Proposing New Equilibrium Concepts in Dynamic Games with Noisy Signals

Jeong-Yoo Kim*
Kyung Hee University

Abstract

In this paper, we consider dynamic games with noisy signals in which due to noises in observations any signal is possible in equilibrium. We argue that in the games with noisy signals, the consistency condition of beliefs required by the perfect Bayesian equilibrium (PBE) or sequential equilibrium (SE) is too *strong* to comprehend some reasonable outcomes, and accordingly propose alternative solution concepts which we will call ϵ-perfect Bayesian equilibrium (ϵ-PBE) and limit perfect Bayesian equilibrium (limit PBE). The two equilibrium concepts rely on a weaker consistency condition that requires beliefs to be updated by Bayes' law only if the likelihood of the signal given the equilibrium strategy was played exceeds $\epsilon > 0$. Our concepts are consistent with empirical observations that show higher deviations from Bayes' law for rare events. We show that under a mild condition, both of ϵ-PBE and limit PBE recover the first mover advantage that disappears with even a slight noise in observing the first mover's strategy. We also show that there exists a fully revealing outcome in a game of cheap talk to an informed receiver with monotone motives if ϵ-PBE or limit PBE is employed as equilibrium concepts.

Key Words: dynamic games with noisy signals, ϵ-likelihood consistency, ϵ-perfect Bayesian equilibrium, limit perfect Bayesian equilibrium, simple likelihood test

JEL Classification Code: C72

*Corresponding author: We are grateful to Kyle Bagwell and Bob Wilson for inspiring this paper. Mailing address: Department of Economics, Kyung Hee University, 1 Hoegidong, Dongdaemunku, Seoul 130-701, Korea, Tel: +822-961-0986, Fax: +822-966-7426, Email: jyookim@khu.ac.kr

1 Introduction

Observability is one of the central issues in a noncooperative game theory. In particular, in dynamic games, an equilibrium outcome must be highly sensitive to whether or not a player can observe the previous action of the other player. Indeed, the issue is closely related to information. If a player can observe all the previous moves of the other players perfectly, we call it a game of perfect information, and if he cannot, it is called a game of imperfect information. Starting from subgame perfection (Selten, 1965), many equilibrium concepts that have been developed in the context of dynamic games pay close attention to the issue. In particular, the concept of subgame perfection may become less appealing in games of imperfect information, even if it is very successful in eliminating all the unreasonable Nash equilibrium outcomes that rely on incredible threat in games of perfect information. However, this problem has been remedied by various subsequent equilibrium concepts including trembling-hand perfect equilibrium in extensive forms (Selten, 1975), sequential equilibrium (Kreps and Wilson, 1982) and perfect Bayesian equilibrium (Fudenberg and Tirole, 1991a) etc. which embrace the possibility of unobservability in equilibrium concepts. It was mainly due to Kreps and Wilson (1982)'s ingenuous idea of introducing beliefs (about what a player cannot observe) into a game.

This idea of using beliefs shifted the attention of game theorists to beliefs of players at an information set in order to define equilibrium concepts for dynamic games. Especially, this approach of using beliefs has been remarkably successful in signaling games in which a player has an opportunity to infer the unobservable type of the other player from observing his previous action. Many equilibrium concepts have been proposed in signaling games to eliminate implausible predictions of the perfect Bayesian equilibrium (PBE) and the sequential equilibrium (SE). Most of them basically rely on defining what is a more reasonable belief at an information set that could be reached out of the equilibrium path.[1] However, all of those concepts except trembling-hand perfect equilibrium which is not a belief-based refinement[2] impose the minimum requirement of beliefs, what we call weak consistency, which is part of the definition of the perfect Bayesian equilibrium.[3] Roughly speaking, weak

[1]To name a few, see Cho and Kreps (1987), Banks and Sobel (1987), Grossman and Perry (1986) and Farrell (1993).

[2]Selten (1975) did not explicitly mention "beliefs" in defining trembling-hand perfect equilibrium but the notion of beliefs is implicit in his definition in the sense that perturbed strategies play the role of beliefs.

[3]Mas-Colell *et al.* (1995) distinguishes PBE and weak PBE, and Fudenberg and Tirole (1991a) defines PBE by imposing the condition of no signaling what you don't know on the off-the-equilibrium belief in

consistency requires that beliefs must be consistent with equilibrium strategies in a Bayesian sense, i.e., beliefs must be updated according to Bayes' law whenever possible.[4] This is the weakest version of consistency that has been known so far. It is clearly weaker than the consistency condition required by sequential equilibrium. In this paper, we argue that this version of consistency is too *strong* to comprehend all reasonable outcomes.

Contrary to what we have described so far, observability is not a matter of all or nothing. Perfect observability and perfect unobservability are not all the possibilities in the real world, although we are most of the times confined into the two extreme cases.[5] Observations by players are often imperfect. They may be able to observe only a noisy signal of what the other player did. We will call such a situation with imperfect observation a dynamic game with noisy signals.[6] It is a game in which the outcome that the second mover (player II) observes provides only an imperfect signal on the private information of the first mover (player I)'s choice which is relevant to the payoff of the second mover.[7] In such games, the second mover can infer the private information of the first mover stochastically by observing the signal that occurs as a result of his action, if he is informed of the probabilistic signal generating rule. However, if the second player knows what the equilibrium strategies are, in other words, what strategy the first player is supposed to play in equilibrium, the inference should be modified by this additional information. Suppose, for example, that player II happens to observe an outcome that can occur only with very low probability if player I takes the equilibrium

---

addition to the requirement of weak PBE. But we do not distinguish between the two concepts for the reason that will be clear soon. Briefly speaking, the additional requirement is not binding in games we consider in this paper.

[4]What "whenever possible" means exactly has never been precisely defined. Watson (2017) is one of recent researches in this direction. He requires conditional-probability updating on separate dimensions of the strategy space.

[5]Fudenberg and Tirole (1991a) defined the perfect Bayesian equilibrium only for games with perfectly observed actions, although their definition can be straightforwardly extended to games with imperfectly observed actions.

[6]Our model of noisy signals should be distinguished from noisy signaling models by Matthews and Mirman (1983) and Hertzendorf (1993) in the sense that no noisy signaling is involved in our model. Noisy signals in our model are just a consequence of the Nature's choice, not a consequence of a player's strategic choice to influence the other player's belief. Neither player has private information of his own type in our model.

[7]In this sense, dynamic games with noisy signals share common features with games of imperfect public information. See Fudenberg and Tirole (1991b) for examples of games of imperfect public information. One difference is that signals in our games need not be public information. They may be private information of the second mover.

action. The notion of weak consistency requires player II to still believe with probability one that player I took the equilibrium action, insofar as the outcome he observed can be generated from the equilibrium action.[8] However, what if this outcome is such that can occur with much higher (*ex ante*) probability[9] as a result of deviating from the equilibrium? Can we say that it is more reasonable to still believe that player I took the equilibrium action but the unlikely outcome has occurred unexpectedly, rather than the first mover took a non-equilibrium action which was very likely to lead to the observation?

In a game with perfect observability, an equilibrium action induces an outcome which is possible in equilibrium, whereas an off-the-equilibrium action leads to an outcome which is impossible in equilibrium. So, it is natural that player II believes that player I took an equilibrium action if he observes an outcome which is possible only on an equilibrium path. However, in a game with imperfect observability that can occur mainly due to noises (with unbounded supports), any outcome is possible, whether player I took an equilibrium action or not. So, it is not necessarily true that player I actually took an equilibrium action when player II observes an outcome which is possible on an equilibrium path. It could be also possible if player I took an off-the-equilibrium action. Moreover, contrary to games with perfect unobservability in which some outcome is possible in equilibrium and some other outcomes are not, every information set (i.e., every outcome) can be reached on an equilibrium path in games with imperfect observability (due to noisy signals), although probabilities of reaching different information sets may differ, depending on whether to choose an equilibrium action or not. The information about this probability difference is completely ignored by the concept of PBE, but it may have an important implication.[10] Then, when should we believe that player II reached a certain information set as a result of an off-the-equilibrium action by player I? If the probability that the information set is reached given that player I took an equilibrium action is very low, it may be more reasonable to reject the hypothesis that player I chose the proposed equilibrium action and to believe that player I deviated from the

[8]This is the case only when we take only pure strategies into account.

[9]By *ex ante* probability, we mean the probability at the point of time before the players get additional information of what the equilibrium strategies are, i.e., what strategies they will agree to play in equilibrium.

[10]For example, let h_1 and h_2 be two information sets. Suppose the probability that h_1 is reached if player I took an equilibrium action is low and the probability is high if she did not take an equilibrium action. On the other hand, suppose the probability that h_2 is reached is high if player I took an equilibrium action and it is low if she did not take an equilibrium action. Then, it seems reasonable that the beliefs that player I took an off-the-equilibrium action at h_1 and h_2 should be updated differently, more specifically, the probability at h_1 should be higher.

equilibrium strategy. At least we should doubt the assumption that she took the equilibrium strategy. In this event, we regard the outcome as an almost off-the-equilibrium signal.

Players sometimes make mistakes (e.g., Selten [1975], Kreps and Wilson [1982]).[11] If a mistake is perfectly observed without noises, the second mover can surely identify where he is at his information set. However, if player I's mistake (generally, her action) is imperfectly observed with noises, player II cannot be sure whether an observation of a deviant action is due to player I's mistake or player II's mistake in observing player I's action. With noises, therefore, equilibrium behavior and off-the-equilibrium behavior are pooled. In this case, one should interpret even an equilibrium outcome by taking account of this small probability of player I's deviation more seriously,[12] rather than taking it for granted that player I took an equilibrium action. Therefore, we believe that the consistency condition required by PBE or SE should be relaxed. A noisy signal that can be observed in equilibrium should be interpreted in some cases as a consequence of an equilibrium action or in other cases as a consequence of a mistake or an intentional deviation especially when the signal is very unlikely to occur, provided that an equilibrium action was actually taken.

Let us take a simple example. Suppose a worker (player I) exerts effort a and then an outcome y is realized as a result of it. Assume that y can take the value of a with probability .9999 and of $a-1$ with probability .0001, if player I chooses action a. Now, suppose that a^* is the equilibrium effort. If a principal (player II) actually observes $y = a^* - 1$, is it still reasonable for player II to believe that player I chose a^*? The perfect Bayesian equilibrium stipulates that the unique reasonable belief is to believe with probability one that player I took a^*, not $a^* - 1$. However, it seems more reasonable (at least) to us to believe that player I chose the action $a^* - 1$, because the likelihood (probability) that $y = a^* - 1$ actually happens given that $a = a^*$ is very small (.0001).

On this ground, we think that it is too demanding to insist on the weak consistency of beliefs that the perfect Bayesian equilibrium requires. The perfect Bayesian equilibrium imposes a too strong restriction on beliefs on the equilibrium path that occurs only with infinitesimally small probability (too stringent!), whereas it imposes no restriction at all on the path that occurs in equilibrium with probability zero, thereby allowing arbitrary beliefs

[11] Players may indeed deviate intentionally from a proposed equilibrium (e.g., Kohlberg [1990]). Considering this possibility allows us to develop a forward induction argument.

[12] Kreps and Wilson (1982) and Kohlberg (1990) considered this small possibility in order to restrict the belief off the equilibrium path, but our approach is different and even in the opposite direction in the sense that we relax the requirement imposed on equilibrium beliefs.

(too lenient!). Furthermore, the perfect Bayesian equilibrium requires that one assign the same beliefs on the path that can occur only with a very low probability in equilibrium as on the path that can occur with a very high probability in equilibrium. It seems reasonable, neither. For example, suppose that player I intended to mean "reminder" by a word "memento" used in equilibrium to mean it. But if the second mover observes a noisy signal "mementa" or "mementu", it would be reasonable for him to believe that the first mover who intended to mean "reminder" just misspelled it, because such misspelling often occurs. However, if he observes a signal "momentu", it appears to be more natural to believe that she intended the word "momentum", because those who intended to spell "memento" are very unlikely to misspell it to "momentu". In this case, it would be more reasonable to believe that she intended to mean a different word "momentum".

Based on this consideration, we propose two equilibrium concepts, the ϵ-perfect Bayesian equilibrium (ϵ-PBE) and the limit perfect Bayesian equilibrium (limit PBE),[13] both of which are slightly weaker than perfect Bayesian equilibrium. The notion of ϵ-PBE imposes the consistency condition only if the likelihood that a noisy signal occurs given the equilibrium action was played is reasonably high (higher than ϵ for some small $\epsilon > 0$), and allows the second mover to believe that the first mover deviated from the proposed equilibrium if the likelihood is not higher than $\epsilon > 0$.[14] We call this consistency ϵ-likelihood consistency. A rationale for ϵ-likelihood consistency is that if the likelihood is very low, it is reasonable for

[13]Similarly, Myerson and Reny (2019) define the perfect conditional ϵ-equilibrium to extend sequential equilibrium to games with infinite sets of signals and actions. They first define the conditional ϵ-equilibrium by a strategy profile such that no player could expect significant gains by unilaterally deviating from it after any event that has positive probability in the equilibrium, and then define the perfect conditional ϵ-equilibrium by a strategy profile such that there exists a pair of a perturbed strategy profile and nature's perturbation in the neighborhood which constitutes a conditional ϵ-equilibrium. Our definition is to apply consistency to an event with a positive *likelihood*, not to an event with a positive *probability*. This is reasonable because a positive likelihood enables us to update the posterior belief by the conditional density function (not the conditional probability). Also, relatedly, ϵ-Nash equilibrium or ϵ-rationality by Radner (1980) permits an ϵ-gap in payoffs rather than in likelihood (not probability). Note that ϵ-PBE is not a special case of ϵ-Nash equilibrium, because a restriction on the belief as in the definition of PBE or ϵ-PBE is to reduce the set of sequentially rational strategies, not to affect the set of optimal strategies by allowing small errors in payoffs.

[14]Ortoleva (2012) proposed a different equilibrium concepts of hypothesis testing equilibrium from a similar motivation. His approach is, however, non-Bayesian. He assumes that players use subjective priors which were initially chosen by them among all possible priors and pick a new prior if they receive an unexpected evidence. For more details of the hypothesis testing equilibrium, see Ortoleva (2012) or Sun (2019).

the second player to believe that the first mover made a mistake rather than believe that the second mover himself made a wrong observation, thereby ignoring the equilibrium message.[15] If an equilibrium is an ϵ-PBE for any small $\epsilon > 0$,[16] we call it a limit PBE.

These concepts turn out to be very useful especially in games with noisy signals that make every outcome possible in equilibrium. By allowing more natural beliefs on the equilibrium path rather than restricting beliefs off the equilibrium path, we can obtain more realistic outcomes in equilibrium in many games with noisy signals. For example, Bagwell (1995) argued that the first-mover advantage is eliminated if there is even a slight amount of noise in observing the first-mover's choice.[17] We show, however, that ϵ-PBE enables the first-mover advantage to be recovered. We also show that a fully revealing outcome is possible in a game of cheap talk to an informed receiver with monotone motives if ϵ-PBE or limit PBE is employed as equilibrium concepts.

There is a large empirical literature arguing that players tend to violate Bayes' law.[18] Among others, Grether (1992) showed that empirical data for rare events are more likely to deviate from Bayes' law. Holt and Smith (2009) also presented the experimental evidence that players do not tend to follow Bayes' law after very unlikely events occur.

The organization of this paper goes as follows; In section 2, we introduce definitions of ϵ-PBE and limit PBE. In section 3, we revisit the example by Bagwell (1995) and show that

[15] Some may think that if this ϵ-chance event occurs, it may be more reasonable to believe that the first mover chose an action maximizing the likelihood function, rather than just ignoring the equilibrium message. If so, the resulting equilibrium concept may not contain the set of PBE, though. Ortoleva (2012) takes the maximum likelihood approach.

[16] A likelihood and a (posterior) belief are both conditional probabilities, but a belief is a probability of reaching a node at an information set conditional on the event that the information set is reached, whereas a likelihood is a probability of reaching an information set conditional on the event that the player chose an equilibrium action. Also, a likelihood is similar to the concept of plausibility used by Bonanno (2013), but plausibility is defined for all histories and not for information sets.

[17] This result relies crucially on the restriction to pure strategy equilibria. Bagwell was aware of the limitation of the claim himself, and later van Damme and Hurkens (1997) generalized the mixed strategy version of the claim. Also, Maggi (1999) showed that the first-mover advantage reappears when the first mover's choice is based on pure private information, i.e., private information that is payoff-irrelevant for the second mover. By focusing only on "pure private information", he abstracts from signaling consideration. If one considers private information that is not pure, the first mover advantage may or may not appear. See Gal-Or (1987), Matthews and Mirman (1983) and Hertzendorf (1993).

[18] See, for example, Kahneman and Tversky (1973), Grether (1992), Griffin and Tversky (1992), and Holt and Smith (2009).

we can recover the first-mover advantage by applying ϵ-PBE, and generalize this claim. In Section 4 and 5, we apply the two concepts to a price competition model with a Stackelberg leader and to a cheap talk model with an informed receiver respectively. Concluding remarks follow in Section 6.

2 Definitions

There are two players, player I ("she") and player II ("he"). Player I first chooses an action a_1 from a set $A_1 \in \mathbb{R}$. Each action induces a probability distribution over observable outcomes y in a set $Y \in \mathbb{R}$. Let $f(y \mid a_1)$ and $F(y \mid a_1)$ denote the density function and the corresponding probability distribution function conditional on the action a_1 is taken. After observing y (not observing a_1), player II responds by choosing an action a_2 from a set $A_2 \in \mathbb{R}$. We assume that $f(y \mid a_1)$ is common knowledge. The support of f given a_1 is defined by $\text{supp}(f \mid a_1) = \{y \mid f(y \mid a_1) > 0\}$. Most of the time, we will assume that $A_1 = Y = \mathbb{R}$, i.e., $\text{supp}(f \mid a_1) = \mathbb{R}$ for any $a_1 \in \mathbb{R}$. We also assume that $f(y \mid a_1)$ first-order stochastically dominates $f(y \mid a_1')$ if $a_1 > a_1'$.

The payoff to player i is given by a function $U^i : A_1 \times A_2 \to \mathbb{R}$. We assume that U^i is twice-continuously differentiable with respect to a_1 and a_2, and that it is concave in a_1 and a_2. This guarantees that for any a_j, U^i is uniquely maximized by $a_i^{BR}(a_j)$, which gives continuous best response functions. We also assume that $a_i^{BR}(a_j)$ is strictly monotonic, implying that we exclude the possibility that $U^i_{ij} = 0$.[19] Notice that U^i does not depend directly on y.[20]

A strategy for player I, σ^I, is identical to his action. A system of beliefs is defined by a map from the set of possible observations (Y) to $\Delta(A_1)$, where ΔA_1 denotes the set of all

[19]If $U^i_{ij} = 0$ for all a_1 and a_2, we will say that a_1 and a_2 are strategically independent. If this is the case, $a_i^{BR}(a_j)$ is the same for any a_j, that is, the best response function $a_i^{BR}(a_j)$ is constant. We ignore this possibility, because there is no strategic interaction in this case. If $U^i_{ij} \gtrless 0$ for any a_1 and a_2, a_1 and a_2 are called strategic complements (strategic substitutes, *resp.*). Our analysis encompasses both cases.

[20]We can imagine many situations in which the payoffs of the players depend directly on y. For instance, in a Cournot game with fluctuating demands, the profit of each firm depends directly on the market price. The output produced by the rival firm affects the profit only indirectly through affecting the market price. Also, in a principal-agent game, the payoff of the principal usually depends directly on the output. The agent's effort level affects the principal's utility only through determining the output level stochastically. It is not difficult to see that our equilibrium concepts can be applied straightforwardly to those situations. In other words, this assumption is not crucial, but just for expositional convenience.

probability distribution functions (or density functions) over A_1. The conditional density function as a belief is denoted by $g(a_1 \mid y)$. If the belief $g(a_1 \mid y)$ has the whole probability mass at a_1^0, we will simply use $\text{supp}(g \mid y)$ interchangeably with $g(a_1 \mid y)$. That is, if the belief is $g(a_1 \mid y) = \delta_{a_1^0}(a_1)$ where $\delta_{a_1^0}(a_1)$ is a Dirac's delta function, it is more convenient to use the notation $b(y) = a_1^0$ for the belief where $b : Y \to A_1$ is a belief function. Finally, a strategy for player II is a function from the set of beliefs (A_1) into A_2, $\sigma^{II} : A_1 \to A_2$. Note that σ^{II} depends on y only indirectly through forming belief b, because we assume that y does not affect U^{II} directly.

An assessment is a pair (σ, b) of a strategy profile σ and a system of beliefs b, where $\sigma \equiv (\sigma^I, \sigma^{II}) = (a_1, a_2(b))$. Note that σ^I and σ^{II} are both defined as pure strategies. For simplicity, we will restrict our attention to pure strategies throughout the paper.

Before we propose two definitions of our main equilibrium concepts, we introduce the definition of the perfect Bayesian equilibrium which is adapted to our model.

Definition 1. *An assessment $(a_1^*, a_2^*(b), b(y))$ is a perfect Bayesian equilibrium (PBE) if (i) it is sequentially rational, i.e., $\int_Y U^I(a_1^*, a_2^*(b(y)))f(y \mid a_1^*)dy \geq \int_Y U^I(a_1, a_2^*(b(y)))f(y \mid a_1)dy, \forall a_1 \in A_1$ and for every $y \in Y$, $U^{II}(a_1^*, a_2^*(b(y))) \geq U^{II}(a_1^*, a_2), \forall a_2 \in A_2$, and (ii) $b(y)$ is weakly consistent, i.e., $b(y) = a_1^*$ if $y \in supp(f \mid a_1^*)$ and $b(y)$ is an arbitrary density function if $y \notin supp(f \mid a_1^*)$.*[21]

To formalize our solution concepts, we need to define a little relaxed notion of consistency as follows.

Definition 2. *An assessment $(a_1^*, a_2^*(b), b(y))$ satisfies ϵ-likelihood consistency iff for any y, $b(y) = a_1^*$ if $L(a_1^*; y) > \epsilon$ and $b(y) \in \mathbb{R}$ can be arbitrary if $L(a_1^*; y) \leq \epsilon$ where $L(a_1; y) = f(y \mid a_1)$ is a likelihood function.*

The likelihood function tells us how likely the occurrence of y is if player I chooses the equilibrium action a_1^*. We may call this ($L(a_1^*; y) \gtreqless \epsilon$) a simple likelihood test. If $L(a_1^*, y) \leq \epsilon$, we reject the hypothesis that player I is playing an equilibrium strategy.[22] In this definition,

[21] If $y \in \text{supp}(f \mid a_1^*)$, the conditional density function $g(a_1^* \mid y)$ is well defined, since $f(y \mid a_1^*) > 0$.

[22] Alternatively, some may want to use a likelihood ratio test but it is not appropriate in this situation to determine whether player I played the equilibrium action or deviated to some action among many possible off-the-equilibrium actions. Also, the test is problematic because the likelihood ratio does not reflect the information that the equilibrium action is more often used than non-equilibrium actions. Therefore, comparing the likelihood without weights will not make much sense.

we slightly extend the meaning of "off the equilibrium" to the case that a random outcome y such that $L(a_1^*; y) \leq \epsilon$ occurs. This event that $L(a_1^*; y) \leq \epsilon$ can be interpreted as almost off the equilibrium path.[23] It is clear that $L(a_1^*, y) \leq \epsilon$, if $y \notin \text{supp}(f \mid a_1^*)$. So, it is obvious that weak consistency implies ϵ-likelihood consistency. Note that we interpret an equilibrium event and an off-the-equilibrium event in terms of the likelihood, not in terms of probability. Even if $L(a_1^*, y) = f(y \mid a_1^*) > 0$, it is possible that the probability that a particular value of y occurs given that a_1^* is chosen is zero, if y is a continuous random variable without any atom in the support of its probability distribution. We believe that this likelihood approach is more relevant to defining an equilibrium event. If the likelihood $L(a_1^*, y)$ is positive (although $\mathbb{P}(y \mid a_1^*) = 0$), it is possible to update the conditional density function (belief) by using the likelihood function, which is all that matters for an equilibrium event.[24]

Now, we can define our solution concepts formally.

Definition 3. *An assessment $(a_1^*, a_2^*(b), b(y))$ is an ϵ-perfect Bayesian equilibrium (ϵ-PBE), for some $\epsilon > 0$, if it satisfies sequential rationality and ϵ-likelihood consistency.*

This concept, ϵ-PBE, is similar to ϵ-perfect equilibrium which can be roughly defined by a strategy profile satisfying the property that if a certain pure strategy yields a strictly lower payoff than another, the strategy should be used with a probability less than $\epsilon(> 0)$, not necessarily with a zero probability.[25] Note that ϵ-PBE is a slight departure from the weak consistency of a *belief*, whereas ϵ-perfect equilibrium is a slight departure from the best response of a *strategy*. Just as the trembling hand perfect equilibrium is defined as the limit of ϵ-perfect equilibrium, we can define a stronger equilibrium concept which can be obtained by making ϵ approach zero.

Definition 4. *A strategy profile, $(a_1^*, a_2^*(b))$, is a limit perfect Bayesian equilibrium (limit PBE), if for any $\epsilon > 0$ such that $\epsilon \leq \bar{\epsilon}$ for some $\bar{\epsilon} > 0$, there exists $b(y; \epsilon)$ such that an assessment $(a_1^*, a_2^*(b), b(y; \epsilon))$ is an ϵ-PBE.*

The difference between the definitions of the ϵ-PBE and the limit PBE is just that the former holds for some small $\epsilon > 0$, while the latter holds for any small $\epsilon > 0$. Although a

[23] We can define an almost off-the-equilibrium event in terms of either the likelihood less than ϵ or the p-value less than ϵ interchangeably.

[24] Myerson and Reny (2019) defined perfect conditional ϵ-equilibrium conditional on positive probability events, not conditional on positive likelihood events.

[25] See Myerson (1978) for the formal definition of the ϵ-perfect equilibrium.

perfect equilibrium is a limit of ϵ-perfect equilibrium,[26] a limit PBE is not necessarily a limit of ϵ-PBE, because the definition of the limit PBE does not require that $\lim_{\epsilon \to 0} b(y;\epsilon)$ exists.[27] If we require that there must exist $b(y)$ such that $b(y) = \lim_{\epsilon \to 0} b(y;\epsilon)$, this additional continuity requirement is so strong that the resulting equilibrium concept, which we will call strong limit PBE, becomes equivalent to PBE, just as a perfect equilibrium is a limit of ϵ-perfect equilibrium.

Definition 5. *An assessment $(a_1^*, a_2^*(b), b(y))$ is a strong limit perfect Bayesian equilibrium (strong limit PBE), if for any $\epsilon > 0$ such that $\epsilon \leq \bar{\epsilon}$ for some $\bar{\epsilon} > 0$, there exist $b(y;\epsilon)$ such that (i) $\lim_{\epsilon \to 0} b(y;\epsilon) = b(y)$ and (ii) $(a_1^*, a_2^*(b), b(y;\epsilon))$ satisfies ϵ-PBE.*

We have the following properties, i.e., inclusion relations among those equilibrium concepts. The proofs are provided in the Appendix.

Proposition 1. *PBE $\subset$ limit PBE $\subset$ ϵ-PBE.*

Proposition 2. *The strong limit PBE is equivalent to PBE.*

In the next two sections, we will present some counterexamples for the cases that the converses do not hold, i.e., ϵ-PBE does not imply limit PBE and that limit PBE does not imply PBE.

3 Stackelberg Model

We consider an example by Bagwell (1995) which captures the main feature of the Stackelberg model. (Figure 1) In this example, the unique Nash equilibrium is (C, C) in a static game (in which neither player can observe the other's choice), while the unique subgame perfect

[26] It is well known that the trembling hand perfect equilibrium requires a_2^* to be a best response to any nearby others' strategies perturbed from the equilibrium strategies (which can be also interpreted as a sequence of beliefs), but the sequential equilibrium only requires a_2^* to be a best response to the limit of the belief sequence. Therefore, limit PBE is in its spirit closer to the trembling hand perfect equilibrium than the sequential equilibrium in the sense that it requires $a_2(b)$ to be a best response to $b(y;\epsilon)$, not to $b(y)$, the limit of $b(y;\epsilon)$.

[27] In an evolutionary game theory, the concept of the limit ESS is defined from a similar motivation. It was proposed by Selten (1983) to alleviate the severe nonexistence problem of ESS. Contrary to the limit PBE, however, the limit ESS is required to be a limit of a sequence of strict ESS in ϵ-perturbed games, although it need not be ESS in the limit.

outcome is (S, S) in a sequential game (in which player 2 can observe player 1's choice), yielding the first-mover advantage.

Now, suppose player 2 can observe player 1 only imperfectly, i.e., if player 1 chooses S, player 2 may receive a signal of either S with probability $1 - \delta$ or C with probability $\delta (> 0)$, and if player 1 chooses C, player 2 receives a noisy signal of either C with probability $1 - \delta$ or S with probability δ. The extensive form of the resulting game is drawn as in Figure 2. Note that due to imperfect observation, any information set can be reached after player 1's choice of either S or C.

Interestingly, the unique PBE outcome of this imperfect information game is (C, C). To see that (C, C) is an equilibrium outcome, note that the weak consistency condition pins down the posterior belief at two information sets, h_1 and h_2, as $(\mu(x|h_1), \mu(z|h_1)) = (0, 1)$ and $(\mu(y|h_2), \mu(w|h_2)) = (0, 1)$, where $\mu(\omega|h_i)$ is the posterior probability that player 2 reaches node ω given that he reaches information set h_i, because he believes that she took the action C regardless of his observation. Thus, it is optimal for player 2 to choose C at both information sets h_1 and h_2. Therefore, player 1 chooses C.

It is also easy to see that the Stackelberg outcome (S, S) is not a PBE. If it is an equilibrium outcome, the only consistent belief is that $((\mu(x|h_1), \mu(z|h_1)), (\mu(y|h_2), \mu(w|h_2))) = ((1, 0), (1, 0))$. Then, player 2 will chooses his best response to S, which is S, whether he observes either S or C. In particular, even when he observes C, he chooses S by reasoning that he received a wrong signal, because player 1 must have chosen S. This implies that player 1 prefers deviating from S to C. Since this holds for any $\delta > 0$, the first-mover advantage disappears when there is even a slight noise ($\delta > 0$) in the observation of player 2.

Intuitively, a noise, no matter how small it is, makes any signal possible when player 1 makes the equilibrium choice S; hence, no off the equilibrium signal. So, no matter what signal player 2 may receive, it should be interpreted as the equilibrium meaning. If player 1's choice cannot affect player 2's belief and his choice due to imperfect observation, the resulting equilibrium outcome must coincide with the Nash outcome. Since the outcome yielding the first mover advantage is not a Nash equilibrium, it cannot be a PBE. This is mainly because the weak consistency condition that PBE requires ignores the information about a difference in probabilities even if the ex ante likelihood that player 1 chose S at information set h_1 (when the signal S was observed) is much higher than the ex ante likelihood at information set h_2 (when the signal C was observed), thereby forcing one to assign identical beliefs at h_1

and h_2 by superseding the information by the information that she can infer from equilibrium behavior (which is supposed to be played). If the information about a difference in ex ante probabilities should not be ignored but taken seriously, the posterior beliefs at h_1 and h_2 need not be identical. In particular, if the ex ante likelihood that S was played given that C was observed at h_2 is very small, player 2 should doubt the presumption that the signal C resulted from player 1's equilibrium behavior of choosing S and be open to all other possibilities by considering all possible scenarios. This is the motivation of ϵ-PBE.

If we use ϵ-PBE as our solution concept, we can see that the first mover advantage reappears in equilibrium. Suppose it is an equilibrium that player 1 chooses S and player 2 chooses S at h_1 and C at h_2. After observing the signal S, player 2 will believe that player 1 chose S because S is the equilibrium meaning.[28] Therefore, he will respond by choosing S based on the belief $(\mu(x|h_1), \mu(z|h_1)) = (1, 0)$. If observing the signal C, however, he will think that it is possible that player 1 chose S because it is an equilibrium action, but that the probability of observing C if player 1 actually chose S is very low, i.e., $f(C \mid S) = \delta$. If $\delta < \epsilon$, ϵ-PBE does not require the belief to be updated by the Bayes' law. In this case, arbitrariness of the belief allows player 2 to believe that player 1 chose C, thereby responding by choosing C. Knowing this, player 1 will choose S because $5(1-\delta)+4\delta > 5\delta+4(1-\delta)$ for small $\delta(< \frac{1}{2})$. However, as ϵ gets smaller so that $\epsilon' < \delta$, this equilibrium (S, S, C) cannot be an ϵ'-PBE because the consistency condition does not allow such behavior as far as $f(C \mid S) = \delta > \epsilon'$. Therefore, (S, S, C) that shows the first mover advantage is an ϵ-PBE for $\epsilon > \delta$, but not a limit PBE.

Can the Nash outcome (C, C) be supported as an ϵ-PBE? It is clear that player 2 chooses C after receiving a signal C (at information set h_2). After receiving S, however, player 2 doubts the possibility that player 1 chose the equilibrium action C, because the probability that he receives S is too small if player 1 chose C, i.e., $f(S \mid C) = \delta < \epsilon$. In this case, arbitrary belief is allowed. If he believes that player I chose C, it is optimal for him to choose C. Therefore, player 1 will choose C. This implies that the Nash outcome is also an ϵ-PBE. Furthermore, for $\epsilon' < \delta$, ϵ'-likelihood consistency allows him only to believe that player 1 chose C at information set h_2, implying that player 1 will choose C. So, (C, C, C) is a limit PBE. In fact, it directly follows from Proposition 1, because (C, C, C) is a PBE.

We can generalize this result to games with any finite actions. Let $n = |A_1|$. For any

[28] Player 2 knows that player 1 is supposed to choose S in equilibrium. So, in this case, the equilibrium meaning of player 1's behavior is S, regardless of what signal player 2 observes.

$a_1 \in A_1$, player 2 is assumed to observe $y = a_1$ with probability $1 - \delta$ and receive a wrong signal $y = a'(\neq a_1)$ with equal probabilities $\frac{\delta}{n-1}$. Let (a_1^N, a_2^N) be the Nash equilibrium in the static game without noises. Also, let $\pi_1^N \equiv \pi_1(a_1^N, a_2^N)$ and $\pi_1^s \equiv \pi_1(a_1^s, a_2^s)$. Then, we have the following Folk Theorem.

Proposition 3 (Folk Theorem)**.** *Define $A^* = \{a_1 \in A_1 \mid \pi_1(a_1, a_2^{BR}(a_1)) > \underline{\pi}_1\}$ where $\underline{\pi}_1 \equiv \max_{a_1} \pi_1(a_1, a_2^{BR}(\underline{a}_1))$ and $\underline{a}_1 = \arg\min \pi_1(a_1, a_2^{BR}(a_1))$. Then, in a game with noises, for any $a_1 \in A^*$, there exists $\bar{\delta} > 0$ such that for any δ such that $0 < \delta \leq \bar{\delta}$, $(a_1, a_2^{BR}(b(y)))$ where $b(a_1) = a_1$ and $b(y) = \underline{a}_1$ if $y \neq a_1$ is an ϵ-PBE for any $\epsilon > 0$ such that $\delta < \epsilon$. In particular, if $\pi_1^N > \underline{\pi}_1$, any a_1 between a_1^N and a_1^s is an ϵ-PBE action for such ϵ.*

The belief $\underline{a}_1 = \arg\min \pi_1(a_1, a_2^{BR}(a_1))$ is the one that gives player 1 the minimum payoff when player 2 responds optimally along the best response function. We call $\underline{a}_1$ the most pessimistic belief and $a_2^{BR}(\underline{a}_1)$ the most severe threat.

The payoff $\underline{\pi}_1 \equiv \max_{a_1} \pi_1(a_1, a_2^{BR}(\underline{a}_1))$ plays a similar role of minmax payoff in the Folk Theorem of repeated games in that player 1 can get at most $\underline{\pi}_1$ when player 2 penalizes her most severely by choosing $a_2^{BR}(\underline{a}_1)$. It is the player 1's payoff from her most profitable deviation when she expects the most severe threat. This reflects the fact that player 1's actual hidden action is not necessarily identical to the belief $\underline{a}_1$. We will just call $\underline{\pi}_1$ the minmax payoff of player 1 in a broad sense.

This Folk Theorem says that if the Nash payoff is higher than the minmax payoff, any outcome (including the Stackelberg outcome) along the best response function of player 2 that yields her a higher payoff than the Nash outcome can be supported as an ϵ-PBE for some small $\delta > 0$, as long as we pick $\epsilon > \delta$. Note that when a_1^* is supported in ϵ-PBE, the equilibrium payoff is not $\pi_1(a_1^*, a_2^{BR}(a_1^*))$, because $\pi_1(a_1^*, a_2^{BR}(\underline{a}_1))$ could be realized with a small probability, depending on the observation y. So, it is better to state this Folk Theorem given in Proposition 3 in terms of strategies rather than in terms of payoffs, unlike the Folk Theorem of repeated games.

We can also generalize the result to games with continuum actions.

Proposition 4. *Assume that $|U_1^I(a_1, a_2)| \leq M$ for some $M \in (0, \infty)$. In a game with noises, $(a_1^s, a_2^{BR}(b))$ is an ϵ-PBE and a limit PBE if the following [UP] condition holds; $U^I(a_1^s, a_2^{BR}(b))$ is strictly monotonic in b and unbounded.*

This proposition has an important implication that the first mover advantage which was forged with even a slight noise in observing the choice of the first mover is re-supported as an equilibrium outcome if we use ϵ-PBE or limit PBE as our equilibrium concept.

The sufficient condition provided in Proposition 4, which we call Unbounded Penalty [UP] condition, is crucial to our result. Intuitively, if a value of y which is very unlikely given that a_1^s is played is realized, player 2 must form a very extreme belief a_1^ϵ which makes it possible for him to rationally respond (to a_1^s), so as to lead to a very small $U^I(a_1, a_2^{BR}(a_1^\epsilon))$ when [UP] condition holds.[29] This punishment strategy $a_2^{BR}(a_1^\epsilon)$ can be called "boiling-in-oil" strategy, often referred to as in the principal-agent literature.[30] Without the condition, player 2's best response even to such an extreme belief might be a mild effect that the resulting utility is bounded below, so a very harsh punishment by player 2 would not be feasible.

Indeed, the [UP] condition is too strong. If this condition does not hold, the Stackelberg action a_1^s may be supported as an ϵ-PBE for some $\epsilon > 0$, but cannot be supported as a limit PBE. This is consistent with the result in Bagwell's example.

An interesting feature of this model is that player 1 knows that player 2's response does not depend on a_1 directly, i.e., her choice of a_1 does not affect player 2's response directly. Most of the times, player 2 behaves as a passive player without responding to a_1. (Actually, player 2 cannot respond to a_1 because what he can observe is y, not a_1.) Then, how can player 1 achieve the Stackelberg outcome? It is possible because player 2 can respond to a signal y which can be affected by a_1. If player 1 deviates from the Stackelberg outcome, player 2 may respond by a punishment action that a value of y falling short of a threshold level triggers with a very small probability. This possibility prevents player 1 from deviating from the Stackelberg action.

It is worthwhile to ponder on the implication of the credibility of a threat to punish in this model. As is well known, if the player 1's move is perfectly observed by player 2, player 2 cannot successfully threaten to choose C, regardless of player 1's choice, because the threat is incredible. In other words, perfect observability makes it in the interest of player 2 to respond to his observation S, thereby choosing S. However, with even a slight noise in observation, player 2 can commit to choosing C. This threat is credible because player 2 can observe nothing to respond to. This is why player 2 cannot achieve the Stackelberg outcome in the Bagwell's example and in the generalized model. Interestingly, employing the concept of ϵ-PBE brings a similar effect to observability. If we employ ϵ-PBE, when player 1 chooses S in equilibrium, player 2 responds by choosing S if he observes a signal S with high probability. This is not because he observes her choosing S, but because he infers

[29]If $A_2 \neq \mathbb{R}$, [UP] condition can be modified by replacing $a_2 \to -\infty$ and $a_2 \to \infty$ with $a_2 \to \underline{a}_2$ and $a_2 \to \overline{a}_2$ respectively, where $\underline{a}_2 = \inf \hat{A}_2$, $\overline{a}_2 = \sup \hat{A}_2$ and $\hat{A}_2$ is the range of a_2^{BR}, i.e., $\hat{A}_2 = a_2^{BR}(A_1) \subset A_2$.

[30]See, for example, Rasmusen (1994).

that she chooses S in equilibrium. Player 2 does not need to believe player 1's threat to choose S in the case that he observes a signal C, because it is a very unlikely event given that player 1 actually chose S. In this case, he may respond by choosing C. This means that player 2's threat to choose C regardless of his observation is not credible even with imperfect observation. Therefore, commitment is impossible.

If a likelihood test in the case of noisy observation can have a similar effect of perfect observability, what is the essential difference that distinguishes the case of noisy observation from the case of perfect observability? Proposition 3 says that if $\pi_1^N > \underline{\pi}_1$, any a_1^* between a_1^N and a_1^s can be an ϵ-PBE action in the case of noisy observation, although it is usually not an equilibrium action in the case of perfect observation. If a_1 is perfectly observable, player 2 responds by $a_2^{BR}(a_1)$. So, if player 1 deviates from a_1^* to a_1^s, it induces player 2 to respond from $a_2^{BR}(a_1^*)$ to $a_2^{BR}(a_1^s)$, which makes player 1 better off. This implies that $a_1^* \neq a_1^s$ cannot be an equilibrium outcome. However, if a_1 is imperfectly observed with a noise, deviating from a_1^* to a_1^s does not induce player 2 to respond to $a_2^{BR}(a_1^s)$ but to $a_2^{BR}(\underline{a}_1)$ with some probability, because such a deviation leads to the belief of either $b(y) = a_1^*$ or $\underline{a}_1$, not $b(y) = a_1^s$. Due to the possibility of the pessimistic belief $b(y) = \underline{a}_1$, player 1 may not deviate from the proposed equilibrium action a_1^*, even if $a_1^* \neq a_1^s$.

4 Application: Price Competition Model

We can apply our equilibrium concepts to various real situations. We consider the following specific model which will be helpful to obtain equilibrium strategies and beliefs explicitly.

Two firms who produce differentiated substitutes compete against each other by choosing prices. They face symmetric demand functions. The demand function for the good produced by firm i is given by $q_i = \alpha - \beta p_i + \gamma p_j$ for $j \neq i$, $i = 1, 2$, where $\alpha, \beta, \gamma > 0$ and $2\beta > \gamma$. For simplicity, we assume that the marginal cost is zero. Then, we can compute the profit function of firm i as

$$\pi^i(p_i, p_j) = p_i(\alpha - \beta p_i + \gamma p_j), \tag{1}$$

and the best response function as

$$p_i^{BR} = \frac{\alpha + \gamma p_j}{2\beta}. \tag{2}$$

Note that p_1 and p_2 are strategic complements, because $p_i^{BR}(p_j)$ function has a positive slope.[31]

If the firms cannot observe the price of each other, they end up with the Nash prices $(p_1^N, p_2^N) = (\frac{\alpha}{2\beta-\gamma}, \frac{\alpha}{2\beta-\gamma})$. On the other hand, if firm 1 first chooses its price as a Stackelberg leader and the price is perfectly observed by firm 2, we can expect the well known Stackelberg outcome (p_1^s, p_2^s) to be realized in a subgame perfect equilibrium. Since firm 1 expects firm 2 to respond optimally according to his best response function, $p_2^{BR}(p_1) = \frac{\alpha+\gamma p_1}{2\beta}$, firm 1 will choose p_1 to

$$\max_{p_1} \pi^I(p_1, p_2^{BR}(p_1)) = p_1 \left[\alpha - \beta p_1 + \frac{\gamma(\alpha + \gamma p_1)}{2\beta}\right],$$

leading to equilibrium prices

$$p_1^s = \frac{\beta + \frac{\gamma}{2}}{2\beta^2 - \gamma^2}\alpha, \tag{3}$$

$$p_2^s = \frac{\alpha + \gamma p_1}{2\beta} = \frac{\beta + \frac{\gamma}{2} - \frac{\gamma^2}{4\beta}}{2\beta^2 - \gamma^2}\alpha, \tag{4}$$

and the resulting profit of firm 1

$$\pi^I(p_1^s, p_2^s) = \frac{\alpha^2(2\beta + \gamma)^2}{8\beta(2\beta^2 - \gamma^2)}. \tag{5}$$

If firm 2 cannot observe p_1 perfectly but only with some noise, however, we can show that the Stackelberg leader's price p_1^s cannot be a PBE outcome with even a slight noise, insofar as the noise has a full support $(-\infty, \infty)$. Since every signal (observation) y is possible in equilibrium, firm 2 must respond to any observation y by $p_2^{BR}(b(y))$ which is just the same as $p_2^{BR}(p_1^s)$, since $b(y) = p_1^s$ in equilibrium, for all y. Knowing that firm 2's response will not be affected by firm 1's choice, firm 1 will deviate from p_1^s by slightly cutting its price, which overturns the Stackelberg equilibrium.

We will now resort to alternative equilibrium concepts, ϵ-PBE and limit PBE. To characterize the equilibrium strategies explicitly, we assume that after firm 1 chooses its price p_1, firm 2 observes a signal $y = p_1 + \eta$ where the noise η is normally distributed with mean zero and variance σ^2, i.e., $\eta \sim N(0, \sigma^2)$. So, the density function of y given p_1 is

$$f(y \mid p_1) = \frac{1}{\sqrt{2\pi}\sigma} e^{-\frac{1}{2}\left(\frac{y-p_1}{\sigma}\right)^2}. \tag{6}$$

[31] It is well known that if the players compete in strategic complements such as prices, the second mover advantage rather than the first mover advantage appears in the sense that the equilibrium profit of the second mover is higher than the equilibrium profit of the first mover. See, for example, Gal-Or (1985). However, throughout this section, we maintain the term of the first mover advantage, simply because the first mover's profit is greater than in Nash equilibrium.

For a given ϵ, we can determine the corresponding cutoff value of the observation y_ϵ for the left tail event by using the conditional density function:

$$L(p_1^s; y_\epsilon) = f(y_\epsilon \mid y_1^s) = \epsilon. \tag{7}$$

A cutoff value has the following meaning and interpretation behind it that if a lower signal than the cutoff value y_ϵ is observed, firm 2 believes that firm 1 did not choose the equilibrium price p_1^s, on the ground that the likelihood $L(p_1^s; y)$ is less than ϵ for any $y \le y_\epsilon$. (See Figure 3.) Since $y \sim N(p_1^s, \sigma)$, y_ϵ is determined from the following equilibrium;

$$f(y_\epsilon \mid p_1^s) \equiv \frac{1}{\sqrt{2\pi}\sigma} \exp\left[-\frac{1}{2}\left(\frac{y_\epsilon - p_1^s}{\sigma}\right)^2\right] = \epsilon, \tag{8}$$

i.e.,

$$\left(\frac{y_\epsilon - p_1^s}{\sigma}\right)^2 = -2\ln(\sqrt{2\pi}\sigma\epsilon) > 0. \tag{9}$$

Therefore, for any given $\epsilon > 0$, the confidence interval for y (equilibrium path) is determined by $p_1^s - \rho_\epsilon < y < p_1^s + \rho_\epsilon$, where $\rho_\epsilon = -\sigma\ln(\sqrt{2\pi}\sigma\epsilon)$ is the maximum permissible error. If $y \le p_1^s - \rho_\epsilon$ or $y \ge p_1^s + \rho_\epsilon$, we regard it as an almost off-the-equilibrium event. Note that the cutoff value of y is $y_\epsilon = p_1^s - \rho_\epsilon$, which determines the left tail event.

We now consider the incentive compatibility condition of firm 1. The expected profit of firm 1 can be computed as follows;

$$E[\pi^I] = \int_{-\infty}^{y_\epsilon} \pi^I(p_1, p_2^\epsilon) f(y \mid p_1) dy + \int_{y_\epsilon}^{\infty} \pi^I(p_1, p_2^s) f(y \mid p_1) dy, \tag{10}$$

where $p_2^s = p_2^{BR}(p_1^s)$ and p_2^ϵ is a threat price that will be applied in case that a signal y such that $y \le y_\epsilon$ is observed. Let $\delta \equiv \int_{-\infty}^{y_\epsilon} f(y \mid p_1^s) = \mathbb{P}(y \le y_\epsilon \mid p_1^s)$.[32] Then, the first order condition characterizing the incentive compatibility condition requires

$$\begin{aligned}
\left.\frac{\partial E(\pi^I)}{\partial p_1}\right|_{p_1 = p_1^s} &= \pi_1^I(p_1^s, p_2^\epsilon)\delta + \pi^I(p_1^s, p_2^\epsilon) \int_{-\infty}^{y_\epsilon} f_{p_1}(y \mid p_1^s) dy \\
&\quad + \pi_1^I(p_1^s, p_2^s)(1-\delta) + \pi^I(p_1^s, p_2^s) \int_{y_\epsilon}^{\infty} f_{p_1}(y \mid p_1^s) dy \\
&= 0,
\end{aligned} \tag{11}$$

[32] This corresponds with the p-value in hypothesis testing.

where $\pi^I(p_1,p_2) = p_1(\alpha - \beta p_1 + \gamma p_2)$ and $\pi_1^I(p_1,p_2) = \alpha + \gamma p_2 - 2\beta p_1$.[33] Intuitively, if firm 1 increases its price, it reduces its profit directly but also decreases the probability the observed price falls into the left tail event thereby triggering a punishment. So, p_1^s is to balance the marginal loss in the profit with the marginal benefit from a lower expected punishment.

By using the following Leibniz's rule;

$$\int_a^b \frac{\partial}{\partial x} f(x;t)dt = \frac{d}{dx}\left(\int_a^b f(x,t)dt\right), \tag{12}$$

we have

$$\begin{aligned}
& \pi^I(p_1^s, p_2^\epsilon) \int_{-\infty}^{y_\epsilon} f_{p_1}(y \mid p_1^s)dy + \pi^I(p_1^s, p_2^s) \int_{y_\epsilon}^{\infty} f_{p_1}(y \mid p_1^s)dy \\
= \;& p_1^s(\alpha - \beta p_1^s + \gamma p_2^s) \int_{-\infty}^{\infty} f_{p_1}(y \mid p_1^s)dy + \gamma p_1^s(p_2^\epsilon - p_2^s) \int_{-\infty}^{y_\epsilon} f_{p_1}(y \mid p_1^s)dy \\
= \;& \gamma p_1^s(p_2^\epsilon - p_2^s) \int_{-\infty}^{y_\epsilon} f_{p_1}(y \mid p_1^s)dy,
\end{aligned}$$

since

$$\int_{-\infty}^{\infty} f_{p_1}(y \mid p_1^s)dy = \frac{d}{dp_1}\left(\int_{-\infty}^{\infty} f(y \mid p_1^s)dy\right) = 0.$$

Also, we have

$$\int_{-\infty}^{y_\epsilon} f_{p_1}(y \mid p_1^s)dy = \frac{d}{dp_1}\int_{-\infty}^{y_\epsilon} f(y \mid p_1^s)dy = \left.\frac{dF(y_\epsilon \mid p_1)}{dp_1}\right|_{p_1=p_1^s}. \tag{13}$$

We know that the normal distribution function is

$$F(y_\epsilon \mid p_1) = \frac{1}{2}\left[1 + \operatorname{erf}\left(\frac{y_\epsilon - p_1}{\sqrt{2}\sigma}\right)\right],$$

where $\operatorname{erf}(y) = \frac{2}{\sqrt{\pi}}\int_0^y e^{-t^2}dt$. This implies that

$$\begin{aligned}
\left.\frac{dF(y_\epsilon \mid p_1)}{dp_1}\right|_{p_1=p_1^s} &= \frac{1}{\sqrt{\pi}} e^{-\left(\frac{y_\epsilon - p_1^s}{\sqrt{2}\sigma}\right)^2} \cdot \left(-\frac{1}{\sqrt{2}\sigma}\right) \\
&= -\frac{1}{\sqrt{2\pi}\sigma} e^{-\left(\frac{y_\epsilon - p_1^s}{\sqrt{2}\sigma}\right)^2} \\
&= -L(p_1^s; y_\epsilon) \\
&= -\epsilon.
\end{aligned} \tag{14}$$

[33]In the Appendix, we prove that the second order condition is satisfied.

Thus, equation (11) is reduced to

$$\pi_1^I(p_1^s, p_2^s) = \gamma(p_2^\epsilon - p_2^s)(\delta + \epsilon p_1^s). \tag{15}$$

The left hand side is the marginal loss in the profit from increasing the price and the right hand side is its marginal gain from a fall in the expected punishment. This leads to

$$p_2^\epsilon = p_2^s + \frac{\pi_1^I(p_1^s, p_2^s)}{\gamma(\delta + \epsilon p_1^s)}. \tag{16}$$

Since $\pi_1^I(p_1^s, p_2^s) < 0$, it must be that $p_2^\epsilon < p_2^s$. Also, as $\epsilon, \delta \to 0$, $p_2^\epsilon \to -\infty$. Is p_2^ϵ a credible threat? It is an optimal response to some belief $b = p_1^\epsilon$ such that $p_2^\epsilon = p_2^{BR}(p_1^\epsilon)$. Since $p_2^{BR}(p_1)$ is surjective, p_1^ϵ, which is a threat belief, always exists. A threat to p_2^ϵ is credible, because it is a best response given the threat belief that firm 1 chose a very low price p_1^ϵ. Such a low price threat can frustrate firm 1's incentive to cut its price secretly. Since equation (16) has a solution for any $\epsilon > 0$, the Stackelberg outcome is an ϵ-PBE for any $\epsilon > 0$; hence, a limit PBE. Note that in this specific model, [UP] condition holds, i.e., $\lim_{p_1^\epsilon \to -\infty} \pi(p_1^s, p_2^{BR}(p_1^\epsilon)) = \lim_{p_2^\epsilon \to -\infty} \pi(p_1^s, p_2^\epsilon) = -\infty$.

It deserves noticing that the Stackelberg outcome cannot be a strong limit PBE, because neither $\lim_{\epsilon \to 0} p_1^\epsilon$ nor $\lim_{\epsilon \to 0} p_2^\epsilon$ exists. This implies that the Stackelberg leader's price cannot be a PBE outcome, either, by Proposition 2.

Figure 4 illustrates equilibrium values of firm 2's threat price, p_2^ϵ, with a change in ϵ for parameter values of $\alpha = 1$, $\beta = 3$, $\gamma = 2$ and $\sigma = 1$. It shows that as ϵ becomes smaller, the threat price declines rapidly. This monotonicity reflects the intuition that a smaller probability of a tail event due to a lower ϵ must be accompanied by a harsher punishment in order to maintain the same deterrence power. The same intuition can be applied to a change in the variance of a noise. Figure 5 shows a monotonic decrease in the threat price as σ^2 gets smaller for the value of $\epsilon = 0.1$. Again, a decrease in firm 1's utility due to an excessively low price by firm 2 must be compensated for a smaller probability of a tail event by more precise information.

5 Cheap Talk Model with an Informed Receiver

Our concepts can be also applied to incomplete information games, although our discussion has been restricted only to imperfect information games. For example, we can consider the following cheap talk game in which a sender observes only a noisy signal of the Nature's choice.

There are a sender S, and a receiver R. The state of nature θ is a random variable which is distributed over $\mathbb{R}$. For example, θ could be the quality of a product that a salesperson sells to a consumer. According to Bayes-Laplace's principle of indifference (insufficient reason),[34] we assume that θ is uniformly distributed over $\mathbb{R}$,[35] i.e., players have no information about θ *a priori*. Although neither the sender nor the receiver knows the accurate value of θ, both of them receive a noisy signal on the state of nature $v_i \in V = \mathbb{R}$ for $i = S$ and R where $v_i = \theta + \epsilon_i$, ϵ_i is stochastically independent with θ, and ϵ_i's are independent. In other words, R is also partially informed.[36] We assume that ϵ_i follows a normal distribution with its mean zero and the precision h_i, $i = S, R$, where $h_S > h_R$.[37] The assumption of the inequality in the precision reflects the feature that the sender (expert) has higher expertise about θ than the receiver (amateur).

The game proceeds as follows. First, the state of nature θ is realized and then a sender and a receiver receive a private signal v_S and v_R respectively without knowing θ. After observing private information v_S, S sends a payoff-irrelevant message (cheap talk) $m \in M = \mathbb{R}$ to R.[38] Then, receiving a message $m \in M$, R updates the posterior belief about v_S, $\hat{b}(m)$, and then forms the belief about θ, $b(m, v_R)$, by using m and v_R, where $b : M \times V \to \mathbb{R}$. Based on the belief $b(m, v_R)$, he chooses an action $a \in A (= \mathbb{R})$. A strategy of the receiver determines the sender's payoff as well as his own payoff.

[34]The principle of indifference, so-named by Keynes (1921), specifies that a uniform prior distribution should be assumed when nothing is known about the true state of nature before observable data are available.

[35]Note that we are assuming an improper prior distribution.

[36]This cheap talk game with a partially informed receiver was first considered by Seidmann (1990), but his analysis and all the subsequent analyses rely crucially on the assumption of finite or bounded supports of the information. See, for example, Watson (1996), Lai (2014) and Ishida and Shimizu (2016). The analysis of this section is unique in that all the noisy variables have unbounded supports.

[37]Assuming a normal distribution is to avoid shifting support of the distribution. Allowing shifting support could make the analysis trivial because some signals off the equilibrium support, which cannot occur in equilibrium, could reveal that S deviated from the equilibrium.

[38]Since the cheap talk message of the sender, m, is payoff-irrelevant by the definition of cheap talk, the payoffs of the players (U^S and U^R) which are described below should not depend on m. This is distinguished from Pitchick and Schotter (1987). In their model, an expert makes a binding recommendation, for example, about the price, so it is not cheap talk, whereas we consider an unbinding recommendation of an expert (for example, about the quality) thereby making the payoff of the receiver not directly depend on the recommendation m.

The payoff to S is given by a continuously differentiable function $U^S : A \to \mathbb{R}$ and the payoff to R is given by twice continuously differentiable function $U^R : A \times \Theta \to \mathbb{R}$. Specifically, we assume that (1) $U^S(a) = u(a)$ where $u' > 0$, $u'' \leq 0$ and $\lim_{a \to -\infty} u(a) = -\infty$, and (2) $U^R(a, \theta) = -(a - \theta)^2$. The assumption that $\lim_{a \to -\infty} u(a) = -\infty$ is a sort of [UP] condition to ensure that punishments are unbounded. The receiver's utility function implies that it has a unique maximum in a for all θ and the maximizer of U^R, denoted by $a^R(\theta)$, is strictly increasing in θ. Independence of the sender's utility function on θ means that S has transparent motives, and the utility which is increasing in a means that S has monotone motives. The monotonic increase of $a^R(\theta)$ in θ means that the receiver will want to buy more units of high θ which can be interpreted as quality. A typical example that corresponds to these assumptions is a situation in which a salesperson gets paid based on the quantity he sells, so that the salesperson's utility is increasing with respect to the consumer's purchasing choice regardless of θ.

A strategy for S specifies a signaling rule given by a measurable function $s : V \to M$. A strategy for R is an action rule given by a function $\alpha : M \times V \to A$. We will resort to several equilibrium concepts, perfect Bayesian equilibrium (PBE), ϵ-perfect Bayesian equilibrium (ϵ-PBE) and limit perfect Bayesian equilibrium (limit PBE).

We are mainly interested in whether a fully revealing equilibrium can be possible in this model in which R has some information about θ. The following proposition says that the fully revealing equilibrium cannot exist, insofar as we use PBE as our equilibrium concept.

Proposition 5. *There is no fully revealing PBE.*

The important insight behind this proposition is that any message m is possible in equilibrium, so that R must believe no matter what message is conveyed by the sender, no matter how unlikely the message is. There is no way to punish and discipline a sender even if it is crystal clear that she has an incentive to lie by inflating m, because the concept of PBE is required to respect any message possible in equilibrium.

We will argue below that this belief is not reasonable in the following sense. Suppose the sender sends a truthful message in equilibrium. Also, suppose that R who has a private signal v_R finds out that the sender's message m is too far from his own information (too high compared to v_R). There are two possibilities. One possibility is that S honestly reported her signal but R happened to receive an exceptionally low value of v_R. The other possibility is that S exaggerated the message above the true value of v_S, although the actual value of v_S was low. Unfortunately, the weak consistency condition of PBE requires us to count on the first scenario, although it is a very unlikely event in the sense that the likelihood $L(m; v_R) \equiv f(v_R \mid m)$ is very low, i.e., the signal v_R is very unlikely to occur if the sender's message m is true. It seems reasonable to us, though, to reject the hypothesis that S sent a truthful message if m is too far from the receiver's signal v_R, because it is very unlikely ($L(m; v_R) \leq \epsilon$ for some small $\epsilon > 0$) even though it is possible ($L(m; v_R) \neq 0$). So, if R rejects the hypothesis that S sent an equilibrium message, the message which R received from the sender can be regarded as an off-the-equilibrium message, and, therefore, the posterior belief need not be pinned down by the Bayes' law; in this case, any arbitrary belief might be allowed. This is what we call ϵ-likelihood consistency. We can adapt the definitions of our equilibrium concepts to this cheap talk game.

Definition 6. *An assessment* $((s^*(v_S)), \alpha^*(b), \hat{b}(m), b(m, v_R, \epsilon))$ *satisfies* ϵ*-likelihood consistency iff for any* m *and* v_R, $b(m; v_R) = \hat{\theta}(m) \equiv \frac{h_S m + h_R v_R}{h_S + h_R}$ *if* $L(m; v_R) > \epsilon$, *and* $b(m, v_R; \epsilon) \in \mathbb{R}$ *can be arbitrary if* $L(m; v_R) \leq \epsilon$, *where* $\hat{\theta}(m)$ *is the maximum likelihood estimator for* θ.[39]

Definition 7. *An assessment* $((s^*(v_S)), \alpha^*(b), \hat{b}(m), b(m, v_R, \epsilon))$ *is an* ϵ*-perfect Bayesian equilibrium (*ϵ*-PBE), for some* $\epsilon > 0$, *if it satisfies sequential rationality and* ϵ*-likelihood consistency, i.e.,*

(2-I) $s^*(v_S) \in \arg\max_m \int_{-\infty}^{\infty} U^{S_i}(\alpha^*(m)) f(v_R \mid v_S) dv_R$,

(2-II) $\alpha^*(m) \in \arg\max_a U^R(a, b(m, v_R; \epsilon))$,

(2-III) $b(m, v_R; \epsilon) = \hat{\theta}(m)$ *if* $L(m; v_R) > \epsilon$ *and* $b(m, v_R; \epsilon)$ *can be arbitrary if* $L(m; v_R) \leq \epsilon$.

Definition 8. *An assessment* $((s^*(v_S)), \alpha^*(b), \hat{b}(m))$ *is a limit PBE if for any small* $\epsilon > 0$ *such that* $\epsilon \leq \bar{\epsilon}$ *for some* $\bar{\epsilon} > 0$, *there exists* $b(m, v_R, \epsilon)$ *such that* $((s^*(v_S)), \alpha^*(b), \hat{b}(m), b(m, v_R, \epsilon))$ *is an* ϵ*-PBE.*

Now, we can apply the equilibrium concepts to our model. To see the possibility of a fully revealing equilibrium, we will concentrate on the following specific form of strategy profile;

(3-I) S with v_S announces $m = v_S$.

(3-II) R chooses $\alpha(b) = b$.

(3-III) R believes $b(m, v_R) = \hat{\theta}(m, v_R) \equiv \frac{h_S m + h_R v_R}{h}$ if $m - v_R < \rho$ and believes $b(m, v_R) = \tilde{\theta}$ if $m - v_R \geq \rho$ for some $\rho > 0$ where $h = h_S + h_R$ and $\tilde{\theta} < m$ is some constant in $\mathbb{R}$.

[39] It is well known that under the normality assumption of error terms, the maximum likelihood estimator is equivalent to the (generalized) Bayesian estimator minimizing the loss function defined by the mean square error, which is the posterior mean.

R's action rule given by (3-II) together with the belief (3-III) will be called a "crosschecking strategy", and ρ will be called permissible deviation between m and v_R. Note that there is no off-the-equilibrium message, because any message can occur even if the sender tells the truth, as long as ϵ_i follows a normal distribution over $(-\infty, \infty)$.

We will now invoke ϵ-PBE in this game. Let us first consider the belief and the optimal strategy of R with private information v_R. Given $\epsilon > 0$, define v_R^ϵ by $L(m, v_R^\epsilon) = \epsilon$ and ρ_ϵ by $\rho_\epsilon = m - v_R^\epsilon$. Then, it turns out that the belief given in (3-III) with $\rho = \rho_\epsilon$ satisfies ϵ-likelihood consistency defined by (2-III), because $m - v_R \geq \rho_\epsilon$ ($v_R \leq v_R^\epsilon$) implies $L(m, v_R) \leq \epsilon$ so that it is not necessary that $b(m, v_R) = \hat{\theta}(m, v_R)$.

Since it is obvious that (3-II) is R's optimal decision, it is enough to focus on the optimal decision of the sender.

S will maximize

$$U^S(m; v_S) = \int_{-\infty}^{m-\rho_\epsilon} u(\tilde{\theta}) f(v_R \mid v_S) dv_R + \int_{m-\rho_\epsilon}^{\infty} u(\hat{\theta}(m, v_R)) f(v_R \mid v_S) dv_R. \tag{17}$$

The economic reasoning behind this formula goes as follows. The first term represents the punishment that the sender would get when v_R is very low ($v_R \leq m - \rho_\epsilon$). The second term indicates her utility when v_R falls into a normal confidence region ($v_R > m - \rho_\epsilon$). Thus, the effect of inflating the message on the sender's utility is

$$\begin{aligned}
\frac{\partial U^S}{\partial m} &= u(\tilde{\theta}) f(m - \rho_\epsilon \mid v_S) + \int_{m-\rho_\epsilon}^{\infty} u'(\hat{\theta}(m, v_R)) \frac{\partial \hat{\theta}}{\partial m} f(v_R \mid v_S) dv_R - u(\hat{\theta}(m, m - \rho_\epsilon)) f(m - \rho_\epsilon | v_S) \\
&= \frac{h_S}{h} \int_{m-\rho_\epsilon}^{\infty} u'(\hat{\theta}(m, v_R)) f(v_R | v_S) dv_R + (u(\tilde{\theta}) - u(\hat{\theta}(m, m - \rho_\epsilon))) f(m - \rho_\epsilon | v_S), \qquad (18)
\end{aligned}$$

since $\frac{\partial \hat{\theta}}{\partial m} = \frac{h_S}{h}$. The first term is the effect of utility increases in normal cases due to the inflated announcement and the second term is the utility loss that she is expected to bear due to an increase in the punishment probability by increasing his announcement marginally.

The incentive compatibility condition requires $\frac{\partial U^S}{\partial m}|_{m=v_S} = 0$, implying that

$$u(\hat{\theta}(v_S, v_S - \rho_\epsilon)) - u(\tilde{\theta}) = \frac{h_S}{h} \frac{\int_{v_S-\rho_\epsilon}^{\infty} u'(\hat{\theta}(v_S, v_R)) f(v_R | v_S) dv_R}{f(v_S - \rho_\epsilon)}. \tag{19}$$

The left hand side is a gain from inflating m, while the right hand side is the loss due to an increase in the expected penalty. Therefore, we have

$$u(\tilde{\theta}) = u(\hat{\theta}(v_S, v_S - \rho_\epsilon)) - \frac{h_S}{h} \frac{\int_{v_S-\rho_\epsilon}^{\infty} u'(\hat{\theta}(v_S, v_R)) f(v_R | v_S) dv_R}{f(v_S - \rho_\epsilon)}. \tag{20}$$

Since u is continuous, $u' > 0$ and $\lim_{a \to -\infty} u(a) = -\infty$, for any v_S, there exists $\tilde{\theta} < \hat{\theta}(v_S, v_S - \rho_\epsilon)$ that satisfies equation (20). Since S has no incentive to lie, the above strategy is ϵ-PBE. Also, since $\tilde{\theta}(\epsilon)$ exists for any $\epsilon > 0$, the fully revealing outcome is a limit PBE.

Proposition 6. *A fully revealing outcome with the crosschecking strategy and the belief given by (3-II) and (3-III) satisfies ϵ-PBE if $\rho_\epsilon = m - v_R^\epsilon$ where $L(m, v_R^\epsilon) = \epsilon$, and furthermore it is a limit PBE.*

The proof is omitted, since the incentive compatibility condition of the sender and the ϵ-likelihood consistency condition of ϵ-PBE are all checked above.

This result implies that truth-telling is possible in equilibrium even in this cheap talk game with transparent and monotone motives and with unbounded support of all signals, if the receiver is somewhat informed. The main message of this proposition is that honesty is not contradictory with the Bayesian approach, because the fully revealing outcome is supported as an ϵ-PBE for any arbitrarily small $\epsilon > 0$.

6 Conclusion

In this paper, we introduced two equilibrium concepts, ϵ-PBE and limit PBE that slightly weaken the consistency requirement of PBE. These concepts turned out to be useful in dynamic games with noisy signals that have unbounded supports. We showed that the first mover advantage that disappeared in those games could be recovered by invoking those equilibrium concepts.

However, we admit that our result of the existence of a limit PBE depends crucially on [UP] condition which makes the "boiling-in-oil" strategy feasible. The "boiling-in-oil" strategy is optimal given the extremely pessimistic belief, but the punisher can be also severely penalized by the belief. We believe that it is indeed a strong sufficient condition but not necessary for the existence of a limit PBE, and thus believe that a milder condition could guarantee its existence. We hope to see more developments in those equilibrium concepts in a near future.

Appendix

Proof of Proposition 1: (i) If a_1^* is a limit PBE, for any $\epsilon > 0$ such that $\epsilon < \bar{\epsilon}$ for some $\bar{\epsilon} > 0$, there exists $b(y; \epsilon)$ such that $(a_1^*, a_2^*(b), b(y; \epsilon))$ is an ϵ-PBE. So, the proof is done.

(ii) If $(a_1^*, a_2^*(b), b(y))$ is a PBE, it is sufficient to show that for any $\epsilon > 0$, one can take $b(y;\epsilon)$ satisfying ϵ-PBE. For any $\epsilon > 0$, take $b(y;\epsilon) = b(y)$. Then, $(a_1^*, a_2^*(b), b(y;\epsilon)) = (a_1^*, a_2^*(b), b(y))$ is an ϵ-PBE, because $b(y;\epsilon) = b(y)$ satisfies consistency, implying ϵ-likelihood consistency. Since this holds for any $\epsilon > 0$, a_1^* is a limit PBE. ■

Proof of Proposition 2: (i) ($\Longleftarrow$) The proof is similar to the proof of Proposition 1(ii). If $(a_1^*, a_2^*(b), b(y))$ is a PBE, for any ϵ_n such that $\epsilon_n \to 0$, take $b(y;\epsilon_n) = b(y)$. Then, by Proposition 1, $(a_1^*, a_2^*(y;\epsilon_n), b(y;\epsilon_n)) = (a_1^*, a_2^*(b), b(y))$ is ϵ_n-PBE. Since it is obvious that $\lim_{n\to\infty} b(y;\epsilon_n) = b(y)$, it is a strong limit PBE.

(ii) ($\Longrightarrow$) If $(a_1^*, a_2^*(b), b(y))$ is a strong limit PBE, for any $\epsilon(>0)$, there exists $b(y;\epsilon)$ such that $\lim_{\epsilon\to 0} b(y;\epsilon) = b(y)$ and $a_2^*(b)$ is a BR to a_1^* given $b(y;\epsilon)$, i.e.,

$$U^{II}(a_1^*, a_2^*(b(y:\epsilon)) \geq U^{II}(a_1^*, a_2), \forall a_2 \in A_2. \tag{21}$$

Since this holds for any $\epsilon > 0$, we take limits to get

$$\begin{aligned}
\lim_{\epsilon\to 0} U^{II}(a_1^*, a_2^*(b(y;\epsilon)) &= U^{II}(a_1^*, \lim_{\epsilon\to 0} a_2^*(b(y;\epsilon))) \qquad \text{(by continuity of } U^{II}) \\
&= U^{II}(a_1^*, a_2^*(\lim_{\epsilon\to 0} b(y;\epsilon)), \qquad \text{(by continuity of } a_2^*(b) \text{ in } b) \\
&= U^{II}(a_1^*, a_2^*(b(y))) \qquad (\because \lim_{\epsilon\to 0} b(y;\epsilon) = b(y)) \\
&\geq U^{II}(a_1^*, a_2), \forall a_2 \in A_2. \qquad \text{(by Inequality (21))}
\end{aligned}$$

(Note that $a_2^*(b)$ is continuous in b because $a_2^*(\cdot) = a_2^{BR}(\cdot)$.) Hence, $(a_1^*, a_2^*(b), b(y))$ is a PBE. ■

Proof of Proposition 3: For any $\epsilon > \delta$, let $a_1^* \in A^*$ be player 1's action of ϵ-PBE. By suppressing the subscript of player 1 in the payoff function, we can compute the equilibrium payoff of player 1 as

$$\Pi^* \equiv (1-\delta)\pi^* + \delta\underline{\pi}(a_1^*),$$

where $\pi^* = \pi(a_1^*, a_2^{BR}(a_1^*))$, since ϵ-likelihood consistency allows $b(y) = \underline{a}_1$ if $y \neq a_1^*$ with probability δ. Her payoff from the most profitable deviation is

$$\max_{a'_1 \neq a_1^*} \Pi \equiv \left[\left(1 - \frac{\delta}{n-1}\right)\pi(a'_1, a_2^{BR}(\underline{a}_1)) + \frac{\delta}{n-1}\pi(a'_1, a_2^{BR}(a_1^*)),\right]$$

because player 2 observes $y = a_1^*$ with probability $\frac{\delta}{n-1}$ if $a_1 \neq a_1^*$.

We have $\pi^N > \underline{\pi}$ since $a_1^* \in A^*$. Therefore, $\pi^* > \pi(a'_1, a_2(\underline{a}_1))$ for any $a'_1 \neq a_1^*$. This implies that there exists small $\delta > 0$ such that $\Pi^* > \Pi$. Take such δ and denote it by $\bar{\delta}$.

Then, for any $\delta \leq \bar{\delta}$, $(a_1^*, a_2^{BR}(b(y)))$ is an ϵ-PBE, since $\epsilon > \delta$, where $b(a_1^*) = a_1^*$ and $b(y) = \underline{a}_1$ if $y \neq a_1^*$. ■

Proof of Proposition 4: Without loss of generality, assume that $\lim_{b \to -\infty} U^I(a_1^s, a_2^{BR}(b)) = -\infty$.

Lemma 1. *$a_2^{BR}(a_1)$ is surjective.*

Proof. It is enough to show that $a_2^{BR}(a_1)$ is not bounded by K for some $K \in (0, \infty)$, i.e., $|a_2^{BR}(a_1)| \leq K$. Suppose $a_2^{BR}(a_1)$ is bounded. By [UP] condition, $\lim_{b \to -\infty} U^I(a_1^s, a_2^{BR}(b)) = -\infty$. Then, by continuity of U^I, we have $\lim_{b \to -\infty} U^I(a_1^s, a_2^{BR}(b)) = U^I(a_1, \lim_{b \to -\infty} a_2^{BR}(b)) = -\infty$. This is not possible, because $U^I(a_1, a_2)$ is continuous for all $a_2 \in \mathbb{R}$ and $a_2^{BR}(b)$ is bounded. Therefore, $a_2^{BR}(a_1)$ must be not bounded. The case that $\lim_{b \to \infty} U^I(a_1^s, a_2^{BR}(b)) = -\infty$ is similar, so we omit the proof for the case. ■

For any fixed $\epsilon > 0$, take y_ϵ such that $L(a_1^s, y_\epsilon) = f(y_\epsilon \mid a_1^s) = \epsilon$ where $y_\epsilon < a_1^s$. Then, ϵ-PBE allows that player 2 punishes player 1 if he observes $y < y_\epsilon$ by assigning the posterior belief $b(a_1 \mid y) = a_1^\epsilon < a_1^s$ and otherwise (if he observes $y \geq y_\epsilon$) $b(a_1 \mid y) = a_1^s$. Let $\delta \equiv \int_{-\infty}^{y_\epsilon} f(y \mid a_1^s) = \mathbb{P}(y \leq y_\epsilon \mid a_1^s)$.

Suppressing the superscript I of U^I, we can compute player 1's expected utility as

$$E[U] = U(a_1, a_2^{BR}(a_1^\epsilon)) \int_{-\infty}^{y_\epsilon} f(y \mid a_1) dy + U(a_1, a_2^{BR}(a_1^s)) \int_{y_\epsilon}^{\infty} f(y \mid a_1) dy. \tag{22}$$

The first order condition requires

$$\left. \frac{\partial E[U]}{\partial a_1} \right|_{a_1 = a_1^s} = \varphi_1 + \varphi_2 = 0, \tag{23}$$

where

$$\varphi_1 = U_1(a_1^s, a_2^{BR}(a_1^\epsilon)) \int_{-\infty}^{y_\epsilon} f(y \mid a_1^s) dy + U_1(a_1^s, a_2^{BR}(a_1^s)) \int_{y_\epsilon}^{\infty} f(y \mid a_1^s) dy,$$

$$\varphi_2 = U(a_1^s, a_2^{BR}(a_1^\epsilon)) \int_{-\infty}^{y_\epsilon} f_{a_1}(y \mid a_1^s) dy + U(a_1^s, a_2^{BR}(a_1^s)) \int_{y_\epsilon}^{\infty} f_{a_1}(y \mid a_1^s) dy.$$

By using $a_2^s = a_2^{BR}(a_1^s)$, $a_2^\epsilon = a_2^{BR}(a_1^\epsilon)$, $\int_{-\infty}^{y_\epsilon} f(y \mid a_1^s) dy = \delta$ and $\int_{-\infty}^{y_\epsilon} f_{a_1}(y \mid a_1^s) dy = -\epsilon$ from (13) and (14), we can simplify equation (23) into

$$\epsilon(U(a_1^s, a_2^s) - U(a_1^s, a_2^\epsilon)) = \delta(U_1(a_1^s, a_2^s) - U_1(a_1^s, a_2^\epsilon)) - U_1(a_1^s, a_2^s). \tag{24}$$

The left hand side (LHS) of (24) is a reduction in the penalty probability (ϵ) due to a marginal increase in a_1 times the magnitude of the penalty ($U(a_1^s, a_2^s) - U(a_1^s, a_2^\epsilon)$), which is a reduction in the expected cost from a marginal increases in a_1. The right hand side (RHS) of (24) is a reduction in the expected gain due to a marginal deviation from a_1^s. It is straightforward to see that (25) follows from (24).

$$U(a_1^s, a_2^s) - U(a_1^s, a_2^\epsilon) = \frac{\delta}{\epsilon}(U_1(a_1^s, a_2^s) - U_1(a_1^s, a_2^\epsilon)) - \frac{U_1(a_1^s, a_2^s)}{\epsilon}. \tag{25}$$

Note that

$$\begin{aligned}
\lim_{\epsilon \to 0} \frac{\delta}{\epsilon} &= \lim_{y_\epsilon \to -\infty} \frac{\int_{-\infty}^{y_\epsilon} f(y \mid a_1^s) dy}{f(y_\epsilon \mid a_1^s)} \\
&= \lim_{z \to -\infty} \frac{\int_{-\infty}^{z} \frac{1}{\sqrt{2\pi}} e^{-\frac{x^2}{2}} dx}{\frac{1}{\sqrt{2\pi}} e^{-\frac{z^2}{2}}} \qquad (\because \text{ by letting } z = \frac{y_\epsilon - a_1^s}{\sigma}) \\
&= \lim_{z \to -\infty} \frac{\frac{1}{\sqrt{2\pi}} e^{-\frac{z^2}{2}}}{-\frac{z}{\sqrt{2\pi}} e^{-\frac{z^2}{2}}} \qquad (\because \text{ by L'Hospital's rule}) \\
&= \lim_{z \to -\infty} -\frac{1}{z} \\
&= 0.
\end{aligned} \tag{26}$$

Since $\frac{\delta}{\epsilon} \to 0$, $|U_1(a_1^s, a_2^\epsilon)| \leq M$ for some $M > 0$ and $\frac{U_1(a_1^s, a_2^s)}{\epsilon}$ is finite for any $\epsilon > 0$, RHS of (25) is bounded, so there must exist $a_2^\epsilon \in \mathbb{R}$ that satisfies (24) by the intermediate value theorem, due to continuity of U with respect to a_2, [UP] condition and Lemma 1.

Since $a_2^{BR}(a_1)$ is surjective by Lemma 1, we can define $a_1^\epsilon \in (-\infty, \infty)$ such that $a_2^\epsilon = a_2^{BR}(a_1^\epsilon)$ for any $a_2^\epsilon \in \mathbb{R}$. Then, for any $\epsilon > 0$, $(a_1^s, a_2^{BR}(b), b(y; \epsilon))$ is an ϵ-PBE where

$$b(y;\epsilon) = \begin{cases} a_1^s & \text{if } L(a_1^s; y) > \epsilon \\ a_1^\epsilon & \text{if } L(a_1^s; y) \leq \epsilon, \end{cases}$$

$$a_2^{BR}(b(y;\epsilon)) = \begin{cases} a_2^s & \text{if } L(a_1^s; y) > \epsilon \\ a_2^\epsilon & \text{if } L(a_1^s; y) \leq \epsilon, \end{cases}$$

Since there exist a_1^ϵ and a_2^ϵ for any ϵ, a_1^s is a limit PBE.

■

Lemma 2. *In the price competition model, the Stackelberg price p_1^* satisfies the second order condition of optimization, i.e.,* $\left.\frac{\partial^2 E(\pi^I)}{\partial p_1^2}\right|_{p_1 = p_1^s} < 0$.

Proof. The second order condition requires

$$\frac{\partial^2 E(\pi^I)}{\partial p_1^2} = \psi_1 + \psi_2 + \psi_3,$$

where

$$\psi_1 = \pi_{11}(p_1^s, p_2^\epsilon) \int_{-\infty}^{y_\epsilon} f(y \mid p_1) dy + \pi_{11}(p_1^s, p_2^s) \int_{y_\epsilon}^{\infty} f(y \mid p_1) dy,$$

$$\psi_2 = 2 \left[\pi_1(p_1^s, p_2^\epsilon) \int_{-\infty}^{y_\epsilon} \frac{\partial f(y \mid p_1)}{\partial p_1} dy + \pi_1(p_1^s, p_2^s) \int_{y_\epsilon}^{\infty} \frac{\partial f(y \mid p_1)}{\partial p_1} dy \right],$$

$$\psi_3 = \pi(p_1^s, p_2^\epsilon) \int_{-\infty}^{y_\epsilon} \frac{\partial^2 f(y \mid p_1)}{\partial p_1^2} dy + \pi(p_1^s, p_2^s) \int_{y_\epsilon}^{\infty} \frac{\partial^2 f(y \mid p_1)}{\partial p_1^2} dy.$$

Tedious calculations yield

$$\psi_1 = -2\beta \int_{-\infty}^{\infty} f(y \mid p_1) dy = -2\beta < 0,$$

$$\psi_2 = 2\gamma(p_2^\epsilon - p_2^s) \int_{-\infty}^{y_\epsilon} f_{p_1}(y \mid p_1^s) dy = -2\epsilon\gamma(p_2^\epsilon - p_2^s) < 0,$$

$$\begin{aligned} \psi_3 &= \pi(p_1^s, p_2^s) \int_{-\infty}^{\infty} \frac{\partial^2 f(y \mid p_1)}{\partial p_1^2} dy + \gamma p_1^s \int_{-\infty}^{y_\epsilon} \frac{\partial^2 f(y \mid p_1)}{\partial p_1^2} dy \\ &= \pi(p_1^s, p_2^s) \frac{\partial}{\partial p_1} \int_{-\infty}^{\infty} \frac{\partial f(y \mid p_1)}{\partial p_1} dy + \gamma p_1^s \frac{\partial}{\partial p_1} \int_{-\infty}^{y_\epsilon} \frac{\partial f(y \mid p_1)}{\partial p_1} dy. \end{aligned}$$

Since $\int_{-\infty}^{\infty} \frac{\partial f(y|p_1)}{\partial p_1} dy = \int_{-\infty}^{\infty} f(y \mid p_1) \left(\frac{y-p_1}{\sigma}\right) dy = E[y - p_1 \mid p_1]/\sigma = 0$, ψ_3 is simplified to

$$\begin{aligned} \psi_3 &= \gamma p_1^s \frac{\partial}{\partial p_1} \int_{-\infty}^{y_\epsilon} f(y \mid p_1) \left(\frac{y - p_1}{\sigma}\right) dy \\ &= \gamma p_1^s \frac{\partial}{\partial p_1} \int_{-\infty}^{\frac{y_\epsilon - p_1^s}{\sigma}} z\phi(z) dz \\ &= \gamma p_1^s \frac{\partial}{\partial p_1} \mathbb{P}[z \leq \frac{y_\epsilon - p_1^s}{\sigma}] \\ &< 0, \end{aligned}$$

where z follows the standard normal distribution and $\phi(z)$ is the standard normal distribution that $z = \frac{y_\epsilon - p_1^s}{\sigma}$ follows. This completes the proof. ■

Proof of Proposition 5: Suppose there is a fully revealing PBE in which $s^*(v_S) = v_S$ for all v_S. Then, the weak consistency condition of PBE requires that $\hat{b}(v_S) = v_S$ and $b(m, v_R) = \hat{\theta}(m) \equiv \frac{h_S m + h_R v_R}{h}$, since any m is possible in equilibrium due to the assumption of the unbounded support of ϵ_S and so any m must be taken seriously to update the belief of θ. Based on this belief, the optimal response of R is $\alpha^* = \hat{\theta}(m) = \frac{h_S m + h_R v_R}{h}$. This is a contradiction to the optimality of $m = v_S$, because S would benefit from deviating to $m' = m + \epsilon$ for $\epsilon > 0$ due to the monotone motive assumption, i.e., $u'(a) > 0$. ■

References

[1] Bagwell, K. 1995, Commitment and Observability in Games, Games and Economic Behavior 8, 271-280.

[2] Banks, J. and J. Sobel, 1987, Equilibrium Selection in Signalling Games, Econometrica 55, 647-662.

[3] Bonanno, G., 2013, AGM-Consistency and Perfect Bayesian Equilibrium. Part I: Definition and Properties, International Journal of Game Theory 42, 567-593.

[4] Cho, I. and D. Kreps, 1987, Signalling Games and Stable Equilibria, Quarterly Journal of Economics 102, 179-221.

[5] Farrell, J., 1993, Meaning and Credibility in Cheap-Talk Games, Games and Economic Behavior 5, 514-531.

[6] Fudenberg, D. and J. Tirole, 1991a, Perfect Bayesian Equilibria and Sequential Equilibria, Journal of Economic Theory 53, 236-260.

[7] Fudenberg, D. and Tirole, J. 1991b, Game Theory, Cambridge, MA, MIT Press.

[8] Gal-Or, E., 1985, First Mover and Second Mover Advantages, International Economic Review 26, 649-653.

[9] Gal-Or, E., 1987, First Mover Advantages with Private Information, Review of Economic Studies 54, 279-292.

[10] Grether, D., 1992, Testing Bayes Rule and the Representativeness Heuristic: Some Experimental Evidence, Journal of Economic Behavior and Organization 17, 31-57.

[11] Griffin, D. and A. Tversky, 1992, The Weighting of Evidence and the Determinants of Confidence, Cognitive Psychology 24, 411-435.

[12] Grossman, S. and M. Perry, 1986, Perfect Sequential Equilibrium, Journal of Economic Theory 39, 97-119.

[13] Hertzendorf, M. N., 1993, I'm Not a High-Quality Firm - But I Play One on TV, Rand Journal of Economics 24, 236-247.

[14] Holt, C. and A. Smith, 2009, An Update on Bayesian Updating, Journal of Economic Behavior and Organization 69, 125-134.

[15] Ishida, J. and T. Shimizu, 2016, Cheap Talk with an Informed Receiver, Economic Theory Bulletin 4, 61-72.

[16] Kahneman, D. and A. Tversky, 1972, Subjective Probability: A Judgement of Representativeness, Cognitive Psychology 3, 430-454.

[17] Keynes, J. M., 1921, A Treatise on Probability, London: Macmillan.

[18] Kohlberg, E., 1990, Refinement of Nash Equilibrium: The Main Ideas, in T. Ichiishi, A. Neyman and Y. Tauman (eds.), Game Theory and Applications, San Diego: Academic Press.

[19] Kreps D. and R. Wilson, 1982, Sequential Equilibria, Econometrica, 50, 863-894.

[20] Lai, E., 2014, Expert Advice for Amateurs, Journal of Economic Behavior and Organization 103, 1-16.

[21] Maggi, G., 1999, The Value of Commitment with Imperfect Observability and Private Information, Rand Journal of Economics 30, 555-574.

[22] Mas-Colell, A., M. Whinston, and J. Green, 1995, Microeconomic Theory, Oxford University Press.

[23] Matthews, S. A. and L. Mirman, 1983, Equilibrium Limit Pricing: The Effects of Private Information and Stochastic Demand, Econometrica 51, 981-995.

[24] Myerson, R., 1978, Refinements of the Nash Equilibrium Concept, International Journal of Game Theory 7, 73-80.

[25] Myerson, R. B. and P. J. Reny, 2019, Perfect Conditional ϵ-Equilibria of Multi-Stage Games with Infinite Sets of Signals and Actions, Mimeo, University of Chicago.

[26] Ortoleva, P., 2012, Modeling the Change of Paradigm: Non-Bayesian Reactions to Unexpected News, American Economic Review 102, 2410-2436.

[27] Pitchik, C. and A. Schotter, 1987, Honesty in a Model of Strategic Information Transmission, American Economic Review 77, 1032-1036.

[28] Radner, R., 1980, Collusive Behavior in Non-Cooperative Epsilon-Equilibria of Oligopolists with Long But Finite Lives, Journal of Economic Theory, 22, 136—154.

[29] Rasmusen, E., 1994, Games and Information, (2nd edition), Cambridge: Blackwell.

[30] Seidmann, D., 1990, Effective Cheap Talk with Conflicting Interests, Journal of Economic Theory 50, 445-458.

[31] Selten, R., 1965, Spieltheoretische Behandlung eines Oligopolmodells mit Nachfragetragheit. Zeitschrift fur die gesamte Staatswissenschaft 12, 201-324.

[32] Selten, R., 1975, Reexamination of the perfectness concept for equilibrium points in extensive games, International Journal of Game Theory 4, 25-55.

[33] Selten, R., 1983, Evolutioinary Stability in Extensive Two-Person Games, Mathematical Social Sciences 5, 269-363.

[34] Sun, L., 2019, Hypothesis Testing Equilibrium in Signaling Games, Mathematical Social Sciences 100, 29-34.

[35] van Damme, E. and Hurkens, S., 1997, Games with Imperfectly Observable Commitment, Games and Economic Behavior 21, 282-308.

[36] Watson, J., 1996, Information Transmission When the Informed Party Is Confused, Games and Economic Behavior 12, 143-161.

[37] Watson, J., 2017, A General, Practicable Definition of Perfect Bayesian Equilibrium, Mimeo, University of California, San Diego.

		II	
		S	C
I	S	$(5,2)$	$(3,1)$
	C	$(6,3)$	$(4,4)$

Figure 1: Bagwell's Example

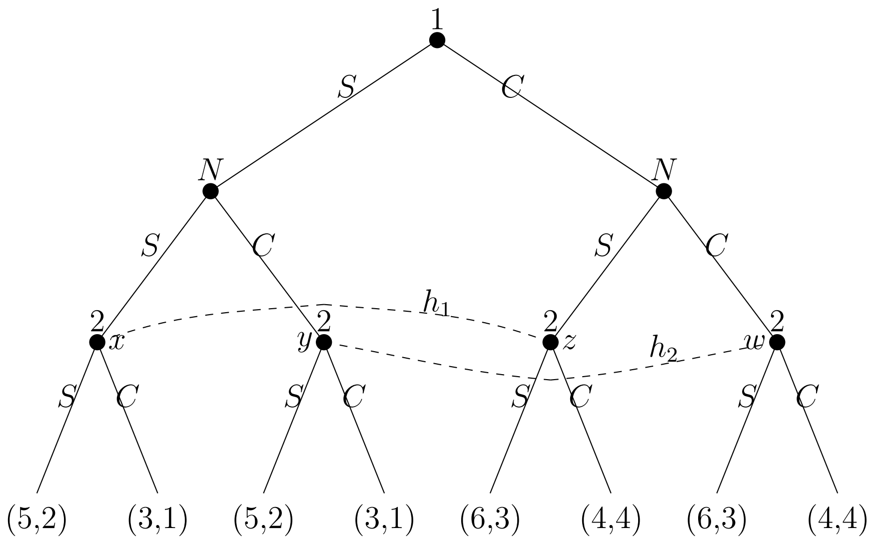

Figure 2. Extensive Form of Bagwell's Example

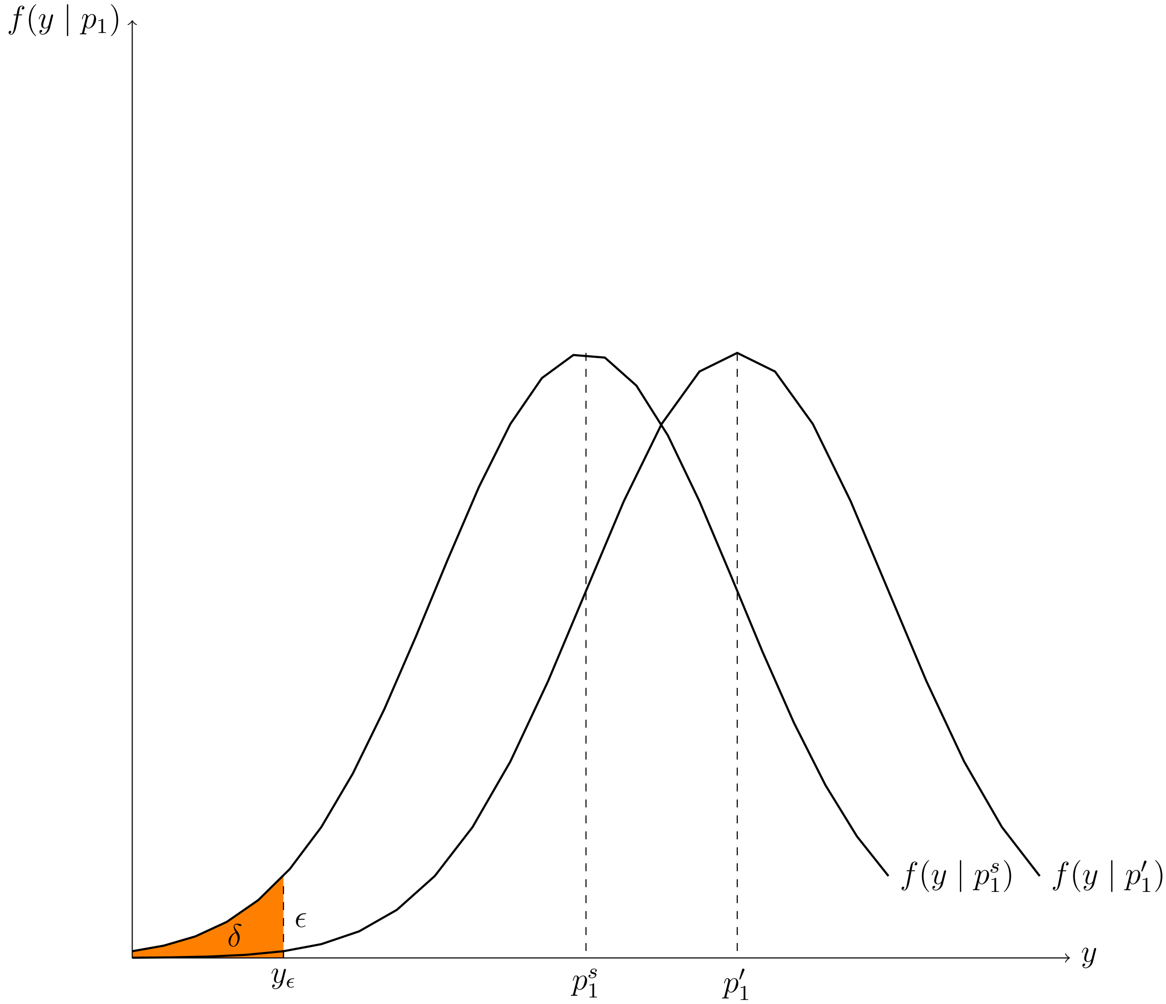

Figure 3. Shift in the Conditional Density Function ($p_1' > p_1^s$)

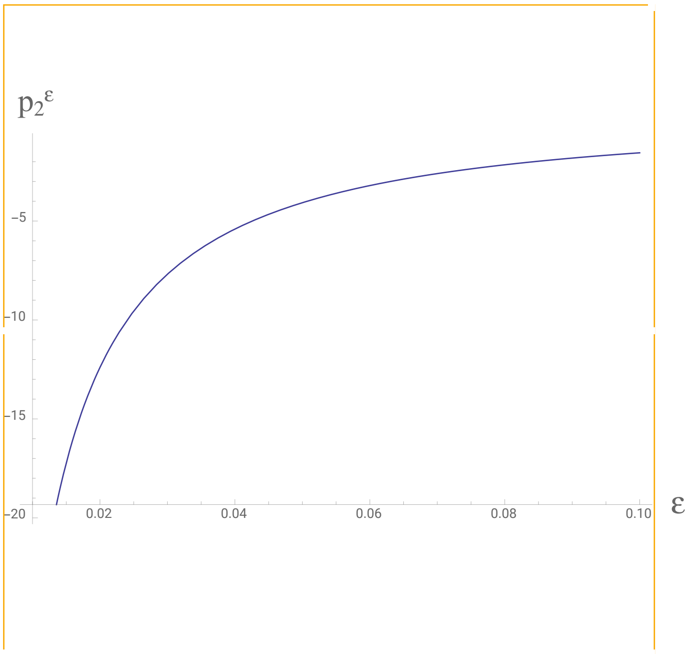

Figure 4. Equilibrium values for p_2^{ε} with a change in ε

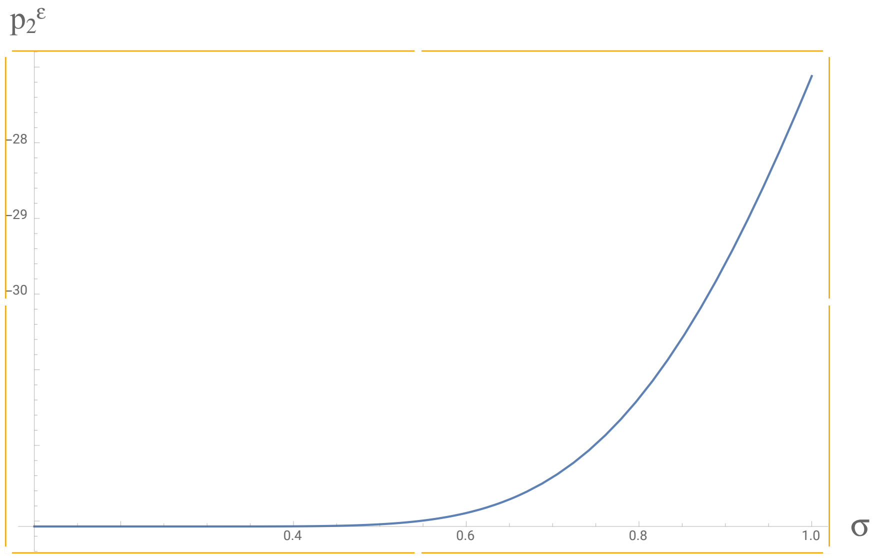

Figure 5. Equilibrium values for p_2^ε with a change in σ