{B}_t)$. Thus, the choice of Player 1 to compete no longer influences the beliefs of

[10] One can define underconfidence in a similar way to overconfidence: Player 1 is underconfident when $\mu_1 < 1$ and $\theta = 1$. A messaging strategy that generates underconfidence would have $\mu_1^L < \mu_1^H$ and $P^1_{\mu^0}(\mu_1^L) > 0$.

Player 2. Sequential rationality implies that $q_2 = 1$ if the information gathering process stops at a belief $\mu_2 = \mu_1 < \mu_2^*$ and $q_2 = 0$ if $\mu_2 > \mu_2^*$. Player 2 is indifferent between competing and quitting at beliefs μ_2^*, but if he competes with positive probability Player 1 would want to move Player 2's beliefs above μ_2^* by an arbitrarily small amount so as to induce him to quit. Therefore, in equilibrium it must be $q_2 = 0$ and $q_1 = 1$ for $\mu_1 = \mu_2 \geq \mu_2^*$, and $q_1 = 0, q_2 = 1$ otherwise. Player 1 chooses the same upper stopping belief, $\mu_1^H = \mu_2^*$.

Since Player 2 quits with probability one when the players beliefs are $\mu_1 = \mu_2 = \mu_2^*$ and Player 1 competes, it is then immediate that by choosing the same information policy as when information is private, Player 1 obtains a higher payoff (strictly higher as long as the flow cost of gathering information is not too high to make gathering no information optimal).

In general, as argued in the next proposition, Player 1 can do even better by gathering more information; that is, by reducing the lower stopping beliefs μ_1^L.

Proposition 7. *If the information signals are also observed by Player 2 (rather than only by Player 1), then in the subgame perfect equilibrium the threshold information cost below which Player 1 gathers information is lower. If Player 1 gathers information in equilibrium, then the lower stopping belief μ_1^L is closer to zero, the upper stopping belief is the same, $\mu_1^H = \mu_2^*$, and Player 1's expected payoff is higher.*

Proof. In the Appendix. □

Proposition 7 shows that the ability to commit to disclose information publicly is, unsurprisingly, valuable. An important message of this paper, as shown by Proposition 5, is that when information gathering is followed by an action that acts as a costly signal of this information as described by Spence (1973), neither commitment power nor costless verifiability of information is required for strategic information acquisition to be beneficial. In any PBE with private information gathering and $q_2 > 0$, Player 1 can credibly signal his belief about his strength through the choice of competing before the opponent reacts, because his expected payoff from competing is increasing in this belief.[11]

[11]We should note that, unlike when it is public, when information is private, it is important that Player 1 chooses his action first and that such a choice is observed by Player 2 and serves as a signal of Player 1's beliefs about his strength. By Assumption 1, if Player 2 moved first, or if the two

5.2 Ex-Ante Information Cost

It is well known (e.g., see Morris and Strack, 2019) that there is a correspondence between Wald's sequential sampling model of information gathering and a static model in which information can be gathered at a cost that depends on the Bayesian feasible posteriors one wishes to obtain. In our setting, we can see from equation (5) in the proof of Proposition 4 that the payoff function of Player 1 can be written as the expected benefit of reaching the upper belief threshold μ_1^H and competing minus the ex-ante expected information cost of reaching a belief distribution with positive mass on the two beliefs μ_1^L and μ_1^H. See footnote 12 in the Appendix.

5.3 Costless Information Gathering

We have already studied what happens in our model if the cost of gathering information goes to zero. In this section, we model gathering information as being costless and, importantly, impose no restriction – besides Bayes plausibility – on the information gathering technology. We show that the limit equilibrium of our costly information gathering model remains an equilibrium of the general costless information gathering model.

Assume that gathering information is costless and consider general information gathering policies. Let S be a countable set of possible signal realizations s. Denote by $\pi : \{0,1\} \to S$ a feasible information policy; for all $s \in S$ and $\theta \in \{0,1\}$, $\pi(s \,|\, \theta)$ denotes the probability that Player 1 observes signal s when the state is θ and π is the chosen information policy. Given signal realization $s \in S$, Player 1 updates his beliefs about the state using Bayes' rule. We denote Player 1's belief that $\theta = 1$ given signal s by $\mu_1(s) = \frac{\pi(s|1)\mu^0}{\pi(s|1)\mu^0+\pi(s|0)(1-\mu^0)}$. The information policy π induces a probability distribution over the posterior beliefs μ denoted by $\phi : [0,1] \to [0,1]$, where $\phi(\mu) = \sum_{s:\mu_1(s)=\mu} \pi(s|1)\mu^0 + \pi(s|0)(1-\mu^0)$. The feasible information policies are those

players moved simultaneously, then Player 2 would choose to compete. Depending on the flow cost of gathering information, Player 1 might or might not gather information, but if Player 1 gathers information the higher stopping belief μ_1^H must be strictly higher than μ_1^*, the belief that makes Player 1 indifferent between competing and quitting when Player 2 competes.

In the limit, as the cost of information converges to zero, Player 1's information policy converges to gathering full information and Player 1's expected payoff converges to $\mu^0(v_1 - c_1)$, which is lower than the expected payoff when Player 1 chooses his action first. Player 1 benefits from being able to signal his type through the choice of competing.

that satisfy Bayes plausibility; that is, the π that induce a ϕ such that $\mathbb{E}_\phi[\mu] = \mu^0$.

The next proposition shows that there is a continuum of payoff equivalent, mixed strategy PBEs.

Proposition 8. *For each $\mu_1^H \in [\mu_2^*, 1]$ there is a PBE characterized by the following:*

1. *An information policy with distribution of posterior beliefs such that:*
   - *Posterior belief $\mu_1^L = 0$ has probability $\phi(0) = \frac{\mu_1^H - \mu^0}{\mu_1^H} > 0$;*
   - *Conditional on $\mu_1 > 0$, the expected posterior belief is $\mathbb{E}_\phi[\mu_1|\mu_1 > 0] = \frac{\mu^0}{1-\phi(0)} = \mu_1^H$.*
2. *Player 1 competes with probability:*
   - *$q_1 = 1 - \frac{\mu_2^* - \mu^0}{\phi(0)\mu_2^*}$ if the posterior belief is $\mu_1^L = 0$;*
   - *$q_1 = 1$ if the posterior belief is $\mu_1 > 0$.*
3. *Player 2 competes with probability $q_2 = \frac{v_1}{v_1 + c_1}$.*
4. *The expected payoffs are $\frac{\mu^0 v_1^2}{v_1 + c_1}$ for Player 1, and $v_2 \frac{v_2(1-\mu^0) - c_2}{v_2 - c_2}$ for Player 2.*

Proof. In the Appendix. □

The least informative PBE, with $\mu_1^H = \mu_2^*$, corresponds to the limit equilibrium of our model with costly information gathering when the unit cost of information ω goes to zero. In such a PBE, after observing that Player 1 competes, the belief of Player 2 is $\mathbb{E}_\phi[\mu_1|\mu_1 > 0] = \mu_2^*$. This highlights an additional advantage of our costly information gathering process. Even if we think that in reality personal information gathering is not very costly: our Wald information gathering process approach allows us to select among all the PBEs of the model with costless, unrestricted (except for Bayes plausibility) information gathering technology.

In the most informative PBE, $\mu_1^H = 1$ and Player 1 uses the full information messaging strategy, but in all other equilibria, one for each $\mu_1^H \in (\mu_2^*, 1)$, the planner-self of Player 1 uses a self-deceptive messaging strategy that induces over-confidence in the doer-self with positive probability. The higher is the equilibrium value of μ_1^H, the lower is the likelihood of self-deception and over-confidence. Self-deception is replaced by the doer-self of Player 1 competing with a positive probability when he knows he is weak (i.e., when $\mu_1 = 0$), that increases with the equilibrium value of μ_1^H. Because Player 2 interprets the action of competing (via Bayes' updating) as Player 1 being strong with probability $\mu_2^* > \mu^0$, we can interpret the choice of C by Player 1 when he knows that he is weak as a form of deception. Under this interpretation, in this costless information setting with no restrictions on the information gathering process, self-deception and deception are substitutes. Self-deception alone can provide Player 1 with the same benefits he would obtain with deception.

It is worth spelling out the difference between the costless information model of this section and Bayesian persuasion. In the standard model of Bayesian persuasion (Kamenica and Gentzkow, 2011), a Sender can influence a Receiver by choosing a costless information policy with publicly verifiable signal realizations. In many real-world settings, however, information is gathered privately by the Sender who then must communicate the relevant information to the Receiver. In these settings, the Sender does not have the commitment power to always reveal the realization of the information policy and moreover, this realization may not be easily verifiable by the Receiver. One might have conjectured that persuasion, or the ability to increase his own payoff, is impossible for the Sender without such commitment. This paper shows that this is not so. Player 1 can also persuade Player 2 when information can only be gathered privately and covertly, by selecting a less than fully informative belief and signaling it through the choice of an action. Not surprisingly, as shown by Proposition 7 persuasion is more effective (raises the Sender's payoff more) when the Sender can design the information observed by the Receiver, than when he can only signal his belief through his actions.

6 Concluding Remarks

We have developed a model of strategic self-deception in a competitive interaction between two rational agents. Our model provides a novel perspective on motivated beliefs and related biases, suggesting that they may arise as rational responses to the strategic incentives present in many social and economic interactions. This is in contrast to the traditional view that self-deception and overconfidence are purely psychological phenomena, driven by a taste for positive self-views or a failure to process information in a Bayesian manner. We have shown that even in the absence of any intrinsic preference for positive self-views, individuals may optimally choose to distort their beliefs and acquire less than fully revealing information about their own strength, in order to gain a strategic advantage over their opponents. This is because confident self-beliefs can lead to take actions that signal strength and influence the behavior of others in advantageous ways.

At the same time, our model also highlights the limits of strategic self-deception. In particular, we have shown that if the prior probability that Player 1 is strong is too low, or if the cost of gathering information is too high, then it may not be worthwhile for Player 1 to engage in self-deception at all. This suggests that strategic self-deception is more likely to occur in environments where the potential gains from influencing others' behavior are large relative to the costs of gathering self-confidence enhancing information.

Our framework could be extended to study other forms of strategic information gathering, such as the selective recall of past experiences or the biased interpretation of feedback. Another possibility is that individuals may receive self-confidence enhancing feedback from others, such as parents, teachers, or managers. This type of *paternalistic validation* can be modeled as an external agent providing Player 1 with a noisy signal about his strength. Our analysis also has potentially important implications for the design of organizations and institutions. In particular, it suggests that in competitive settings where the strategic value of self-confidence is high, there may be a role for policies or practices that facilitate the gathering of confidence-enhancing information. This could include the provision of motivational feedback, the celebration of success stories, or the use of confidence-building exercises in training and development programs. Having said that, we should not ignore the potential risks of such practices, as they may lead to overconfidence and excessive risk-taking if not properly calibrated.

References

Anderson, C., Brion, S., Moore, D. A., and Kennedy, J. A. (2012). A status-enhancement account of overconfidence. *Journal of Personality and Social Psychology*, 103(4):718.

Arrow, K. J., Blackwell, D., and Girshick, M. A. (1949). Bayes and minimax solutions of sequential decision problems. *Econometrica*, pages 213–244.

Baumeister, R. F. (1998). The self. *The Handbook of Social Psychology*, 1:680–740.

Bénabou, R. and Tirole, J. (2002). Self-confidence and personal motivation. *The Quarterly Journal of Economics*, 117(3):871–915.

Benoît, J.-P. and Dubra, J. (2011). Apparent overconfidence. *Econometrica*, 79(5):1591–1625.

Bergemann, D. and Morris, S. (2019). Information design: A unified perspective. *Journal of Economic Literature*, 57(1):44–95.

Bogacz, R., Brown, E., Moehlis, J., Holmes, P., and Cohen, J. D. (2006). The physics of optimal decision making: a formal analysis of models of performance in two-alternative forced-choice tasks. *Psychological Review*, 113(4):700.

Brunnermeier, M. K. and Parker, J. A. (2005). Optimal expectations. *American Economic Review*, 95(4):1092–1118.

Burks, S. V., Carpenter, J. P., Goette, L., and Rustichini, A. (2013). Overconfidence and social signalling. *Review of Economic Studies*, 80(3):949–983.

Caplin, A. and Leahy, J. V. (2019). Wishful thinking. Technical report, National Bureau of Economic Research.

Chew, S. H., Huang, W., and Zhao, X. (2020). Motivated false memory. *Journal of Political Economy*, 128(10):3913–3939.

Compte, O. and Postlewaite, A. (2004). Confidence-enhanced performance. *American Economic Review*, 94(5):1536–1557.

Crawford, V. P. and Sobel, J. (1982). Strategic information transmission. *Econometrica*, pages 1431–1451.

Ditto, P. H. and Lopez, D. F. (1992). Motivated skepticism: Use of differential decision criteria for preferred and nonpreferred conclusions. *Journal of Personality and Social Psychology*, 63(4):568–584.

Fréchette, G. R., Lizzeri, A., and Perego, J. (2022). Rules and commitment in communication: An experimental analysis. *Econometrica*, 90(5):2283–2318.

Fudenberg, D., Newey, W., Strack, P., and Strzalecki, T. (2020). Testing the drift-diffusion model. *Proceedings of the National Academy of Sciences*, 117(52):33141–33148.

Fudenberg, D., Strack, P., and Strzalecki, T. (2018). Speed, accuracy, and the optimal timing of choices. *American Economic Review*, 108(12):3651–84.

Gold, J. I. and Shadlen, M. N. (2002). Banburismus and the brain: decoding the relationship between sensory stimuli, decisions, and reward. *Neuron*, 36(2):299–308.

Hagenbach, J. and Saucet, C. (2024). Motivated skepticism. *Review of Economic Studies*, May(rdae058).

Henry, E. and Ottaviani, M. (2019). Research and the approval process: The organization of persuasion. *American Economic Review*, 109(3):911–55.

Kamenica, E. (2019). Bayesian persuasion and information design. *Annual Review of Economics*, 11:249–272.

Kamenica, E. and Gentzkow, M. (2011). Bayesian persuasion. *American Economic Review*, 101(6):2590–2615.

Kant, I. (1996). *The Metaphysics of Morals.* (M.J. Gregor, Trans.). Cambridge University Press. (Originally published 1797).

Köszegi, B. (2006). Ego utility, overconfidence, and task choice. *Journal of the European Economic Association*, 4(4):673–707.

Kunda, Z. (1990). The case for motivated reasoning. *Psychological Bulletin*, 108(3):480–498.

Lipster, R. S. and Shiryayev, A. N. (1977). *Statistics of Random Processes I: General Theory.* Springer-Verlag.

McClellan, A. (2022). Experimentation and approval mechanisms. *Econometrica*, 90(5):2215–2247.

Min, D. (2021). Bayesian persuasion under partial commitment. *Economic Theory*, 72(3):743–764.

Möbius, M. M., Niederle, M., Niehaus, P., and Rosenblat, T. S. (2022). Belief updating and the demand for information. *Experimental Economics*, pages 1–32.

Morris, S. and Strack, P. (2019). The wald problem and the relation of sequential sampling and ex-ante information costs. Technical report, Available at SSRN: https://ssrn.com/abstract=2991567 or http://dx.doi.org/10.2139/ssrn.2991567.

Schwardmann, P., Tripodi, E., and Van der Weele, J. J. (2022). Self-persuasion: Evidence from field experiments at international debating competitions. *American Economic Review*, 112(4):1118–1146.

Schwardmann, P. and Van der Weele, J. (2019). Deception and self-deception. *Nature human behaviour*, 3(10):1055–1061.

Shiryaev, A. N. (2007). *Optimal Stopping Rules*, volume 8. Springer Science & Business Media.

Sobel, J. (2020). Lying and deception in games. *Journal of Political Economy*, 128(3):907–947.

Solda, A., Ke, C., Page, L., and Von Hippel, W. (2020). Strategically delusional. *Experimental Economics*, 23(3):604–631.

Spence, M. (1973). Job market signaling. *The Quarterly Journal of Economics*, 87(3):355–374.

Trivers, R. (1976). Foreword to r. dawkins,"the selfish gene"(pp. xiv-xx).

Von Hippel, W. and Trivers, R. (2011). The evolution and psychology of self-deception. *Behavioral and Brain Sciences*, 34(1):1–16.

Wald, A. (1945). Sequential tests of statistical hypotheses. *The Annals of Mathematical Statistics*, 16:117–186.

Wald, A. (1992). Sequential tests of statistical hypotheses. In *Breakthroughs in Statistics: Foundations and Basic Theory*, pages 256–298. Springer.

Webb, R. (2019). The (neural) dynamics of stochastic choice. *Management Science*, 65(1):230–255.

Zimmermann, F. (2020). The dynamics of motivated beliefs. *American Economic Review*, 110(2):337–363.

Appendix

Proof of Lemma 1

First, note that it follows directly from the definitions of the payoffs of the game that $u_1(C, 1, q_2) > u_1(Q, 1, q_2)$ holds for any $q_2 \in [0, 1]$, since the inequality is equivalent to $v_1 - q_2 c_1 > 0$. Second, suppose, contrary to the lemma's claim, that $q_2 \leq \underline{q}_2 < 1$. Then $u_1(Q, 0, q_2) \leq u_1(C, 0, q_2)$ and since $u_1(Q, 1, q_2) < u_1(C, 1, q_2)$ for all q_2, Player 1 would prefer to choose C over Q for any $\mu_1 \in (0, 1)$ regardless of the state, and would hence stop gathering information at $t = 0$. It follows that in equilibrium $\mu_2 = \mu^0$. By Assumption 1, $\mu^0 < \mu_2^*$ and therefore $u_2(C, \mu^0) > u_2(Q, \mu^0)$, which implies that it is optimal for Player 2 to choose $q_2 = 1$, contradicting the supposition that $q_2 < 1$. This implies that in any PBE $q_2 > \underline{q}_2 = \frac{v_1}{v_1 + c_1}$ and, as a consequence, $u_1(Q, 0, q_2) = 0 > u_1(C, 0, q_2) = (1 - q_2)v_1 - q_2 c_1$. □

Proof of Proposition 3

Let $\langle(\tau^*, q_1^*), \mu_{1,t}^*; q_2^*, \mu_2^*\rangle$ be a PBE where $\tau^*(\mathcal{B}) = 0$ for all $\mathcal{B}$. From Proposition 2, for Player 1 to not gather any information, either $\mu^0 > \mu_1^H(q_2^*)$ or $\mu^0 < \mu_1^L(q_2^*)$. Suppose first that it is the case that $\mu^0 > \mu_1^H(q_2^*)$. Then by Proposition 2 Player 1 competes, $q_1(\mathcal{B}) = 1$ and Player 2's must choose whether to compete on the equilibrium path. When making such a choice Player 2 must correctly anticipate that Player 1 did not gather information and thus have beliefs equal to the prior beliefs, $\mu_2 = \mu^0$. By Assumption 1, $\mu^0 < \mu_2^*$ and therefore it must be that $q_2^* = 1$. However, by Assumption 1 $\mu_1^* > \mu_2^*$, meaning that $q_1^* = 1$ cannot be a best response action rule for Player 1. Therefore, there can be no PBE with $\mu^0 > \mu_1^H(q_2^*)$.

Second, suppose that $\mu^0 < \mu_1^L(q_2^*)$. Then by Proposition 2, player 1 quits, $q_1(\mathcal{B}) = 0$. Because Player 1 never chooses C on the equilibrium path, then μ_2^* is not disciplined by Bayes' rule. There is an equilibrium in this case exactly when $\mu_2 \leq \mu_2^*$ and q_2^* is such that $\mu^0 < \mu_1^L(q_2^*)$. If $\mu_2 < \mu_2^*$, then sequential rationality off the equilibrium path requires that $q_2 = 1$. □

Proof of Proposition 4

To find the stopping thresholds and the probability q_2 that Player 2 chooses C that are part of an equilibrium, it is convenient to re-parametrize these cutoffs and players' equilibrium beliefs with log-odds belief ratios:

$$\lambda^0 = \ln \tfrac{\mu^0}{1-\mu^0}; \quad \lambda_t = \ln \tfrac{\mu_{1,t}}{1-\mu_{1,t}}; \quad \lambda_1^L = \ln \tfrac{\mu_1^L}{1-\mu_1^L}; \quad \lambda_1^H = \ln \tfrac{\mu_1^H}{1-\mu_1^H}; \quad \lambda_2^* = \ln \tfrac{\mu_2^*}{1-\mu_2^*}.$$

As shown in the next lemma, it follows from B_t being a Brownian motion that λ_t also follows a Brownian motion.

Lemma 2. *The log-odds process follows a Brownian motion with variance $\frac{4\rho^2}{\sigma^2}$, positive drift $\frac{2\rho^2}{\sigma^2}$ if $\theta = 1$, and negative drift $-\frac{2\rho^2}{\sigma^2}$ if $\theta = 0$:*

$$\lambda_t = \lambda^0 + \tfrac{2\rho}{\sigma^2} B_t = \lambda^0 + \tfrac{2\rho^2}{\sigma^2}(2\theta - 1)t + \tfrac{2\rho}{\sigma} W_t.$$

Proof. Consider the belief process λ_t generated by the discovery process B_t. Given a path $\mathcal{B}_t$, let $P^1(\mathcal{B}_t) = \Pr(\mathcal{B}_t \mid \theta = 1)$, $P^0(\mathcal{B}_t) = \Pr(\mathcal{B}_t \mid \theta = 0)$, and $P^W(\mathcal{B}_t) = \Pr(\mathcal{B}_t \mid \theta = 1/2)$. The latter is the probability of observing the path if the stochastic process were a driftless Weiner process with variance σ. By Theorem 7.1, equation (7.5), in Lipster and Shiryayev (1977) the Radon-Nykodim derivative of the measures P^1 and P^0 with respect to P^W are:

$$\frac{dP^W}{dP^0} = \exp\left\{\frac{\rho}{\sigma^2} B_t + \frac{1}{2}\frac{\rho^2}{\sigma^2}\right\}; \qquad \frac{dP^W}{dP^1} = \exp\left\{-\frac{\rho}{\sigma^2} B_t + \frac{1}{2}\frac{\rho^2}{\sigma^2}\right\}.$$

Hence we can write the likelihood ratio as:

$$\frac{dP^1}{dP^0} = \exp\left\{\frac{2\rho}{\sigma^2} B_t\right\}.$$

By Bayes' rule, the posterior log-odds λ_t are equal to:

$$\lambda_t = \lambda^0 + \frac{2\rho}{\sigma^2} B_t = \lambda^0 + \frac{2\rho^2}{\sigma^2}(2\theta - 1)t + \frac{2\rho}{\sigma} W_t.$$

The log-odds process follows a Brownian motion with variance $\frac{4\rho^2}{\sigma^2}$, positive drift $\frac{2\rho^2}{\sigma^2}$ if $\theta = 1$, and negative drift $-\frac{2\rho^2}{\sigma^2}$ if $\theta = 0$. □

We now use this result to complete the proof of Proposition 4. In the PBE, Player 1 chooses log-odds stopping thresholds λ_1^L and λ_1^H to maximize his expected payoff given Player 2's strategy of playing C with probability q_2. We use the following definitions:

1. $\mathbb{E}^{\theta}_{\lambda^0}[\tau; \lambda_1^L, \lambda_1^H]$ is the expected time the belief process first reaches either λ_1^L or λ_1^H starting from λ^0, conditional on $\theta \in \{0, 1\}$;

2. $P^{\theta}_{\lambda^0}(\lambda_1^H) = P^{\theta}_{\mu^0}(\mu_1^H)$ is the probability that the belief process starting at λ^0 stops by reaching λ_1^H, conditional on $\theta \in \{0, 1\}$.

Player 1's expected payoff at his prior belief $\mu^0 = \frac{e^{\lambda^0}}{1+e^{\lambda^0}}$ can be written as follows:[12]

$$\begin{aligned} V(\lambda^0; \lambda_1^L, \lambda_1^H) =& \frac{e^{\lambda^0}}{1+e^{\lambda^0}}\Big[P^1_{\lambda^0}(\lambda_1^H) u_1(q_2; \mu_1^H) - \gamma \mathbb{E}^1_{\lambda^0}[\tau; \lambda_1^L, \lambda_1^H]\Big] \\ &+ \frac{1}{1+e^{\lambda^0}}\Big[P^0_{\lambda^0}(\lambda_1^H) u_1(q_2; \mu_1^H) - \gamma \mathbb{E}^0_{\lambda^0}[\tau; \lambda_1^L, \lambda_1^H]\Big] \end{aligned} \tag{5}$$

The following formulas are well known (see Stokey, 2009, equation (5.16) in Proposition 5.4 and equation (5.23)).

$$P^1_{\lambda^0}(\lambda_1^H) = \frac{e^{-\lambda_1^L} - e^{-\lambda^0}}{e^{-\lambda_1^L} - e^{-\lambda_1^H}} = P^1_{\mu^0}(\mu_1^H) = \frac{1 - \frac{\mu_1^L}{\mu^0} \cdot \frac{1-\mu^0}{1-\mu_1^L}}{1 - \frac{\mu_1^L}{\mu_1^H} \cdot \frac{1-\mu_1^H}{1-\mu_1^L}} \tag{6}$$

$$P^0_{\lambda^0}(\lambda_1^H) = \frac{e^{\lambda^0} - e^{\lambda_1^L}}{e^{\lambda_1^H} - e^{\lambda_1^L}} = P^0_{\mu^0}(\mu_1^H) = \frac{\frac{\mu^0}{\mu_1^L} \cdot \frac{1-\mu_1^L}{1-\mu^0} - 1}{\frac{\mu_1^H}{\mu_1^L} \cdot \frac{1-\mu_1^L}{1-\mu_1^H} - 1} \tag{7}$$

$$\mathbb{E}^1_{\lambda^0}[\tau; \lambda_1^L, \lambda_1^H] = \frac{\sigma^2}{2\rho^2} \cdot \frac{e^{-\lambda_1^L}(\lambda_1^H - \lambda^0) + e^{-\lambda_1^H}(\lambda^0 - \lambda_1^L) - e^{-\lambda^0}(\lambda_1^H - \lambda_1^L)}{e^{-\lambda_1^L} - e^{-\lambda_1^H}}$$

$$\mathbb{E}^0_{\lambda^0}[\tau; \lambda_1^L, \lambda_1^H] = \frac{\sigma^2}{2\rho^2} \cdot \frac{e^{\lambda_1^L}(\lambda_1^H - \lambda^0) + e^{\lambda_1^H}(\lambda^0 - \lambda_1^L) - e^{\lambda^0}(\lambda_1^H - \lambda_1^L)}{e^{\lambda_1^H} - e^{\lambda_1^L}}$$

[12] We can think of the payoff of Player 1 as the expected benefit from reaching belief μ_1^H and competing minus the ex-ante information cost of ending either with belief μ_1^L or μ_1^H:

$$C(\mu_1^L, \mu_1^H) = \gamma\left[\mu^0 \mathbb{E}^1_{\lambda^0}[\tau; \lambda_1^L, \lambda_1^H] + (1-\mu^0)\mathbb{E}^0_{\lambda^0}[\tau; \lambda_1^L, \lambda_1^H]\right]$$

Player 1's expected payoff (5) thus can be written as:

$$
\begin{aligned}
V(\lambda^0;\lambda_1^L,\lambda_1^H) &= \tfrac{e^{\lambda^0}}{1+e^{\lambda^0}}\left[\tfrac{u_1(q_2;\mu_1^H)\left(e^{-\lambda_1^L}-e^{-\lambda^0}\right)-\frac{\gamma\sigma^2}{2\rho^2}\left[e^{-\lambda_1^L}(\lambda_1^H-\lambda^0)+e^{-\lambda_1^H}(\lambda^0-\lambda_1^L)-e^{-\lambda^0}(\lambda_1^H-\lambda_1^L)\right]}{e^{-\lambda_1^L}-e^{-\lambda_1^H}}\right] \\
&\quad+\tfrac{1}{1+e^{\lambda^0}}\left[\tfrac{u_1(q_2;\mu_1^H)\left(e^{\lambda^0}-e^{\lambda_1^L}\right)-\frac{\gamma\sigma^2}{2\rho^2}\left[e^{\lambda_1^L}(\lambda_1^H-\lambda^0)+e^{\lambda_1^H}(\lambda^0-\lambda_1^L)-e^{\lambda^0}(\lambda_1^H-\lambda_1^L)\right]}{e^{\lambda_1^H}-e^{\lambda_1^L}}\right] \\
&= \tfrac{1}{1+e^{\lambda^0}}\left[\tfrac{u_1(q_2;\mu_1^H)e^{\lambda_1^H}\left(e^{\lambda^0}-e^{\lambda_1^L}\right)-\frac{\gamma\sigma^2}{2\rho^2}\left[e^{\lambda^0}e^{\lambda_1^H}(\lambda_1^H-\lambda^0)+e^{\lambda^0}e^{\lambda_1^L}(\lambda^0-\lambda_1^L)-e^{\lambda_1^L}e^{\lambda_1^H}(\lambda_1^H-\lambda_1^L)\right]}{e^{\lambda_1^H}-e^{\lambda_1^L}}\right] \\
&\quad+\tfrac{1}{1+e^{\lambda^0}}\left[\tfrac{u_1(q_2;\mu_1^H)\left(e^{\lambda^0}-e^{\lambda_1^L}\right)-\frac{\gamma\sigma^2}{2\rho^2}\left[e^{\lambda_1^L}(\lambda_1^H-\lambda^0)+e^{\lambda_1^H}(\lambda^0-\lambda_1^L)-e^{\lambda^0}(\lambda_1^H-\lambda_1^L)\right]}{e^{\lambda_1^H}-e^{\lambda_1^L}}\right] \\
&= \tfrac{1}{\left(1+e^{\lambda^0}\right)\left(e^{\lambda_1^H}-e^{\lambda_1^L}\right)}\Bigg\{u_1(q_2;\mu_1^H)(1+e^{\lambda_1^H})\left(e^{\lambda^0}-e^{\lambda_1^L}\right) \\
&\quad-\tfrac{\gamma\sigma^2}{2\rho^2}\left[\left(e^{\lambda_1^L}+e^{\lambda^0+\lambda_1^H}\right)(\lambda_1^H-\lambda^0)+\left(e^{\lambda_1^H}+e^{\lambda^0+\lambda_1^L}\right)(\lambda^0-\lambda_1^L)-\left(e^{\lambda^0}+e^{\lambda_1^L+\lambda_1^H}\right)(\lambda_1^H-\lambda_1^L)\right]\Bigg\}
\end{aligned}
$$

To lighten the notation, it is convenient to do a change of variables. Let $y^0=e^{\lambda^0}, y^L=e^{\lambda_1^L}, y^H=e^{\lambda_1^H}, \omega=\frac{\gamma\sigma^2}{2\rho^2}$, and note that $u_1(q_2;\mu_1^H)=q_2v_1\frac{y^H}{1+y^H}+v_1-q_2(v_1+c_1)$. Then, we can write the expected payoff as:

$$
\begin{aligned}
V(y^0;y^L,y^H) &= \tfrac{1}{1+y^0}\Bigg\{q_2v_1y^H\left(\tfrac{y^0-y^L}{y^H-y^L}\right)+\left[v_1-q_2(v_1+c_1)\right](1+y^H)\left(\tfrac{y^0-y^L}{y^H-y^L}\right) \\
&\quad+\omega\tfrac{(y^0-y^L)(1-y^H)}{y^H-y^L}\ln y^H-\omega(1-y^0)\ln y^0+\omega\tfrac{(y^H-y^0)(1-y^L)}{y^H-y^L}\ln y^L\Bigg\}
\end{aligned}
\tag{8}
$$

We can then represent Player 1's problem as one of choosing y^L, y^H to maximize his expected profit subject to the constraints: $y^L \leq y^0$ and $y^H \geq y^0$. We solve the problem by finding what the solution would be if we ignored the constraints. If the constraints are not satisfied by the solution, then one of them must bind, which implies that Player 1 will not gather costly information.

After some algebra simplification, the first order conditions for Player 1 (unconstrained) maximization problem can be written as follows:

$$\frac{\partial V}{\partial y^H} = \frac{y^0 - y^L}{(1+y^0)(y^H - y^L)^2}\Big\{ - q_2 v_1 y^L - [v_1 - q_2(v_1 + c_1)](1 + y^L) \tag{9}$$
$$+\omega\left(\frac{1-y^H}{y^H}\right)\left(y^H - y^L\right) - \omega\left(1 - y^L\right)\left(\ln y^H - \ln y^L\right)\Big\} = 0$$
$$\frac{\partial V}{\partial y^L} = \frac{y^H - y^0}{(1+y^0)(y^H - y^L)^2}\Big\{ - q_2 v_1 y^H - [v_1 - q_2(v_1 + c_1)](1 + y^H) \tag{10}$$
$$+\omega\left(\frac{1-y^L}{y^L}\right)\left(y^H - y^L\right) - \omega\left(1 - y^H\right)\left(\ln y^H - \ln y^L\right)\Big\} = 0$$

We now verify that the second order conditions are satisfied at a solution. Note that the first order conditions imply that the two expressions in curly brackets are equal to zero. Using this, we can write the derivatives at the solution as follows:

$$\frac{\partial^2 V}{\partial (y^H)^2} = \frac{y^0 - y^L}{(1+y^0)(y^H - y^L)^2}\left\{ - \omega\left(\frac{1+y^H}{(y^H)^2}\right)\left(y^H - y^L\right) \right\} < 0$$

$$\frac{\partial^2 V}{\partial (y^L)^2} = \frac{y^H - y^0}{(1+y^0)(y^H - y^L)^2}\left\{ - \omega\left(\frac{1+y^L}{(y^L)^2}\right)\left(y^H - y^L\right) \right\} < 0$$

$$\frac{\partial^2 V}{\partial y^L \partial y^H} = \frac{y^0 - y^L}{(1+y^0)(y^H - y^L)^2}\left\{ - q_2 v_1 - [v_1 - q_2(v_1 + c_1)] - \omega\left(\frac{1-y^H}{y^H}\right) + \omega\left(\frac{1-y^L}{y^L}\right) + \omega\Big(\ln y^H - \ln y^L\Big) \right\} = 0$$

$$\frac{\partial^2 V}{\partial y^H \partial y^L} = \frac{y^H - y^0}{(1+y^0)(y^H - y^L)^2}\left\{ - q_2 v_1 - [v_1 - q_2(v_1 + c_1)] - \omega\left(\frac{1-y^H}{y^H}\right) + \omega\left(\frac{1-y^L}{y^L}\right) + \omega\Big(\ln y^H - \ln y^L\Big) \right\} = 0,$$

where the equalities in the last two expressions follow by noting that the term in the curly bracket of both expressions is equal to $y^H - y^L$ times the term obtained by subtracting the expression in the curly brackets of (9) from the expression in the curly brackets of (10), which are both zero at the solution of the first order conditions. Thus the second order conditions for Player 1's maximization problem are satisfied.

Noting that at an equilibrium with information gathering the two expressions in curly brackets in (9) and (10) must be equal to zero, we can rewrite the first order conditions as follows:

$$q_2 = \frac{v_1}{c_1} - \frac{\omega}{c_1}\left(\frac{1}{y^L} - \ln y^L - \frac{1}{y^H} + \ln y^H\right) \tag{11}$$
$$q_2 = \frac{\omega}{v_1}\left(\frac{1}{y^L} - y^L - \frac{1}{y^H} + y^H\right) + \frac{2\omega}{v_1}\left(\ln y^H - \ln y^L\right) \tag{12}$$

The first equality is obtained by subtracting the expression in curly brackets in (10) from the expression in curly brackets in (9) and rearranging. The second expression is obtained by replacing (11) in the expression in curly brackets in (10).

Denote by $q_2^A(y^L, y^H)$ the RHS of (11) and by $q_2^B(y^L, y^H)$ the RHS of (12). Then note that $q_2^A(y^L, y^H)$ is increasing in y^L and decreasing in y^H, while $q_2^B(y^L, y^H)$ is decreasing in y^L and increasing in y^H. As a consequence, $\Delta q_2(y^L, y^H) = q_2^A(y^L, y^H) - q_2^B(y^L, y^H)$ is increasing in y^L and decreasing in y^H. Moreover, taking $y^H \geq y^L$ as given, we have:

$$\lim_{y^L \to 0} \Delta q_2(y^L, y^H) < 0 \quad \text{and} \quad \lim_{y^L \to y^H} \Delta q_2(y^L, y^H) > 0$$

It follows that there is a unique $y^L \in (0, y^H)$ at which $\Delta q_2(y^L, y^H) = 0$.

By Proposition 2, upon observing that Player 1 has chosen C, Player 2 infers that the posterior belief of Player 1 is μ_1^H. As a consequence, Player 2's belief is μ_2 is also equal to μ_1^H. Since Player 2 must be indifferent between C and Q, it must be $\mu_1^H = \mu_2^*$ at an equilibrium with information gathering; that is, it must be $y^H = \frac{\mu_2^*}{1-\mu_2^*}$. We can thus conclude that the unique candidate stopping beliefs for an equilibrium with information gathering are (μ_1^L, μ_2^*) where μ_1^L is the unique solution to $\Delta q_2\left(\frac{\mu_1^L}{1-\mu_1^L}, \frac{\mu_2^*}{1-\mu_2^*}\right) = 0$. The first-order conditions in the proposition are obtained by changing variables again, with $y^H = \frac{\mu_1^H}{1-\mu_1^H}$, $y^L = \frac{\mu_1^L}{1-\mu_1^L}$, and $\omega = \frac{\gamma\sigma^2}{2\rho^2}$. □

Proof of Corollary 1

For any ω, v_1, c_1, and μ_2^* there is a unique pair $(q_2^*, \mu_1^L(q_2^*))$ that solve (3) and (4). If $\mu_1^L(q_2^*) < \mu^0$, the conditions for existence of a PBE with information gathering as characterized in Proposition 4 hold. If $\mu_1^L(q_2^*) \geq \mu^0$, the conditions are satisfied for the existence of a PBE with no information gathering as characterized in Proposition 3. □

Proof of Proposition 5

Adding up equations (11) and (12) yields:

$$\frac{v_1}{c_1} = \left[\frac{v_1 + c_1}{v_1 c_1}\left(\frac{1}{y^L} - \ln y^L - \frac{1}{y^H} + \ln y^H\right) + \frac{1}{v_1}\left(\ln y^H - \ln y^L + y^H - y^L\right)\right]\omega \quad (13)$$

and totally differentiating with respect to y^L and ω shows that y^L, and hence μ^L, is monotonically increasing in ω (recall, y^H does not change with ω):

$$0 = \frac{v_1}{c_1 \omega} d\omega - \left[\frac{\omega(v_1 + c_1)}{v_1 c_1} \left(\frac{1}{\left(y^L\right)^2} + \frac{1}{y^L} \right) + \frac{\omega}{v_1} \left(\frac{1}{y^L} + 1 \right) \right] dy^L$$

Since, by (11) and (12), y^L and q_2 are continuous function of ω , it follows that there exists $\overline{\omega}$ such that Player 1 gathers information if $\omega < \overline{\omega}$.

Equation (13) implies that for all values of ω it is $y^L < y^H$ and $\lim_{\omega \to \infty} y^L = y^H$. It follows that for all ω there exists $\overline{y}$, and hence $\overline{\mu}$, such that if the prior belief ratio is $y^0 \in (\overline{y}, y^H)$, and hence $\mu^0 \in (\overline{\mu}, \mu^H = \mu_2^*)$, then Player 1 gathers information in equilibrium.

As $\omega \to 0$, the RHS of (12) converges to

$$\frac{\omega}{v_1 y^L} - \frac{2\omega \ln y^L}{v_1}.$$

For this to be equal to q_2 it must be that $\lim_{\omega \to 0} y^L = \lim_{\omega \to 0} \mu_1^L = 0$.

Furthermore, since $\frac{1}{y^L}$ converges to $+\infty$ faster than $\ln y^L$ when $y^L \to 0$, by (12) it must be $\lim_{\omega \to 0} \frac{\omega}{v_1 y^L} = q_2$. Using this, we conclude that as $\omega \to 0$, the RHS of (11) converges to $\frac{v_1}{c_1} - \frac{v_1 q_2}{c_1}$. For this to be equal to q_2, it must be $\lim_{\omega \to 0} q_2 = \frac{v_1}{v_1 + c_1}$.

As $\omega \to 0$, $\mu_1^L \to 0$. Thus, the probability that Player 1 ends with belief $\mu_1^H = \mu_2^*$ is $\frac{\mu^0}{\mu_2^*}$ by Bayes law. It follows that Player 1 expected payoff is $\frac{\mu^0}{\mu_2^*} u_1 \left(\frac{v_1}{v_1 + c_1}; \mu_2^* \right) = \frac{\mu^0 v_1^2}{v_1 + c_1}$. □

Proof of Proposition 6

It is immediate that the full-disclosure strategy is not self-deceptive, as the stopping belief μ_1^L is never reached when $\theta = 1$ and the stopping belief μ_1^H is never reached when $\theta = 0$.

It is also immediate that the no-disclosure strategy is not self-deceptive, as only one stopping belief, the prior belief μ^0, is ever reached, irrespectively of the state θ.

Now suppose Player 1 adopts a messaging strategy that is informative, but not fully informative. Then $\mu_1^L < \mu_1^H$, and either $\mu_1^L > 0$ or $\mu_1^H < 1$. By (6) and (7), it follows that if $\mu_1^L > 0$, then with positive probability the messaging strategy chosen by the planner-self of Player 1 reaches the stopping belief μ_1^L (instead of μ_1^H), when the true state is $\theta = 1$. Similarly, if $\mu_1^H < 1$, then with positive probability the messaging strategy chosen by the planner-self of Player 1 reaches the stopping belief μ_1^H (instead of μ_1^L), when the true state is $\theta = 0$.

Since the PBE described in Proposition 5 is such that the messaging strategy of the planner-self of Player 1 has stopping beliefs $\mu_1^L = 0$ and $\mu_1^H = \mu_2^* < 1$, it follows that the messaging strategy is self-deceptive and that, with positive probability, it makes the doer-self of Player 1 overconfident in his ability to win. □

Proof of Proposition 7

As in the proof of Proposition 4, Player 1 chooses y^L and y^H to maximize his expected payoff. However, q_2 is now a function of y^H; it is equal to zero if $y^H \geq \frac{\mu_2^*}{1-\mu_2^*}$ and it is equal to one otherwise. It follows that the LHS $\frac{\partial V}{\partial y^H}$ of the first order condition (9) must be non-positive for $q_2 = 0$ and non-negative for $q_2 = 1$. On the other hand, the LHS $\frac{\partial V}{\partial y^L}$ of the first order condition (10) must be equal to zero for $q_2 = 0$. As the solution y^H remains equal to $\frac{\mu_2^*}{1-\mu_2^*}$, we can totally differentiate (10) with respect to y^L and q_2 to derive that y^L is an increasing function of q_2. Since when Player 2 observes the information gathered by Player 1 it is $q_2 = 0$, it follows that y^L and hence the stopping belief μ_1^L is lower. This in turn implies that the threshold information cost below which Player 1 gathers information is lower (as the interval of priors (μ_1^L, μ_1^H) for which Player 1 gathers information is larger). That Player 1's expected payoff is higher follows from the fact that if Player 1 competes, than Player 2 quits (rather than mix) and the fact that Player 1 competes with higher probability (because the threshold belief μ_1^L that induces Player 1 to quit is lower). □

Proof of Proposition 8

Note that, since Player 1's expected payoff from C is $(\mu_1 v_1 - c_1)q_2 + v_1(1 - q_2)$, the belief that makes Player 1 indifferent between choosing C and Q is [13]

$$\mu_1^{q_2} := \frac{v_1 + c_1}{v_1} - \frac{1}{q_2}. \tag{14}$$

We first show, by contradiction, that in any PBE, $q_2 = q_2^* \equiv \frac{v_1}{v_1+c_1}$.

[13] If $\mu_1^{q_2} < 0$, then Player 1 prefers C for all beliefs μ_1.

If $q_2 > q_2^*$, then by (14) the belief that makes Player 1 indifferent between choosing C and Q is $\mu_1^{q_2} > \mu_1^{q_2^*} = 0$ and Player 1 prefers Q and to obtain a zero payoff for all $0 \leq \mu_1 < \mu_1^{q_2}$ and to choose C and obtain a positive payoff for $\mu_1^{q_2} \leq \mu_1^* < 1$. The concave closure of $u(\mu_1; q_2)$ is the linear function $(v_1 - c_1 q_2)\mu_1$, which takes higher values than $u(\mu_1; q_2)$ everywhere on $(0, 1)$. Therefore, the optimal information policy for Player 1 is such that $\phi(\mu_1) > 0$ only for $\mu_1 \in \{0, 1\}$; that is, it is optimal for Player 1 to choose an information policy that fully reveals the state θ. Sequential rationality then implies that the choice of Player 1 must be $q_1 = 0$ if $\theta = 0$ and $q_1 = 1$ if $\theta = 1$, which in turn implies that it is sequentially rational for Player 2 to choose $q_2 = 0$ (i.e., to choose Q if Player 1 has chosen C), a contradiction.

If $q_2 < q_2^*$, then by (14) $\mu_1^{q_2} < 0$ and by choosing C Player 1 obtains a positive payoff for all μ_1. By sequential rationality it follows that for any information policy and resulting posteriors Player 1 chooses $q_1 = 1$. Therefore, Bayesian consistency implies that $\mu_2 = \mu^0$ and by sequential rationality and Assumption 2 Player 2 chooses $q_2 = 1$, a contradiction.

When $q_2 = q_2^*$, by (14) the belief that makes Player 1 indifferent between C and Q is $\mu_1^{q_2^*} = 0$. All Bayes plausible posteriors give the same expected payoff $\frac{\mu^0 v_1^2}{v_1 + c_1}$ to Player 1. It follows that any information policy is optimal for Player 1. Let $q_1(\mu_1)$ be Player 1 probability of choosing C when the posterior belief is μ_1. By sequential rationality, Player 1 chooses $q_1(\mu_1) = 1$ for any $\mu_1 > 0$, and $q_1(0) \in [0, 1]$ for $\mu_1 = 0$. To satisfy sequential rationality for Player 2, it must be $\mu_2 = \mu_2^*$. Then the information policy and competition strategy for Player 1 must satisfy $\frac{\mathbb{E}_\phi[\mu_1 q_1(\mu_1)]}{\mathbb{E}_\phi[q_1(\mu_1)]} = \mu_2^*$. For a given posterior distribution ϕ, and using $q_1(\mu_1) = 1$ for $\mu_1 > 0$, this equality is equivalent to $\mathbb{E}_\phi[\mu_1 \mid \mu_1 > 0](1 - \phi(0)) = [q_1(0)\phi(0) + (1 - \phi(0))]\mu_2^*$. Since, by Bayes plausibility, the term on the lhs is equal to μ^0, this equality is satisfied when $q_1(0) = \frac{\mu^0 - \mu_2^*(1-\phi(0))}{\phi(0)\mu_2^*} = 1 - \frac{\mu_2^* - \mu^0}{\phi(0)\mu_2^*} \in [0, 1]$. For this to be true it must be $\phi(0) \geq \frac{\mu_2^* - \mu^0}{\mu_2^*}$ and $\mathbb{E}_\phi[\mu_1 \mid \mu_1 > 0] \geq \mu_2^*$.

Therefore any PBE can be characterized by $\phi(0) = \frac{\mu_1^H - \mu^0}{\mu_1^H}$ for $\mu_1^H \in [\mu_2^*, 1]$ and $\mathbb{E}_\phi[\mu_1 | \mu_1 > 0] = \mu_1^H = \frac{\mu^0}{1 - \phi(0)}$. In these equilibria, for belief $\mu_1 = 0$ Player 1 chooses $q_1 = 1 - \frac{\mu_2^* - \mu^0}{\phi(0)\mu_2^*}$ and Player 2 selects $q_2 = \frac{v_1}{v_1 + c_1}$. All equilibria are payoff equivalent for the two players with expected payoffs equal to $\frac{\mu^0 v_1^2}{v_1 + c_1}$ for Player 1 and $\frac{(\mu_2^* - \mu^0) v_2}{\mu_2^*}$ for Player 2. □