UNIVERSITY OF Southampton
School of Social Sciences

Economics Division
University of Southampton
Southampton SO17 1BJ, UK

Discussion Papers in
Economics and Econometrics

Dynamic Game under Ambiguity: the Sequential Bargaining Example, and a New "Coase Conjecture"

Authors: David Besanko (Northwestern University), Jian Tong (University of Southampton) & Jason Jianjun Wu (Compass Lexecon)

No. 1606

ISSN 0966-4246

Dynamic Game under Ambiguity: the Sequential Bargaining Example, and a New "Coase Conjecture"[1]

David Besanko[2]
Northwestern University

Jian Tong[3]
University of Southampton

Jianjun Wu[4]
Compass Lexecon

December 16, 2016

[1]The authors wish to thank Charles Manski, Alessandro Pavan, Itzhak Gilboa, David Schmeidler, Larry Epstein, Larry Samuelson, Sujoy Mukerji, Peter Klibanoff, David Levine, Joel Sobel, Martin Cripps, Hamid Sabourian, John Aldrich, Spyros Galanis and Miltos Makris for stimulating discussions.

[2]Department of Management and Strategy, Kellogg School of Management, Northwestern University, Evanston, IL 60208. Email: d-besanko@kellogg.northwestern.edu

[3]Department of Economics, Faculty of Social, Human and Mathematical Sciences, University of Southampton, Southampton, SO17 1BJ, UK. Email: j.tong@soton.ac.uk

[4]Compass Lexecon, 707 State Rd, Suite 223, Princeton, NJ 08540. Email: jwu@compasslexecon.com.

Abstract

Conventional Bayesian games of incomplete information are limited in their ability to represent severe incompleteness of information. Using an illustrative example of (seller offer) sequential bargaining with one-sided incomplete information, we analyze a dynamic game under ambiguity. The novelty of our model is the stark assumption that the seller has complete ignorance—represented by the set of all plausible prior distributions—over the buyer's type. We propose a new equilibrium concept—Perfect Objectivist Equilibrium (POE)—in which multiple priors and full Bayesian updating characterize the belief system, and the uninformed player maximizes the infimum expected utility over non-weakly-dominated strategies. We provide a novel justification for refining POE through Markov perfection, and obtain a unique refined equilibrium. This results in a New "Coase Conjecture"—a competitive outcome arising from an apparent monopoly, which does not require the discount rate to approach zero, and is robust to reversion caused by reputation equilibria.

Key words: ambiguity, dynamic game with incomplete information, bargaining power, bargaining indeterminacy, Markov perfection, deflationary expectation, Coase Conjecture

JEL classification: D82, D83, C7, C78

"Acknowledging what is known as known, what is not known as unknown, that is knowledge."

—Analects of Confucius

"What has now appeared is that the mathematical concept of probability is inadequate to express our mental confidence or diffidence in making such inferences, and that the mathematical quantity which appears to be appropriate for measuring our order of preference among different possible populations does not in fact obey the laws of probability. To distinguish it from probability, I have used the term 'likelihood' to designate this quantity."

—Sir R. A. Fisher (1925)

1 Introduction

Numerous economic models, especially dynamic Bayesian games, routinely start with the presumption that a decision maker is endowed with a unique prior probabilistic belief about an unknown state of the world. A crucial implication of this presumption is to permit the notions of optimal risk taking and expected utility maximization to be well defined and employed. From an epistemic point of view, since inferences must start from somewhere, a unique prior seems a reasonable starting point, as is often argued by defenders of the Savage-Bayesian paradigm.

At a fundamental level, this line of argument has been increasingly questionable in light of the recent literature on multiple priors and ambiguity (Gilboa, Maccheroni, Marinacci and Schmeidler (GMMS), 2010; Manski, 2008)—according to which a multiple-prior belief system should be the rule and the unique-prior belief system is just a special case. Epstein and Schneider (2007) further illustrate that in a multiple-priors setting, some form of likelihood inference can suitably extend or replace the familiar Bayesian inference based on a single prior. Once the generic presumption of a single prior is removed, the notion of optimal risk taking can become ill-defined, which in turn invalidates behavior predictions built upon this notion.

Using a (seller offer) sequential bargaining game with one-sided incomplete information[1] and common interest,[2] we demonstrate how epistemic subjectivity can be avoided in modelling players' prior beliefs in a dynamic game of incomplete information, and predict its consequence. By allowing multiple priors, we show that arbitrary restrictions on the set of priors, such as uniqueness, are unnecessary. To fix ideas, consider a seller who has an exclusive right to propose the price of a good to a buyer, whose valuation of the good is his private information and thus his type. Suppose the seller has complete ignorance about the buyer's type. This can happen either because the seller is venturing into a totally unfamiliar territory or field, or because her intellectual and epistemic confidence has been shaken after major financial or economic crises. She thus justifiably doubts the relevance and applicability of any established statistical regularity to the characterization of the buyer's type. Nevertheless, the conventional Bayesian will have to force the

[1] Bargaining power relation between transacting parties lies at the heart of social and economic interactions. Among its various determinants, the effect of incompleteness of information may be of great importance. For example, if one transacting party has an apparent monopoly position but lacks commitment to her proposed terms of transaction (like a durable good monopolist, or a seller with incomplete information about the buyer's valuation), the willingness of the monopolist to make future concession may create competition with herself, and in effect lose her monopoly status. This is (the essence of) the famous Coase Conjecture (see Coase, 1972; Gul, Sonnenschein and Wilson, 1986). Therefore bargaining games are an extremely interesting context for studying the implications of ambiguous (multiple-prior) belief system on bargaining power relation.

[2] The common interest component in the bargaining relation entails a minimal level of congruence between the parties' interests.

assumption that the seller forms a unique prior probability measure, e.g., a uniform distribution, over the buyer's type. This is epistemically flawed since any probability measure, including the uniform distribution, is informative and therefore cannot adequately represent complete ignorance.[3] This restrictive assumption understates the degree of incompleteness of information. In contrast, the assumption of ignorance can be adequately represented by the set of all plausible priors within the multiple-prior approach, ruling nothing out.

When ambiguity (i.e., large set of priors) is modelled to represent a high degree of incompleteness of the information, a decision marker's optimal choice may become ambiguous in the sense that the set of optimum is too large (e.g., a continuum) to render the term "optimal risk taking behavior" meaningful[4]. In such cases, the uninformed player's behavior becomes unpredictable on the basis of individual rationality alone, and hence may result in multiplicity of equilibria in strategic games. This possibility is particularly relevant in the context of bargaining—since if an uninformed player who has the right to make the offer is ambiguous about what her preferred term of transaction is, many arbitrary bargaining strategies can be rationalized—which can cause bargaining indeterminacy. This novel cause of potential bargaining indeterminacy and its resolution are an interesting application—and motivation—of our theory of dynamic games under ambiguity.

Our illustrative example is based on seller offer sequential bargaining with one-sided incomplete information. In contrast to the conventional Bayesian model, the main novelty of our model is the stark assumption that the seller has complete ignorance—represented by the set of all plausible prior distributions—over the buyer's type. For the sake of connecting with the "Coase Conjecture"[5], we choose the standard assumption that the lower bound of total surplus from trade is negative ("no gap" as known in the literature). We also add another novel feature to the model—the assumption that if trade occurs, besides the price of the good, the seller also benefits from some positive externality[6] which is proportional to the valuation of the good by the buyer. Even at the limit—as the magnitude of the positive externality becomes infinitesimal—this qualitatively important feature can make the seller a bit keener to trade with the seller.

As a standard assumption, the seller is an expected utility maximizer if she is endowed with a unique probability measure about the state of the world. Her incompleteness of information about the true state of the world means that her expected utility is not well-defined in general so she may not form complete ordering over alternative strategies. How this potential indecision (incompleteness of preferences) is resolved depends critically on the seller's attitude toward ambiguity. Following GMMS (2010) we consider the uninformed player's objective is to maximize her infimum expected utility over non-weakly-dominated strategies. This decision rule has two criteria: The first is the max-inf expected utility criterion, which is a slight variant

[3]It is now well known (see Edwards 1992, pp. 57-61) that a single probability distribution cannot accurately represent complete ignorance.

[4]As a specific application of this point, the notion of "optimal second degree price discrimination" also becomes (practically) meaningless if the monopolist seller knows that she does not know the distribution of the buyers' type. Similarly, even the familiar notion of "monopoly pricing" can become (practically) meaningless if the monopolist seller knows that she has no clue about the actual demand function.

[5]The original Coase Conjecture (Coase 1972) envisages a surprising phenomenon of competitive outcome arising from durable good monopoly. Later literature on conventional Bayesian game of (seller offer) sequential bargaining with one-sided incomplete information identifies an equivalent outcome, which we also informally refer to as the "Coase Conjecture". Here we use the scare quotes to indicate this informal references to both the original durable good monopoly and the later sequential bargaining versions.

[6]For example, this can be (user generated) knowledge about the commercial value of the product or technology the seller (firm) is developing, from which the firm can benefit in the future. It seems reasonable to assume that this positive externality is higher if the buyer has a higher valuation of the product or technology. The magnitude of this externality has no significance for our main result which still holds as the magnitude approaches zero.

of the familiar max-min expected utility criterion.[7] The second is the criterion that no weakly dominated strategy should be chosen if there exist multiple max-inf solutions.[8] Although these decision criteria help guarantee the completeness of preferences, in the bargaining context, they may prove too weak to help resolve a potential bargaining indeterminacy. This is where the refinement of equilibria by Markov perfection fits in well and seems to have a particular justification.

Our paper makes five main contributions. First, we formulate and analyze a dynamic game of incomplete information with a multiple-prior belief system—as opposed to the conventional single-prior Bayesian game, *a la* Harsanyi (1967, 1968a,b). Second, we complement the emerging literature on learning under ambiguity (Epstein and Schneider, 2003, 2007) by incorporating strategic interactions between rational players into learning under ambiguity. Third, we present a new equilibrium concept—the Perfect Objectivist Equilibrium (POE)—which extends and contrasts the familiar Perfect Bayesian equilibrium (PBE)—that is the standard in dynamic single-prior Bayesian games. The term 'objectivist' emphasizes the importance of epistemic objectivity in the learning process. The new equilibrium concept also allows us to deal with out-of-equilibrium updating of beliefs in a systematic way. Fourth, we show that the epistemic under-determination of probabilistic belief adds a new cause for potential bargaining indeterminacy—i.e., multiplicity of POE—and provide a novel justification for its resolution through equilibrium refinement based on Markov perfection. We fully characterize—and establish the existence and uniqueness of—a Markov Perfect Objectivist Equilibrium (MPOE) for our sequential bargaining model. Fifth, we contrast the Markov Perfect Bayesian Equilibria (MPBE) and the Perfect Bayesian Equilibria (PBE) of the conventional Bayesian version of the model with the MPOE of our model. Comparing the equilibria reveals how severe incompleteness of information impacts bargaining power, and allows us to offer new insights into the renowned "Coase Conjecture".

Our model of a multiple-prior belief system is built upon the axiomatic foundation laid by GMMS (2010)—which, in turn, is a synthesis of the pioneering works of Bewley (1986, 2002) and Gilboa and Schmeidler (1989). Crucially, GMMS (2010) interpret the former strand as representing objectively rational preferences, and the latter, subjectively rational preferences. The contrast between the two is: the former does not involve personal attitude toward ambiguity, but cannot assure the completeness of the preferences relation; while the latter does rely on personal taste (over ambiguity) to achieve the completeness of the preferences relation. In comparison with the axiomatic foundation of the subjective probability theory—notably, the popular version by Anscombe and Aumann 1963)—GMMS (2010) propose less restrictive sets of axioms for a pair of rational preference relations (representing objective rationality and subjective rationality respectively).[9] Based on these more reasonable sets of rationality axioms, GMMS (2010) prove that rational beliefs of a decision maker—as revealed by hypothetical betting behavior—are represented by a set of multiple prior probability measures (hence ambiguity)—as opposed to the Savage-Bayesian unique-prior belief system.

[7] It could be argued that the max-min (max-inf) criterion for decision making under ambiguity may be too cautious or pessimistic. Now there is a rich variety of models of decision making under ambiguity that consider weaker representations of aversion to ambiguity. (See Gilboa and Marinacci (2011) for a survey of axiomatic decision theoretic models.) While we appreciate that the modelling choice of the max-min (max-inf) criterion cannot claim to be compelling and unrestrictive, we suggest that this weakness may not be so critical in the context of games under ambiguity because—as will transpire through the analysis of the bargaining under ambiguity—a key problem to be tackled is the multiplicity of equilibria. In applications, the (potential) restrictiveness of the max-min (max-inf) criterion *per se* can turn out to be not restrictive enough—because the "indifference curve" remains too "fat"—to help resolve the problem of multiplicity of equilibria.

[8] This criterion was advocated by Manski (2008). We suggest that GMMS (2010)—when suitably extended—could provide an axiomatic foundation for this condition.

[9] The subjective (or personal) probability theory was pioneered by Ramsey (1926), de Finetti (1937) and Savage (1954).

Since the multiple prior belief system makes the subjective expected utility theory unworkable, GMMS (2010) also provide an axiomatic foundation for the minimum expected utility theory—which is a substitute for the subjective expected utility theory. Our model of decision under ambiguity is more explicit than GMMS (2010) in separating beliefs from tastes over ambiguity—we require that the rational beliefs should be free from epistemic subjectivity even at the cost of increased ambiguity. In explicitly adopting lexicographic preferences, we require the objectively rational preferences to be primary and the subjectively rational to be secondary. The former is based on the weak dominance partial order—therefore action plan A is strictly preferred to plan B if and only if A weakly dominates B. If A and B cannot be primarily ranked according to the former, then they will be secondarily ranked according to the latter—which is complete and represented by the infimum expected utility function.

Our model of learning under ambiguity is closely related to the seminal works of Epstein and Schneider (2003, 2007), which study belief updating and likelihood inference under multiple priors. Our paper adopts the likelihood inference as a generic extension or replacement of Bayesian inference to avoid epistemic subjectivity in beliefs. Starting with a set of objectively plausible priors, the learning process is modelled as a prior-by-prior Bayesian updating—on the basis of a likelihood function—to derive a set of objectively plausible posteriors. In a game-theoretic context, we assign a central role to the deduced equilibrium path in determining the likelihood function and belief updating. The multiple-prior likelihood inference we adopt can be seen as a synthesis of ideas from the three competing philosophies of statistics: Bayesian (Bayes' Theorem, prior-by-prior Bayesian updating), frequentist (hypothesis testing theory) and Fisherian (likelihood).[10] A central theme of this synthesis is the realization that a single prior (or posterior) probability measure is not a sufficient statistic to summarize existent information[11]—while multiple prior distributions plus likelihoods more adequately represent knowledge and ignorance.[12]

Our paper also contributes to the literature of sequential bargaining (see Rubinstein 1982, 1987 for origin and survey), particularly, sequential bargaining under incomplete information.[13] (For surveys, see Fudenberg and Tirole 1991, Chapter 10; Kennan and Wilson 1993; Ausubel, Cramton and Deneckere 2000.) Rubinstein (1982) formulated and characterized an alternating offers bilateral bargaining game with complete information. An important insight that emerged from the Rubinstein dynamic bargaining theory is that—taking bargaining procedure as given—the player's time preferences (i.e., discount rates) are a key determinant of bargaining power and outcome: the higher the discount rate, the more costly to reject the opponent's offer, the weaker one's bargaining power.[14] It is an amazing discovery—in sequential bargaining games with complete information, positive discount rates help resolving bargaining indeterminacy which is prevalent in static bargaining games. Models of sequential bargaining under incomplete information allow the surplus from trade to be private information of informed players, therefore are useful to describe many realistic bargaining situations. Unfortunately, the thorny problem of bargaining indeterminacy reappears in many of

[10]According to Efron (1998): "the development of modern statistical theory has been a three-sided tug of war between the Bayesian, frequentist and Fisherian viewpoints." "In many ways the Bayesian and frequentist philosophies stand at opposite poles from each other, with Fisher's ideas being somewhat of a compromise." "The world of applied statistics seems to need an effective compromise between Bayesian and frequentist ideas."

[11]This statement is inspired by the quote from Fisher (1925) presented before the introduction of the paper.

[12]To adequately represent knowledge, it is important to indicate both what is known and what is (known to be) unknown (see the quote from Analects of Confucius presented at the beginning of the paper).

[13]For original contributions to the large Bayesian literature of this bargaining game, see Sobel and Takahashi (1983), Fudenberg, Levine and Tirole (1985), Gul, Sonnenschein and Wilson (1986), Ausubel and Deneckere (1989a,b).

[14]See Sutton (1986) for a concise exposition and proof.

these models—usually through the "Folk Theorem" type of results for sufficiently small but positive discount rates. In a survey, Rubinstein (1987) made an insightful remark on the "state of the art" of the literature:

> In my opinion, we are far from having a definitive theory of bargaining with incomplete information for use in economic theory. The problems go deeper than bargaining theory and appear in the literature of refinement of S.E. [i.e., Sequential Equilibrium or Perfect Bayesian Equilibrium[15] for that matter], an issue explored thoroughly in the last few years. My intuition is that something is basically wrong in our approach to games with incomplete information and that the 'state of the art' of bargaining reflects our more general confusion.

The difficult problem of refining perfect Bayesian equilibrium remains unresolved. The issue centers on whether the famous "Coase Conjecture"—the outcome in which an apparent monopolist seller loses her bargaining power entirely in the face of incomplete information and diminishing discount rate—should obtain, or whether it can be reversed by a continuum of reputation equilibria that help alleviate the commitment problem that causes the "Coase Conjecture". Revealingly, the prominent equilibrium refinement criterion—Markov perfection, which eliminates all reputation equilibria in the context of sequential bargaining (on the ground of lacking simplicity)—seems obviously objectionable. When the incompleteness of information is recognized as much more severe than the conventional Bayesian model can represent—e.g., reaching complete ignorance—the nature of potential bargaining indeterminacy changes. Our contribution to this literature is the finding that the epistemic under-determination of probabilistic belief adds a previously unrecognized and more generic cause of (potential) bargaining indeterminacy. Furthermore, this seeming curse turns out to a blessing in disguise—we find a novel justification for using the criterion of Markov perfect to refine equilibrium, which can suitably resolve this type of bargaining indeterminacy. As a result, we obtain a unique solution to the bargaining problem. This solution results in competitive outcome similar to the "Coase Conjecture", but it does not require the discount rate to approach zero. Also, it is robust to reversion that may be caused by a reputation equilibrium.

The paper is organized in eight sections. Section 2 describes the illustrative model of sequential bargaining that forms the basis of our analysis. Section 3 explains formally what we mean by ambiguity, and how the uninformed player makes likelihood inferences and updates beliefs; and describes the criteria for decision making under ambiguity. Section 4 defines the Perfect Objectivist Equilibrium and establishes the foundations—including refinement—for characterizing it in our illustrative model. Sections 5 proceeds through two steps to characterize Markov Perfect Objectivist Equilibrium. To build the preliminaries for the main analysis, we first characterize a benchmark model with complete information. We then establish the existence and uniqueness of a Markov perfect objectivist equilibrium. In Section 6, we illustrate the (potential) problem of bargaining indeterminacy by establishing the multiplicity of POE, and provide a (novel) justification for using MPOE as an equilibrium refinement. In Section 7 we compare the insights of the current model with its counterpart: perfect Bayesian equilibrium. Particularly, we highlight the main results in relation to the famous "Coase Conjecture", and introduce a New "Coase Conjecture" which is based on the MPOE analysis. Section 8 begins with a summary, we then make some concluding remarks; particularly, we draw a striking implication about the plausibility of *secular stagnation* in a post-crisis era—it is based on the

[15] In this paper, our discussion is focused on a closely related solution concept—Perfect Bayesian Equilibrium—to which Rubinstein's comments also apply.

analysis of ambiguity and the New "Coase Conjecture"—their impact on bargaining power, monopoly rent (Schumpeterian entrepreneurial profit), innovation and growth. Proofs of all results are in the Appendix.

2 Sequential Bargaining with One-Sided Incomplete Information

Consider a sequential bargaining game in which a seller has the exclusive right to make an offer to produce and sell a single unit of a good at a price w_t to a buyer, whose valuation of the good, $\theta\Pi_B$, is his private information and thus his type is denoted by $\theta \in [0,1]$.[16] If the offer price w_t is accepted, trade occurs and the game ends; otherwise, the seller can revise (or maintain) the offer in the next period. Trading with a buyer who values the good, or the underlying technology, yields to the seller a positive externality of magnitude $\theta\Pi_S$. Such positive externality can be in forms of technology spillover or surplus unappropriated by the buyer. The production cost is $a > 0$ for the seller. To make an interesting economic problem, we maintain the following assumption which says for the highest possible type buyer, the trade is socially worthwhile.

Assumption 1 *$\Pi_B > 0$, $\Pi_S > 0$ and $\Pi_B + \Pi_S > a$.*

The buyer has the following static utility function for period $t = 0, 1, \ldots$:

$$v_t = (\theta\Pi_B - w_t)\, k_t, \tag{1}$$

and the seller's static utility is

$$u_t = (\theta\Pi_S + w_t - a)\, k_t, \tag{2}$$

where $w_t \in \mathcal{W} := [-\Pi_S, \Pi_B]$ is the price offered by the seller in period t; $k_t \in \mathcal{K} := \{0,1\}$ is buyer's decision in period t, with $k_t = 1$ denoting acceptance and $k_t = 0$, rejection. Both w_t and k_t are commonly observable. The game continues from period t to period $t+1$ if and only if the indicator of trade $K_t := \sum_{\tau=0}^{t} k_t = 0$.

Denote by $\mathcal{H}^t = \mathcal{H}_S^t \times \mathcal{H}_B^t$ the set of all possible histories from period 0 till t as long as the game has not ended by t, where

$$\mathcal{H}_S^t := \left\{ \{(w_\tau, k_\tau)\}_{\tau=0}^{t-1} | w_\tau \in \mathcal{W}, k_\tau \in \mathcal{K} \right\},$$

$$\mathcal{H}_B^t = \left\{ \left(\theta, \{(w_\tau, k_\tau)\}_{\tau=0}^{t-1}, w_t\right) | w_\tau, w_t \in \mathcal{W}, k_\tau \in \mathcal{K} \right\},$$

$t = 0, 1, \ldots$.[17] The realized history is given by $h^t = (h_S^t, h_B^t) \in \mathcal{H}^t$, where $h_S^t = \{(w_\tau, k_\tau)\}_{\tau=0}^{t-1} \in \mathcal{H}_S^t$ and $h_B^t = \left(\theta, \{(w_\tau, k_\tau)\}_{\tau=0}^{t-1}, w_t\right) \in \mathcal{H}_B^t$ are the information accessible to the seller and buyer respectively.

A pure strategy for the seller is the function $s_S : \mathcal{H}_S^t \to \mathcal{W}$ for $t = 0, 1, \ldots$ (such that the game has not ended by t); while $s_B : \mathcal{H}_B^t \to \mathcal{K}$ for $t = 0, 1, \ldots$ is the counterpart for the buyer. A pure strategy profile $s = (s_S, s_B)$ is a map

$$s : \mathcal{H}^t \to \mathcal{W} \times \mathcal{K}, \text{ for } t = 0, 1, \ldots.$$

For brevity and without loss of clarity, we write $s_S(h^t)$ (and respectively $s_B(h^t)$) instead of $s_S(h_S^t)$ (and respectively $s_B(h_B^t)$). Notice that $s_S(\cdot)$ (and respectively $s_B(\cdot)$) is only measurable in $\mathcal{H}_S^t$ (and respectively

[16] This formulation uses two parameters θ and Π_B. Since Π_B is a parameter which denotes the upper bound of the valuation, it is without loss of generality that θ has the range of $[0,1]$.

[17] Throughout the paper, we use subscript t to index players' actions taken in given period t, while we use superscript t to index histories of actions up to (excluding) the players' respective actions in period t.

$\mathcal{H}_B^t$). Thus, the seller's move $s_S(h^t) \in \mathcal{W}$ does not depend on parameter θ, which is not contained in h_S^t. In contrast, the buyer's strategy is a function of θ. Note that the seller moves first at the beginning of each period to offer w_t, which is observed by the buyer in his turn to move. So w_t is contained in h_B^t while k_t is not contained in h_S^t, an asymmetry reflecting the sequential relation between w_t and k_t. The seller can commit to w_t in period t, but the buyer cannot commit to a specific response prior to observing w_t but is motivated to play the best response. This asymmetry gives the seller a first-mover advantage. The buyer's potential advantage is his hidden information about θ, which is contained in h_B^t but not in h_S^t.

Let $\mathcal{S}_S := \{s_S : \mathcal{H}_S^t \to \mathcal{W} | t = 0, 1, \ldots\}$ and $\mathcal{S}_B := \{s_B : \mathcal{H}_B^t \to \mathcal{K} | t = 0, 1, \ldots\}$ be the sets of all possible pure strategies of the seller and the buyer respectively, and let $\mathcal{S} := \mathcal{S}_S \times \mathcal{S}_B$ denote the set of all pure strategy profiles.

The buyer's type space is $\Theta := [0, 1]$. The seller's belief about θ is inferred from observational data on history h_S^t according to an endogenous set of inference rules. The seller's set of inference rules within an equilibrium of the game involves a description of that equilibrium. Therefore there exist multiple sets of inference rules across multiplicity of equilibrium outcomes; but within each pure strategy equilibrium, there is a unique set of inference rules consistent with that pure strategy equilibrium. The posterior set of inferred types of buyer may (partially) preserve the ambiguity about θ originated from the multiple priors. That is to say, the inference from observational data on history h_S^t may allow only partial identification of $\theta \in \Theta$.

3 Ambiguity in a Dynamic Game

With multiple priors, the sequential bargaining game embodies ambiguity. We start by laying the ground work for modeling ambiguity in two steps. First, we develop the analytical approach for how a seller with multiple priors would learn over time. We do so by formulating the seller's multiple prior probabilistic beliefs over the buyer's private type $\theta \in \Theta$ and then explaining the likelihood inference process that the seller uses to update her multiple probabilistic beliefs. Second, we specify the seller's objective in light of the ambiguity she faces, and indicate what it means for the seller to follow an optimal strategy.

3.1 Interactive Learning under Ambiguity

3.1.1 The Seller's Multiple Priors Over θ

Each of the seller's multiple priors over $\theta \in \Theta$ is a (countably additive) probability measure, a set function $\mu_0 : \Sigma(\Theta) \to [0, 1]$, where $\Sigma(\Theta)$ is the Borel σ-algebra over the buyer type space Θ. To make the illustration stark, we make the extreme assumption that the seller is completely ignorant about what the true prior is, therefore entertains the set of all priors, denoted by $\mathcal{M}_0 := \triangle(\Sigma(\Theta))$.[18] We also make the standard assumption that the seller is an expected utility maximizer when she knows the true unique probability measure. For any random variable or real-valued function $g : (\Theta, \Sigma(\Theta), \mu) \to \mathbb{R}$ such that g is Borel-measurable and bounded, the (mathematical) expectation of g under $\mu \in \triangle(\Sigma(\Theta))$ is well defined as the Lebesgue integral:

$$E_\mu[g] := \int_\Theta g d\mu. \tag{3}$$

[18] The measure-theoretic formulation allows great generality, so that each $\mu \in \mathcal{M}_0$ can be either a continuous or a discrete distribution, or a convex combination of both.

When μ is a Dirac measure (i.e., an indicator function), for any $\mathcal{A} \in \Sigma(\Theta)$ we have

$$\mu(\mathcal{A}) = \delta_\theta(\mathcal{A}) := 1_{\mathcal{A}}(\theta) := \begin{cases} 1 & \text{if } \theta \in \mathcal{A}, \\ 0 & \text{otherwise,} \end{cases} \tag{4}$$

and

$$E_\mu[g] = E_{\delta_\theta}[g] = g(\theta). \tag{5}$$

Since the state space $\Theta = [0, 1]$ is a set of real numbers, in general we can define a μ-induced cumulative distribution function $F_\mu(x) := \mu((-\infty, x])$ for all $x \in \Theta$.[19] Thus, $F_\mu(x)$ is everywhere well-defined. It is non-decreasing, right-continuous, $F_\mu(0_-) = 0$ and $F_\mu(1) = 1$. So $E_\mu[g]$ can be rewritten as a Lebesgue-Stieltjes integral:

$$E_\mu[g] = \int_0^1 g dF_\mu. \tag{6}$$

Notice that $\int_0^1 g dF_\mu$ agrees with the Riemann-Stieltjes integral whenever the latter is defined. Particularly, for applications where $F_\mu(x)$ is continuously differentiable and therefore there exists a continuous probability density function $f_\mu(x) := \frac{dF_\mu(x)}{dx}$, if $g(x)$ is continuous then $E_\mu[g]$ equals the well-defined Riemann integral:

$$E_\mu[g] = \int_0^1 [g(x) f_\mu(x)] dx. \tag{7}$$

3.1.2 Equilibrium and Likelihood Inference

The seller's inference about the (unobservable) type of the buyer from observation has two phases. First, an inference from each given type of the buyer to a (deduced) history of the game on the premise of a putative equilibrium path determines a likelihood function. The second phase is an inference running from an observed history to a posterior characterization of types of the buyer by prior-by-prior Bayesian updating on the basis of the likelihood function. The likelihood function, denoted by l, which plays the central role in the seller's learning process, is determined by the common knowledge among the players about the structure of the game. Define the function $l : \mathcal{H}_S^t \times \mathcal{S} \times \Theta \rightarrow [0, 1]$ such that $l(\cdot; s, \theta)$ is a conditional probability function over $\mathcal{H}_S^t$. The argument s is a putative pure strategy equilibrium in question. For example, $l(h_S^t; s, \theta)$ is the probability of history h_S^t conditional on (the assumption) (s, θ), that is, the data h_S^t is generated by the equilibrium s with buyer type θ.[20] Notice that $l(h_S^t; s, \theta)$ is the likelihood of the parameter value of θ; and, importantly, $l(h_S^t; s, \cdot)$ is neither a probability measure over Θ, nor dependent of an prior probability measure $\mu \in \mathcal{M}_0$.

In this paper, we confine our theoretical construct and analysis to the special case of pure strategy equilibrium s, on the premise of which the seller can deduce a unique history from the hypothesis θ. We therefore have $l(h_S^t; s, \theta) = 0$ for all $h_S^t \in \mathcal{H}_S^t$ except for one unique history. Conversely, given s, the seller can (partially) identify from the observed history h_S^t the types of the buyer that are consistent with the

[19] Observe that μ is also the unique Lebesgue-Stieltjes measure induced by the cumulative distribution function $F_\mu(x)$. Particularly, $\mu((a, b]) = F_\mu(b) - F_\mu(a)$ for $0 \leq a < b \leq 1$ and $\mu(\{b\}) = F_\mu(b) - F_\mu(b_-)$. (See Royden, 1988, chapter 12.3 for a textbook treatment.)

[20] Thus, $l\left(h_S^t; s, \theta\right)$ does not obey the laws of probability when h_S^t is fixed and θ is treated as a variable, e.g., the integration over θ does not add up to one. That is why we use the term "likelihood" as opposed to "probability." Likelihood measures the evidential support for parameter value θ from data h_S^t conditional on the assumption of s.

history.

Within each putative pure strategy equilibrium s, the likelihood inference is conditioned on the assumption that the observed history is on the equilibrium path of s. This is captured by the fact that the conditional likelihood function $l\left(h_S^t; s, \cdot\right): \Theta \to [0,1]$ depends on argument s.

The deduction of the history from the assumption of (s,θ) can be represented by the map $\Gamma: \mathcal{S}\times\Theta \to \mathcal{H}_S^{\infty}$ such that for $\forall (s,\theta) \in \mathcal{S}\times\Theta$, $\Gamma(s,\theta) = h_S^{\infty} \in \mathcal{H}_S^{\infty}$ predicts an entire (seller-observable part of) equilibrium path of the game (conditional on the game having not ended). Let $\Gamma^t(s,\theta)$ be the restriction of $\Gamma(s,\theta)$ to the set (of periods) $\{0,1,\ldots,t\}$. Therefore, $\Gamma^t(s,\theta)$ predicts the (seller-observable part of) on-equilibrium-path history $h_S^t \in \mathcal{H}^t$. This part of reasoning is hypothetico-deductive inference, and it determines the unique likelihood function l such that

$$l\left(h_S^t; s, \theta\right) = 1_{\Gamma^t(s,\theta)}\left(h_S^t\right) := \begin{cases} 1 & \text{if } h_S^t = \Gamma^t(s,\theta) \\ 0 & \text{otherwise} \end{cases}, \tag{8}$$

where $1_{\Gamma^t(s,\theta)}\left(h_S^t\right)$ is a indicator function.[21]

Conversely, from the seller-observed history of the game h_S^t, the seller can (partially) identify the plausible types of the buyer. Define the (partial) identification correspondence, $\mathcal{I}: \mathcal{H}_S^t \times \mathcal{S} \to \Sigma(\Theta)$, such that for any $h_S^t \in \mathcal{H}_S^t$ and $s \in \mathcal{S}$,

$$\mathcal{I}\left(h_S^t; s\right) := \left\{\theta \in \Theta | \Gamma^t(s,\theta) = h_S^t\right\}. \tag{9}$$

The set $\mathcal{I}\left(h_S^t; s\right)$ is the updated (plausible) type space of the buyer.[22] We confine our attention to well-behaved pure strategy equilibrium s such that $\mathcal{I}\left(h_S^t; s\right) \in \Sigma(\Theta)$.[23] Since $\mathcal{I}\left(h_S^t; s\right)$ is rarely a singleton, the identification is usually partial.[24] In typical applications, the set $\mathcal{I}\left(h_S^t; s\right)$ is an interval, in which case the infimum and supremum of $\mathcal{I}\left(h_S^t; s\right)$, denoted by $\inf\left(\mathcal{I}\left(h_S^t; s\right)\right)$ and $\sup\left(\mathcal{I}\left(h_S^t; s\right)\right)$ are the boundary points of the (updated) buyer type space.

Conditional on the assumption that the pure strategy equilibrium s is the underlying data generating mechanism, we can treat $\mathcal{I}\left(h_S^t; s\right) \in \Sigma(\Theta)$ as the conditioning event upon which the seller update her belief.

3.1.3 Belief Updating

Since the seller is the only uninformed player in the game who has to make inferences about the opponent's type, the (inference-related) likelihood function $l(\cdot; s, \theta)$ and identification correspondence $\mathcal{I}(\cdot; s)$ are only measurable in $\mathcal{H}_S^t$ instead of $\mathcal{H}^t$. Therefore without causing confusion and for the sake of brevity, for the rest of the paper we will drop subscript $_S$ and write h_S^t and $\mathcal{H}_S^t$ simply as h^t and $\mathcal{H}^t$.

If the seller entertains a unique prior μ_0, belief updating will simply follow Bayes' Rule, using history h^t

[21] This feature of the likelihood being degenerated to an indicator function is caused by the confinement of the analysis to pure strategy equilibrium.

[22] Thus, a pure strategy equilibrium defines an endogenous filtration (sequential partition of the state space), which determines the information structure of the interactive learning under ambiguity. This is in contrast with the non-interactive learning under ambiguity modelled by Epstein and Schneider (2003), where the filtration is exogenously given.

[23] Practically this (technical) assumption is always satisfied. Intuitively, this assumption requires that the (pure strategy) equilibrium must not be so complicated (or contrived) that it induces pathological (partial) identification (by the standard of Borel measurability).

[24] The term partial identification is borrowed from Manski (1995). In the current context, the pure strategy profile-based inferential problem is abstracted from the statistical inference problem, and therefore is purely a problem of identification *a la* Manski (1995).

as the conditioning event. Denote by $\mu_t(\cdot|h^t;\mu_0,s) \in \triangle(\Sigma(\Theta))$ the posterior probability measure. From Bayes' Rule, we have

$$\mu_t\left(\mathcal{A}|h^t;\mu_0,l\right) = \frac{\int_{\mathcal{A}} l\left(h^t;s,\theta\right)d\mu_0}{\int_0^1 l\left(h^t;s,\theta\right)d\mu_0}, \text{ for } \forall \mathcal{A} \in \Sigma(\Theta), \tag{10}$$

where the numerator (of the right-hand side of the equation) is the probability of the joint event that the observed history is h^t and the buyer type is $\theta \in \mathcal{A}$, and the denominator is the probability of the event that the observed history is h^t. Given the special restriction that s is a pure strategy equilibrium, the information contained in the history is sufficiently summarized by the updated buyer type space $\mathcal{I}(h^t;s)$, which, as an equivalent conditioning event to the history h^t, can be used for Bayesian updating as follows:

$$\mu_t\left(\mathcal{A}|h^t;\mu_0,s\right) = \frac{\mu_0\left(\mathcal{A}\cap\mathcal{I}\left(h^t;s\right)\right)}{\mu_0\left(\mathcal{I}\left(h^t;s\right)\right)}. \tag{11}$$

Note, the posterior is well defined only if the denominator $\mu_0(\mathcal{I}(h^t;s))$ is positive.

As a key difference from a conventional Bayesian dynamic game, the seller is not epistemically committed to any particular single prior as she entertains multiple priors. Consequently she also entertains multiple posteriors—the Bayesian updates based on all her priors. Let $\mathcal{M}_t(h^t;s)$ denote the set of all her posteriors, which is defined by

$$\mathcal{M}_t\left(h^t;s\right) := \left\{\mu_t\left(\cdot|h^t;\mu_0,s\right)|\mu_0 \in \mathcal{M}_0\right\}, \tag{12}$$

where $\mu_t(\cdot|h^t;\mu_0,s)$ is determined by (10). This is called prior-by-prior Bayesian updating, or Full Bayesian Updating (FBU).[25].

Since we assume that the seller is completely ignorant about the true prior probability measure, her set of priors $\mathcal{M}_0$ thus does not exclude any possible prior, meaning it is thus the largest set possible. In this case, the belief updating turns out to be extremely simple, as is shown by the next proposition, the seller has complete ignorance about the true probability measure over the updated buyer type space $\mathcal{I}(h^t;s) \in \Sigma(\Theta)$.

Proposition 1 *Suppose s is the strategy profile of a pure-strategy equilibrium such that the updated buyer type space is given by $\mathcal{I}(h^t;s) \in \Sigma(\Theta)$ for all history $h^t \in \mathcal{H}^t$, all $t = 0,1,\ldots$. For any history h^t such that $\mathcal{I}(h^t;s) \neq \varnothing$, the set of posteriors is given by*

$$\mathcal{M}_t\left(h^t;s\right) = \left\{\mu_0 \in \mathcal{M}_0|\mu_0\left(\mathcal{I}\left(h^t;s\right)\right) = 1\right\}, \tag{13}$$

i.e., the set of posteriors is the set of all possible probability measures whose supports are included in the updated type space $\mathcal{I}(h^t;s)$.

So far we have only dealt with histories h^t such that $\mathcal{I}(h^t;s) \neq \varnothing$, which is a necessary condition for on-equilibrium path histories. For any history h^t such that $\mathcal{I}(h^t;s) = \varnothing$ (which implies $\mu_0(\mathcal{I}(h^t;s)) = 0$ for all $\mu_0 \in \mathcal{M}_0$), the equilibrium-paths-based identification process fails and indicates that the history is off the equilibrium paths. To complete the updating of beliefs for such histories, we impose the following basic assumption.

[25] The latter term is borrowed from Gilboa and Marinacci (2011), which is a recent comprehensive survey of the ambiguity literature.

Assumption 2 *If a history* $h^t \in \mathcal{H}^t$ *is such that* $l\left(h^t; s, \theta\right) = 0$ *for all* $\theta \in \Theta$ *(i.e.,* $\mathcal{I}\left(h^t; s\right) = \varnothing \in \Sigma\left(\Theta\right)$*), then* $\mathcal{M}_t\left(h^t; s\right)$ *is given by*

$$\mathcal{M}_t\left(h^t; s\right) = \mathcal{M}_0. \tag{14}$$

This assumption captures the idea that the identification failure (i.e., $\mathcal{I}\left(h^t; s\right) = \varnothing$) indicates that the data h^t is (known to be) generated by off-equilibrium path behavior, which implies that the equilibrium-paths-based likelihood function is not valid for explaining the data. As a result, learning or belief-updating based on the invalid likelihood function should be undone, and the beliefs should revert to the initial priors. The seller maintains complete ignorance about the true probability measure over the original buyer type space Θ.

Overall, the key to belief updating is the updating of the buyer type space. In a pure strategy equilibrium, the seller always remains agnostic about the true probability measure over the (plausible) buyer type space, unless it degenerates to a singleton.

3.2 Seller's Objective and Best Response

In Bayesian games in which a player has a unique prior probability measure about the state of the world, the specification of the player's objective is straightforward: conditional on the history of play, the player forms a mathematical expectation of the payoff it seeks to maximize based on the posterior beliefs implied by that history. However, if the player does not have a unique prior, the formulation of the player's objective is more complex. Our specification of the seller's objective is inspired by Manski (2008) and GMMS (2010).

These papers deal with the situation faced by the seller in our model: what are the criteria for decision making under uncertainty when the decision maker lacks the information to quantify uncertainty using a single probability measure? Manski (2008) formulates a two-step procedure in which the first step is to eliminate all weakly-dominated actions, and the second step is to maximize the minimum expected utility function or minimize the maximum regret function (over non-weakly-dominated actions). GMMS (2010) axiomatize the problem of a decision maker who has a pair of preference relations: objectively rational preferences and subjectively rational preferences. If the decision maker chooses on the basis of objectively rational preferences, he can "solidly justify" and hence defend his choice to others because these preferences do not depend on her personal taste over ambiguity; if the decision maker chooses on the basis of subjectively rational preferences, he cannot be convinced by others that his choice was wrong because these preferences depend on her personal taste over ambiguity in a way that is regarded as subjectively rational. Objectively rational preferences generate a unanimous but incomplete ordering of actions, while subjectively rational preferences generate a complete ordering of actions that can be represented by a minimum expected utility function (with respect to all priors in the set of the decision maker's possible priors). GMMS (2010) demonstrate that given two plausible conditions (*consistency* and *caution*), there exists a common set of priors that can be used to represent both the objectively rational preferences and the subjectively rational preferences of the decision maker. This provides an epistemically intuitive interpretation for decision making based on the max-min rule. Basing decisions on a minimum expected utility function can be thought of as a way of completing an otherwise incomplete preference ordering based on objective rationality.

Our formulation is a slight variant of GMMS (2010) and applied to a dynamic game-theoretic setting. To develop this formulation, we define, derive, and characterize the seller's infimum expected utility function—

which we call the *worst-case value function*. To begin, note that the seller's static conditional utility for $\hat{s} \in \mathcal{S}$ (conditional on θ) is in general given by

$$u_\tau(\hat{s};\theta) = (\theta \Pi_S + w_\tau(\hat{s};\theta) - a)\, k_\tau(\hat{s};\theta). \tag{15}$$

The seller's conditional value function in general can thus be expressed as

$$\begin{aligned} &W_t(\hat{s};\theta) = \textstyle\sum_{\tau=t}^{\infty} u_\tau(\hat{s};\theta)\,(1+r)^{-(\tau-t)}\, 1_{K_{\tau-1}=0} \\ &= u_t(\hat{s};\theta)\, 1_{K_{t-1}=0} + (1+r)^{-1}\, W_{t+1}(\hat{s};\theta), \end{aligned} \tag{16}$$

where $1_{K_{\tau-1}=0}$ is the indicator for whether the game continues into period τ, and $r > 0$ is the common discount rate used by the seller and the buyer.

Generally, we only impose a weak restriction on the pure strategy profile $\hat{s}$ such that $W_t(\hat{s};\theta)$ must be Borel-measurable in θ. Since $W_t(\hat{s};\theta)$ is bounded, its expectation with respect to probability measure μ_0 is well defined as has been formulated in section 3.1.1. To maintain a simple and clear interpretation of the likelihood inference, we formulate the worst-case value function in terms of infimum (in stead of minimum) of expected payoff. This greater generality removes the necessity for proof of existence of the minimum.[26] We can now define the seller's worst-case value function:

Definition 1 *The seller's worst-case value function $U_t(\hat{s};h^t,s)$ is the infimum of the seller's expected value for $\hat{s} \in \mathcal{S}$, given the history h^t, the pure strategy profile $s \in \mathcal{S}$ and the set of plausible probabilistic beliefs $\mathcal{M}_t(h^t;s)$, i.e.,*

$$U_t\left(\hat{s};h^t,s\right) := \inf_{\mu \in \mathcal{M}_t(h^t;s)} \int_0^1 W_t(\hat{s};\theta)\, d\mu. \tag{17}$$

To characterize the seller's best response conditional on her set of plausible beliefs, we begin by defining what it means for one pure strategy profile to weakly dominate another from the seller's perspective:

Definition 2 *For a pair of pure strategy profiles $\hat{s}, \check{s} \in \mathcal{S}$, the profile $\hat{s}$ weakly dominates $\check{s}$ conditional on $\mathcal{M}_t(h^t;s)$—denoted by $\hat{s} \succ^*_{\mathcal{M}_t(h^t;s)} \check{s}$—iff*

$$\int_0^1 W_t(\hat{s};\theta)\, d\mu \geq \int_0^1 W_t(\check{s};\theta)\, d\mu,$$

for all $\mu \in \mathcal{M}_t(h^t;s)$ and

$$\int_0^1 W_t(\hat{s};\theta)\, d\mu > \int_0^1 W_t(\check{s};\theta)\, d\mu,$$

for some $\mu \in \mathcal{M}_t(h^t;s)$.

[26] As a simple mathematical fact, the infimum (of the expected payoff) equals the minimum if the latter exists. The standard way to ensure the existence of the minimum would be to represent the ambiguous beliefs with the closure of the convex hull of $\mathcal{M}_t\left(h^t;s\right)$ instead of $\mathcal{M}_t\left(h^t;s\right)$ itself. This transformation would complicate the epistemic interpretation of likelihood inference and belief updating. From an epistemic point of view, there is no reason why the set of probabilistic beliefs has to be compact. In fact, for example, in our sequential bargaining model, it is possible for $\mathcal{M}_t\left(h^t;s\right)$ to be the set of all (Borel measurable) probability measures whose supports are included in $[0,\theta^{**})$ for some $\theta^{**} \in (0,1)$ and $t = 1,2,\ldots$; both $\mathcal{M}_t\left(h^t;s\right)$ and its support $[0,\theta^{**})$ are not compact, since θ^{**} is a limit point of $[0,\theta^{**})$ but not included in it, and the Dirac measure $\delta_{\theta^{**}}$ is a limit point of $\mathcal{M}_t\left(h^t;s\right)$ but not included therein either.

We can now define the seller's set of pure strategy best responses:

Definition 3 *For a given $\mathcal{M}_t(h^t;s)$ and buyer strategy $s_B \in \mathcal{S}_B$, the seller's set of all best responses conditional on $\mathcal{M}_t(h^t;s)$ is given by*

$$\mathcal{S}_S^*\left(s_B;\mathcal{M}_t\left(h^t;s\right)\right) := \left\{\hat{s}_S \in \mathcal{S}_S \;\middle|\; \begin{array}{l} \hat{s}_S \in \arg\max_{s'_S \in \mathcal{S}_S} U_t\left(s'_S, s_B; h^t, s\right) \text{ and} \\ \left(\check{s}, s_B\right) \not\succ^*_{\mathcal{M}_t(h^t;s)} \left(\hat{s}, s_B\right), \forall \check{s}_S \in \mathcal{S}_S \end{array}\right\},$$

where $(\check{s}_S, s_B) \not\succ^*_{\mathcal{M}_t(h^t;s)} (\hat{s}_S, s_B)$ expresses the negation of $(\check{s}_S, s_B) \succ^*_{\mathcal{M}_t(h^t;s)} (\hat{s}_S, s_B)$.

Given the buyer's strategy $s_B \in \mathcal{S}_B$ and the seller's belief represented by $\mathcal{M}_t(h^t;s)$, the seller's objective is to play one of her best responses $\tilde{s}_S \in \mathcal{S}_S^*(s_B;\mathcal{M}_t(h^t;s))$. By definition of a best response, the worst-case value of strategy $\tilde{s}_S$ must be the maximum in comparison with all alternative strategies in $\mathcal{S}_S$. If there are multiple strategies that have the maximum worst-case value, then $\tilde{s}_S$ also must not be conditionally weakly dominated by any alternative strategy in $\mathcal{S}_S$ (conditional on $\mathcal{M}_t(h^t;s)$).[27] For this reason, the seller's preferences are lexicographical because they cannot entirely be represented by a utility function (such as the infimum expected utility function). Notice that our formulation of the seller's preferences gives a significance to weak dominance (among the objectively rational preferences). To see this, note that if strategies $\hat{s}_S, \check{s}_S \in \mathcal{S}_S$ have the same infimum expected value, but $\hat{s}_S$ conditionally weakly dominates $\check{s}_S$, then the seller must not be indifferent between $\hat{s}_S$ and $\check{s}_S$; instead, she must strictly prefer $\hat{s}_S$ to $\check{s}_S$. Consequently, the conditionally weakly dominated strategy $\check{s}_S$ cannot be optimal. This criterion might help to reduce the multiplicity of best responses if they exist.

Before turning to an analysis of equilibrium, we specify the buyer's value function. Using (1), the buyer's value function in general is given by

$$\begin{aligned} V_t(s;\theta) &= \textstyle\sum_{\tau=t}^{\infty} v_\tau(s;\theta)(1+r)^{-(\tau-t)} 1_{K_{\tau-1}=0} \\ &= v_t(s;\theta) 1_{K_{t-1}=0} + (1+r)^{-1} V_{t+1}(s;\theta). \end{aligned} \tag{18}$$

4 Solution Concept

4.1 Perfect Objectivist Equilibrium (POE)

Our solution concept is an extension of the familiar perfect Bayesian equilibrium (PBE) to a setting in which the uninformed player is allowed to entertain multiple priors. We call it the *perfect objectivist equilibrium* (POE). The term '*objectivist*' is to emphasize that the ambiguity in the probabilistic beliefs is entailed by epistemic objectivity when there is insufficient information to uniquely determine the probabilistic belief as prescribed by the PBE.

Definition 4 *A POE of the game is a tuple $(s, \mathcal{M}(s))$, where $\mathcal{M}(s)$ represents the beliefs of the seller such that*

$$\mathcal{M}(s) = \left\{\mathcal{M}_t\left(h^t;s\right) | h^t \in \mathcal{H}^t, t = 0,1,\ldots\right\},$$

[27] In the Appendix, we provide an example of a weakly dominated strategy which satisfies the max-inf expected utility criterion. However, to understand the example, one must first understand the equilibrium analysis developed in Section 5. For this reason, the example is presented immediately after the proof of Theorem 1.

and $s = (s_S, s_B)$ *is the pure strategy profile represented by a map*

$$s : \mathcal{H}^t \to \mathcal{W} \times \mathcal{K},\ \textit{for all } t = 0, 1, \ldots$$

such that,

(i) Both s_S *and* s_B *are* sequentially rational, *namely, for any period* $t = 0, 1, \ldots$ *and realized history* h^t*, given the conditional posteriors* $\mathcal{M}_t(h^t; s)$*, the continuation strategies derived from* s_B *and* s_S *are mutual best responses. That is,*

$$s_S \in \mathcal{S}_P^* \left(s_B; \mathcal{M}_t \left(h^t; s\right)\right);$$

for all $\mathcal{M}_t(h^t; s)$ *such that* $\mathcal{I}(h^t; s) \in \Sigma(\Theta)$*, and*

$$V_t(s_S, s_B, \theta) \geq V_t(s_S, s'_B, \theta)$$

for all $s'_B \in \mathcal{S}_B$ *for all* $\theta \in \Theta$*.*

(ii) The priors are given by $\mathcal{M}_0\left(h^0; s\right) = \mathcal{M}_0$*. The belief updating is through the "likelihood inference" based on* s*, which is described in Sections 3.1.2 and 3.1.3. Specifically, (a) for all* $t = 1, 2, \ldots$ *the set of all plausible posteriors* $\mathcal{M}_t(h^t; s)$ *against any history* h^t *such that* $\mathcal{I}(h^t; s) \neq \varnothing$ *is derived from full Bayesian updating—equations (10) and (12)—based on likelihood function* $l(h^t; s, \cdot)$*; and (b)*$\mathcal{M}_t(h^t; s) = \mathcal{M}_0$ *against any history* h^t *such that* $\mathcal{I}(h^t; s) = \varnothing$*.*

In Definition 4, condition (i) requires *sequential rationality*, that is, given the beliefs fit for any given history (either $\mathcal{I}(h^t; s) \neq \varnothing$ or not), all subsequent strategies of the seller must be best responses to the buyer's subsequent strategies, and all subsequent strategies of the buyer must be best responses to the seller's subsequent strategies too.

Condition (ii) requires that belief updating follows the "likelihood inference" based on the putative equilibrium s, as is described in Sections 3.1.2 and 3.1.3. In particular, it requires full Bayesian updating whenever Bayesian updating is feasible, i.e., $\mathcal{I}(h^t; s) \neq \varnothing$; and when $\mathcal{I}(h^t; s) = \varnothing$, that is, for any history such that every plausible equilibrium-based prediction of the play of the game is contradicted by the data, the set of all posteriors reverts to the set of all priors.

Our definition of POE mirrors the definition of PBE, but it differs from the definition of the PBE in two ways. First, in the POE the seller's strategy must maximize her infimum expected utility over non-conditionally-weakly-dominated strategies (Definition 3), which addresses the situation of ambiguity and reflects the seller's ambiguity aversion. By contrast, in a PBE, the seller's equilibrium strategy maximizes her expected utility function given the posterior beliefs that follow from its unique prior beliefs through Bayesian updating. Second, this (single prior) Bayesian updating is replaced in the POE by likelihood inference and full Bayesian updating on a set of multiple priors.

4.2 Refinement – Markov Perfect Objectivist Equilibrium (MPOE)

In the POE, in principle the strategies of the players can depend on the full history h^t in period t. This may permit various complicated strategic interaction (and feedback loops) between players and result in a large multiplicity of equilibria and diminished predictive power of the game-theoretic model. As a novel contribution to the study of this issue, we emphasize that epistemic under-determination of probabilistic beliefs may add a new cause to the problem of multiplicity of equilibria in general, and bargaining indeterminacy in particular. The call for equilibrium refinement therefore has some added urgency.

Applied game theorists conventionally employ the criterion of Markov perfection to tackle the problem of multiplicity of equilibria . The central idea of this principle is that an equilibrium that is Markov perfect represents the simplest form of (rational) strategic interaction. In practice, Markov perfection (typically) requires that actions are conditioned upon the smallest number of commonly known payoff-relevant state variables. Reducing the size of the space of states that the players have to keep track of and respond to can generate cognitive economies. That is, (everything else being equal) playing an equilibrium that is Markov perfect requires less (scarce) cognitive resources, therefore is more likely to emerge in reality. Given that a key motivation for inventing the POE solution concept is to emphasize the epistemic (cognitive) impact on behavior, we find it appealing to combine POE with Markov perfection whenever it is applicable. We call a POE which is also Markov perfect a *Markov perfect objectivist equilibrium* (MPOE).

Methodologically, to solve for MPOE, one can start with a conjectured small set of (commonly known) state variables which the players' strategies are conditioned upon. If there exist multiple POE formulated in terms of these state contingent strategies, then one can refine the equilibrium by reducing the numbers of (commonly known) state variables upon which POE still exist(s). It is necessary that MPOE is such a POE that further reduction of state variables is impossible.

5 Characterization of MPOE

In this sequential bargaining model, there is a clear list of candidate state variables for an MPOE. They include w_t and θ, which are payoff-relevant for the buyer. Although θ affects the seller's payoff, it is not known to the seller, so she cannot condition her strategy on θ. We therefore will start with a candidate MPOE for which the buyer's strategy is formulated in terms of state variables (w_t, θ), and the seller's strategy is independent of state variables. Because the set of commonly known state variables is empty and cannot be reduced further, this candidate equilibrium will be an MPOE if it is proven to be a POE.

To isolate the effect of asymmetric information, we first consider a benchmark case in which the seller knows the buyer's type θ. This benchmark analysis will highlight how information can allow the seller to leverage her exogenous right of making all offers into the endogenous bargaining power of hers.

5.1 The Benchmark: Seller Knows θ

From a social planner's perspective, $\Pi_B + \Pi_S$ represents the maximum social surplus from trade. Define $\theta^{**} := \frac{a}{\Pi_B+\Pi_S}$ to be the type of the buyer that would make a social planner just indifferent between sanctioning trade and not. Since for any $\theta \in [0, \theta^{**})$ it is the case that $\theta\,(\Pi_B + \Pi_S) < a$—that is, the social surplus from trade is negative—and trade does not occur in any SPNE. There exist many SPNE with an outcome of no trade, for example, $w_t \in [\theta^{**}\Pi_B, \Pi_B]$ and $k_t = 0$. Such multiplicity of SPNE is trivial and uninteresting. Thus we will focus on the case with $\theta \in (\theta^{**}, 1]$.

Proposition 2 *Suppose the buyer's type is* $\theta \in (\theta^{**}, 1]$*, which is common knowledge. Then there exists a unique SPNE, in which the seller's strategy is*

$$w(\theta) = \theta \Pi_B \tag{19}$$

and the buyer's strategy is

$$k(w_t; \theta) = \begin{cases} 1 & \text{if } w_t \leq \theta \Pi_B, \\ 0 & \text{otherwise.} \end{cases} \tag{20}$$

On equilibrium path there is always trade between the seller and the buyer. This SPNE is Pareto efficient.

We can define the *bargaining power* of a player in this bargaining game as her (or his) ability to carry out her (or his) preferred term of transaction despite the (potential or actual) resistance of the opponent. Naturally, each player's most preferred term of transaction—subject to that trade occurs voluntarily—is the one that maximizes her (or his) surplus from trade. In this sequential bargaining game, the seller has the right to make all offers, which is an exogenously endowed advantage in the bargaining game. For the simple case of perfect knowledge and positive discount rate r, the unique SPNE price $w(\theta) = \theta \Pi_B$ is indeed the most preferred by the seller—i.e., the seller has all of the total surplus from a Pareto-efficient trade—so we can say she has full bargaining power. Furthermore, the seller has sufficient information to assess her bargaining power, and she knows that her bargaining strategy allows her to achieve full bargaining power. In essence, this result corresponds to a first-degree price discrimination by a monopoly seller, which can be achieved only under complete information.

5.2 MPOE: Existence and Uniqueness

Theorem 1 *There exists a unique MPOE, denoted by* $(s^{**}, \mathcal{M}(s^{**}))$*. In this MPOE, the seller offers the price:*

$$s_B^{**}\left(h^t\right) = w^{**} := \theta^{**} \Pi_B \text{ for } t = 0, 1, \ldots, \tag{21}$$

and the buyer's strategy is

$$s_S^{**}\left(h^t\right) = k(w_t; \theta) = \begin{cases} 1 & \text{if } \theta \Pi_B - w_t \geq \max\{0, \frac{1}{1+r}\left(\theta \Pi_B - w^{**}\right)\}, \\ 0 & \text{otherwise,} \end{cases} \tag{22}$$

or equivalently

$$s_S^{**}\left(h^t\right) = k(w_t; \theta) = \begin{cases} 1 & \text{if } \theta \geq \max\left\{\frac{w_t}{\Pi_B}, \hat{\theta}(w_t)\right\}, \\ 0 & \text{otherwise,} \end{cases} \tag{23}$$

where $\hat{\theta}(w_t) = \frac{(1+r)w_t - w^{**}}{r \Pi_B}$ *is the cut-off type of the buyer who is indifferent between buying now at price* w_t *or buying at price* w^{**} *in the next period. On an equilibrium path the buyer buys if and only if he is of type* $\theta \in [\theta^{**}, 1]$*. In the limit, the MPOE price tends to the cost, i.e.,* $w^{**} \to a$ *as* $\Pi_S \to 0$.[28]

[28] Note, as a main result of the current paper, this theorem is sensitive to the assumption of positive Π_S. Under the alternative assumption $\Pi_S = 0$, we have $w^{**} = a$. By always offering the price $w^{**} = a$, the seller's payoff is always be zero. This strategy is weakly-dominated by the strategy of always offering price $a + \varepsilon$, for some small $\varepsilon > 0$, which gives a payoff of at least zero, and strictly positive under some plausible beliefs, e.g., the prior δ_θ for $\theta = 1$. So under the alternative assumption $\Pi_S = 0$, the propositional content of this theorem will be false.

The MPOE entails the seller offering the same price w^{**} each period. That price coincides with the price that an uninformed social planner would set in a *static* setting—i.e., it is the price that induces a buyer to purchase if and only if the net social benefit from trade is non-negative. We can interpret the offer price $w^{**} = a - \theta^{**}\Pi_S$ as a "cost-minus" contract in which the seller essentially offers the product to the buyer "at cost" minus a Pigouvian-like subsidy $\theta^{**}\Pi_S$, which is intended to induce the buyer to internalize the externality. In other words, the seller in the MPOE acts in the same way that a static social planner would act, taking into account the positive externality through Pigouvian subsidization.

Under the limiting condition: $\Pi_S \to 0$—i.e., as the positive externality becomes infinitesimal—the time path of price offer is a flat line at the competitive price a. Note that the game of seller offer sequential bargaining with one-sided incomplete information is well known to embed the Coase Problem—the limits on monopoly power created by monopolist's own future competition—which is based on the insight that the future self of the seller cannot commit to not undercutting her current price offer unless it is already at the cost level a. It implies that the time path of the price offer is either a price deflation process, or a flat line at the level of a. Now we see that the criterion of Markov perfection seems to have selected the latter against the former for the reason of simplicity.

The uniqueness of the MPOE is a striking result. It denies the possibility that the price offer is a flat line at a higher level $w > w^{**}$. The cause of this outcome is the same as the cause of the Coase problem—the seller's inability to commit her future self to not undercutting the current price w. To see why, note that after w is rejected, the seller must infer from the rejection that the buyer's type is such that given his strategy the price w is not acceptable. As is shown in detail in the proof of Theorem 1, with the updated belief about the buyer type, while the constant price w means no trade will occur ex post, a deviation to the price w^{**} will, guarantee a non-negative expected payoff for the seller, and—conditional on some posterior beliefs—also allow trade to occur and give the seller positive surplus from trade. Therefore every flat line of price offer at a higher price $w > w^{**}$ is ex post weakly dominated, and consequently the seller has no ability to commit to it because it is not sequentially rational.

A comparison between Proposition 2 and Theorem 1 sheds important light on the effect of information—or the lack of information—on affecting the endogenous distribution of bargaining power. Recall that the seller, who is endowed with the exogenous right to make all offers, has the upper hand—she gains full bargaining power—when bargaining under complete information. If her complete information about buyer type is replaced by complete ignorance, then, as predicted by the unique MPOE, the buyer has the upper hand in the bargaining—even full bargaining power for the limiting case $\Pi_S \to 0$—although the seller still keeps her (exogenous) right to make all the offers. Clearly, it takes adequate information for the seller to leverage the exogenously endowed bargaining advantage—e.g., the right to make all offers—into endogenous bargaining power. Thus, the possession of private information under severe informational asymmetry does seem to be a source of significant bargaining power. In the model the seller has severe information disadvantage, as is represented by her ambiguous probabilistic beliefs about buyer type. Because of the complete ignorance of the seller about buyer type, the seller has insufficient information to carry out any meaningful assessment of her bargaining power. Notice that she cannot even calculate her expected payoff from her bargaining strategy because she cannot determine a unique prior probability measure upon which the expected payoff can be defined. Since she cannot carry out an unambiguous assessment of her bargaining power, she cannot be meaningfully motivated to maximize her bargaining power. As a result, she has no obstacle (based on a concern about bargaining power) to the pragmatic pursuit of cognitive economies, and she should be willing to coordinate on equilibrium selection through the criterion of Markov perfection.

The buyer knows that the seller lacks incentive to experiment with higher prices to enhance her bargaining power. Therefore, this can epistemically and strategically induce the buyer to form a *deflationary expectation* in the face of high prices (i.e., $w_t > w^{**}$). The effect of this *deflationary expectation* can be demonstrated dramatically from the buyer's cut-off strategy (23). If the price exceeds a certain threshold, i.e., $w_t > w^{**} + \frac{r(\Pi_B - w^{**})}{1+r}$, then the cut-off type is $\hat{\theta}(w_t) > 1$, implying the price will be rejected by all buyer types in expectation of a price deflation in the next period from w_t to w^{**}. This kind of strategic resistance to a high price by the buyer under the influence of a *deflationary expectation* certainly does not encourage the seller to experiment with higher prices, but can strengthen the force of cognitive economies in facilitating a coordinated selection of the MPOE strategies. This is consistent with our novel argument for the use of Markov perfection for equilibrium refinement, which applies particularly when there exist multiplicity of POE associated with multiplicity of probabilistic beliefs. In the next section, we elaborate on these multiplicities.

6 Potential Bargaining Indeterminacy and Refinement

6.1 Multiplicity of POE

So far our equilibrium analysis of the sequential bargaining has been founded upon equilibrium refinement through Markov perfection. In the following, we approach this non-trivial issue of refinement by establishing two crucial points. First, there exists a multiplicity of POE and hence a potential for bargaining indeterminacy, which calls for refinement. Second, the refinement criterion—Markov perfection—is justified in a way immune to some well-argued objections informed by the related literature.

To start, under multiplicity of prior probability measures, the problem of optimal risk-taking is ambiguously defined. This can make the uninformed decision maker indifferent among multiplicity of solutions that can be justified among ambiguous beliefs. The following proposition establishes that there exists a continuum of POE for $r > 0$ (finitely).[29]

Proposition 3 *For any $w_0 \in [w^{**}, \bar{w}]$, where $\bar{w} := w^{**} + \frac{r(\Pi_B - w^{**})}{1+r} \in (w^{**}, \Pi_B)$, there exists a POE such that w_0 is the seller's initial price offer at $t = 0$. In a strategy profile which can support this POE outcome, the seller's strategy is given by:*

$$w(t) = \begin{cases} w_0 & \text{for } t = 0, \\ w^{**} & \text{otherwise,} \end{cases} \tag{24}$$

and the buyer plays the following cut-off strategy:

$$k(w_t; \theta) = \begin{cases} 1 & \text{if } \theta \geq \max\left\{\frac{w_t}{\Pi_B}, \hat{\theta}(w_t)\right\}, \\ 0 & \text{otherwise,} \end{cases} \tag{25}$$

where $\hat{\theta}(w_t) = \frac{(1+r)w_t - w^{**}}{r\Pi_B}$.

[29] For our purpose of establishing the existence of multiplicity of POE, we do not need to and do not claim that the continuum of POE specified in Proposition 3 include the full spectrum of POE. For example, the following proposition does not claim to rule out any POE which has its initial price offer in the interval $(\bar{w}, \Pi_B]$.

Intuitively, the exclusive right to make offers—if combined with perfect knowledge of the buyer's type—allows the seller to make the price offer such that the buyer is indifferent between accepting and rejecting, and thus extracts maximum surplus from the buyer without causing costly rejection. This is what Proposition 2 says. In contrast, the seller who is completely ignorant about the buyer's type risks causing costly rejection by offering high prices. Furthermore, the severe lack of information prevents the seller from assessing the risks objectively, and therefore permits her to entertain the possibilities of either high or low risks arbitrarily, simply because of her ambiguous beliefs over the buyer type.

Technically, the multiplicity of priors that represents the ambiguity gives the seller a "fat indifference curve" despite of the restrictiveness of the mix-inf expected utility criterion of choice. This "fat indifference curve" admits a multiplicity of best responses on the seller's behalf to the same strategy of the buyer. To see this in sharp focus, notice that all the multiple POE explicitly specified in Proposition 3 have the same buyer strategy as given by (25), which is identical to the buyer strategy in the unique MPOE as specified by (23). The multiplicity of seller strategies specified by (24) as well as in the unique MPOE are best responses to this same buyer strategy, and they constitute this multiplicity of POE. The insight of this observation is that we can precisely identify ambiguity—the multiplicity of priors—as a cause of this (particular) multiplicity of POE, through the transmission mechanism of the "fat indifference curve". For the same reason—the multi-prior-induced "fat indifference curve"—the seller is also indifferent among all the multiplicity of POE ex ante at the level of individual rationality.

6.2 Cognitive Economies as an Equilibrium Selection Force – When Does It Work?

The analysis in the pervious section thus calls for equilibrium refinement in order to tackle the problem of bargaining indeterminacy. Among the multiplicity of POE bargaining outcomes, some are less demanding on cognitive effort to acquire, store and process information about the state of the game than others. The unique MPOE outcome—the flat line of price offer at the level w^{**}—is actually the cognitively simplest and most economical among the continuum of POE—the other POE all involve some more complex process of price deflation. Since the seller is indifferent among the multiplicity of POE (at the individual rationality level), the cognitive economies can be seen as a "secondary benefit" (at the collective equilibrium selection level) for the seller to "break the tie" and be willing to select the unique MPOE among the multiplicity of POE.[30]

To anticipate Theorem 3 that will be presented in Section 7.1, it is well known that from the literature on (seller offer) sequential bargaining with one-sided incomplete information (Sobel and Takahashi 1983, Fudenberg, Levine and Tirole 1985, Gul, Sonnenschein and Wilson 1986, Ausubel and Deneckere 1989a, 1989b) that Markov perfection blocks the reputation equilibria, which are a main cause of multiplicity of equilibria and potential bargaining indeterminacy. In comparison with the Markov perfect Bayesian equilibrium (MPBE), the seller's most preferred reputation equilibrium can best enhance her bargaining power against the buyer. This gives the seller an incentive not to coordinate on playing the MPBE, which has the benefit of cognitive economies but is not favorable in terms of bargaining power. In this context, the

[30] This "secondary benefit" does not affect the choice at the individual rationality level, therefore does not affect the definition of POE.

argument for giving cognitive economies a decisive role of equilibrium refinement is not compelling, to say the least. As a result, refinement of PBE remains an unresolved problem.

A further discussion on using Markov Perfection for equilibrium refinement appears in Section 7.3. In its foreshadow, we argue that Markov Perfection is not a suitable general-purpose equilibrium refinement criterion. This entails that we have to provide a novel specific-purpose justification for the refinement of POE by the Markov Perfection criterion.

In Sections 7.1 and 7.2, we explain in detail a key difference between POE and PBE from the epistemic point of view, which justifies why (in the current sequential bargaining context) refining POE through Markov perfection is appealing while refining PBE through Markov perfection is objectionable. A formal principle—or a necessary and sufficient condition—for selecting an equilibrium which is Markov perfect is developed in the remainder of the current section, with its immediate implication also articulated.

In the previous section we identify a novel cause of potential indeterminacy in bargaining power relation between a seller and buyer. This cause stems from extraordinarily severe incompleteness of information (about the type of the buyer) accessible to the seller who has the right to make all offers. Since the ambiguity in probabilistic beliefs prevents the uniformed player from assessing her bargaining power meaningfully, it also prevents her from conducting a trade off between a gain in cognitive economies and a gain in bargaining power (against the opponent), thereby allowing the gain in cognitive economies to prevail in affecting the bargaining power relation. For cognitive economies to work as an effective equilibrium selection force, we propose that the following general formal necessary and sufficient condition:

Assumption 3 *A necessary and sufficient condition for an equilibrium that is Markov-perfect to be selected by the players collectively is that there is no countervailing incentive (in terms of payoff) for any player of the game to oppose this equilibrium in favor of any alternative equilibrium ex ante.*

Assumption 3 trivially implies the following lemma:

Lemma 1 *(i) If an equilibrium which is Markov-perfect is (weakly) Pareto dominant among all alternative equilibria ex ante, then it is selected by the players collectively. (ii) If an equilibrium which is Markov-perfect is (weakly) Pareto dominated by any alternative equilibrium ex ante, then it is not selected by the players collectively.*

7 "Coase Conjecture" Revisited

We now compare our analysis of the MPOE with its conventional counterpart—perfect Bayesian equilibrium. In particular, we highlight the main results in relation to the famous "Coase Conjecture". We also offer further thoughts on Markov perfection-based equilibrium refinement.

7.1 Conventional Dynamic Bayesian Game

In this section we investigate how the results based on our new equilibrium concept differ from the conventional perfect Bayesian equilibrium analysis. The conventional Bayesian approach assigns a unique prior probability measure over the state space. This implies in our example of seller offer sequential bargaining with one-sided incomplete information, the seller must be endowed with a prior probability measure over the unknown type θ of the buyer. If the unique probability measure is to be established with some sense of epistemic objectivity, then it must summarize a large amount of information. When the informational asymmetry is severe, the seller is unlikely to have access to such a large amount of information to ascertain a single prior probability measure. In the extreme, the seller can be completely ignorant about the true probability measure over the buyer's type. Our following discussion again focuses on this extreme case, with which we intend to highlight an important point—when incompleteness of information is severe, the conventional Bayesian approach has serious weakness, and in fact, its prediction can be misleading.

Let us start with the single prior probability measure that the Bayesian approach presumes—even when complete ignorance is the true epistemic situation. With the Bayesian approach, a uniform probability distribution is typically used to approximate complete ignorance—though the foundations of this approach are logically dubious[31] (see Edward, 1992). In this section we explore perfect Bayesian equilibrium based on the uniform prior. When interpreting the results obtained here one needs to distinguish the two opposing interpretations of the unique prior probability distribution: (1) the uniform distribution represents epistemically objective knowledge about θ. (2) the uniform distribution represents epistemically subjective pretense of knowledge against the backdrop of the epistemic state of complete ignorance about θ. In case the second interpretation is the valid one, we argue that the prediction of the conventional PBE and the problem identified thereby may be misleading, and an alternative ambiguity-based approach—e.g., the MPOE prediction—should be pursued.

To ease comparison with the existing literature, in this section we abstract the common interest component of the current model by focusing on the limiting $\Pi_S \to 0$. Consequently we have $w^{**} := \frac{\Pi_B a}{\Pi_B + \Pi_S} \to a$ and $\theta^{**} := \frac{a}{\Pi_B + \Pi_S} \to \frac{a}{\Pi_B}$; and the analyses of the conventional MPBE and PBE can be directly based on the existing results from the literature (Fudenberg, Levine and Tirole, 1985; Gul, Sonnenschein and Wilson, 1986; Ausubel and Deneckere, 1989a,b). Although the results are well known, for the sake of developing useful intuition we present the analyses in considerable detail.

We first characterize a pure strategy Markov perfect Bayesian equilibrium of the seller offer sequential bargaining game. This analysis is based on (Fudenberg, Levine and Tirole, 1985). Suppose in this MPBE, the seller's Markov strategy is given by $w\left(q_t^{\sup}\right) \in [0, \Pi_B]$, where state variable $q_t^{\sup} \in [0,1]$ is an estimate of the supremum of the updated buyer's type space. The buyer's Markov strategy is given by the cut-off strategy:

$$k\left(w_t, q_t^{\sup}; \theta\right) = \begin{cases} 1 & \text{if } \theta \geq \hat{\theta}\left(w_t, q_t^{\sup}\right), \\ 0 & \text{otherwise,} \end{cases} \tag{26}$$

where $\hat{\theta}\left(w_t, q_t^{\sup}\right)$ is the cut-off type of the buyer, who is indifferent between accepting w_t and accepting $w\left(\hat{\theta}\left(w_t, q_t^{\sup}\right)\right)$ in the next period; $\hat{\theta}\left(w_t, q_t^{\sup}\right)$ is such that

$$\hat{\theta}\left(w_t, q_t^{\sup}\right)\Pi_B - w_t = \frac{\hat{\theta}\left(w_t, q_t^{\sup}\right)\Pi_B - w\left(\hat{\theta}\left(w_t, q_t^{\sup}\right)\right)}{1+r}. \tag{27}$$

[31] Since the single prior probability measure is by assumption not based on real information, it must be based on pure epistemic subjectivity. Nevertheless it appears to be informative because it is indistinguishable from the identical prior distribution which is supported by genuine information; therefore it cannot logically represent complete ignorance, but some subjective pretense of knowledge.

This equation implies a relationship between function $\hat{\theta} : \mathcal{W} \times \Theta \to \Theta$ and function $w : \Theta \to \mathcal{W}$ in equilibrium. Taking the function w (i.e., the seller's strategy) as given, the above equation defines a functional equation to which the function $\hat{\theta}$ is its solution. In general, if the function w varies, the functional equation and its solution $\hat{\theta}$ also vary. When the buyer's strategy function k is taken as given, then the cut-off type function $\hat{\theta}$ must be taken as given; consequently, the underlying function w must be taken as given. Here the function w thus has another interpretation: it is the buyer's expectation (or representation) of the seller's strategy.

Conditional on the game has not ended, the evolution of the state variable $q_t^{\sup}$ is given by the following law of motion:

$$
\begin{aligned}
q_{t+1}^{\sup} &= \min \left\{ q_t^{\sup}, \hat{\theta}\left(w_t, q_t^{\sup}\right) \right\} \text{ for } t = 0, 1, \ldots, \\
q_0^{\sup} &= 1.
\end{aligned}
$$

Note that if the seller believes (with certainty) that the real type $\theta \leq q_t^{\sup} \in (\theta^{**}, 1]$, it will be suboptimal to choose w_t such that $\hat{\theta}\left(w_t, q_t^{\sup}\right) > q_t^{\sup}$ because the offer will be rejected with certainty, leaving the state variable unchanged, i.e., $q_{t+1}^{\sup} = q_t^{\sup}$. Such an offer has the cost of delaying trade without any benefit. So for $q_t^{\sup} \in (\theta^{**}, 1]$, the optimal choice must be such that $q_{t+1}^{\sup} = \hat{\theta}\left(w_t, q_t^{\sup}\right) \leq q_t^{\sup}$. If the function $\hat{\theta}$ is strictly monotone in w_t then there is a one-to-one mapping between w_t and $q_{t+1}^{\sup}$, and the seller's move can be seen as a choice of the next period value of state variable $q_{t+1}^{\sup} = \hat{\theta}\left(w_t, q_t^{\sup}\right)$, which is implemented by price w_t. The relation between w_t and $q_{t+1}^{\sup}$ is determined by (27), and takes the following form:

$$
w_t = \frac{r\Pi_B q_{t+1}^{\sup} + w\left(q_{t+1}^{\sup}\right)}{1+r}, \tag{28}
$$

where function w represents the buyer's expectation of the seller's strategy, which affects the definition of the cut-off type and the buyer's strategy. It is also obvious that a choice of $q_{t+1}^{\sup} < \theta^{**}$ is suboptimal since trading with any type $\theta < \theta^{**}$ has negative payoff for the seller.[32]

Taking as given the buyer's strategy, which is captured by equation (28), the seller's optimal strategy can be formulated as the solution to a dynamic programming problem: it chooses the next period value of state variable $q_{t+1}^{\sup}$ given the current value $q_t^{\sup}$. As necessary for sequential rationality, taking its implied value function W as given, the optimal strategy must maximize the right-hand side of the following Bellman equation for all $q_t^{\sup} \in [\theta^{**}, 1]$,

$$
W\left(q_t^{\sup}\right) = \max_{q_{t+1}^{\sup} \in [\theta^{**}, q_t^{\sup}]} \left\{ \left[F\left(q_t^{\sup}\right) - F\left(q_{t+1}^{\sup}\right)\right] \left[\frac{r\Pi_B q_{t+1}^{\sup} + w\left(q_{t+1}^{\sup}\right)}{1+r} - a\right] + \frac{1}{1+r} W\left(q_{t+1}^{\sup}\right) \right\}, \tag{29}
$$

where $F(\theta)$ is the cumulative distribution function for the uniform prior; and the seller's (expected) value function W must be the solution to the above functional equation for all $q_t^{\sup} \in [\theta^{**}, 1]$. For $q_t^{\sup} < \theta^{**}$, we must have $W\left(q_t^{\sup}\right) = 0$ and it is obvious that the constant price offer $w\left(q_t^{\sup}\right) = w^{**}$ can attain this result. The closed-form solution of an MPBE is given below.

Proposition 4 *Suppose $\Pi_S \to 0$ and the seller is an expected utility maximizer endowed with a uniform prior probability measure with the probability density function $\frac{dF(\theta)}{d\theta} = 1$ for $\theta \in [0, 1]$. Then there exists an*

[32] Recall that by definition θ^{**} is the cutoff type of buyer with whom the social surplus from trade is zero. If the true type is $\theta < \theta^{**}$ then the social surplus from trade is negative; because the buyer does not accept negative surplus from trade, the seller's surplus from trade must be negative.

MPBE for the sequential bargaining game such that the seller plays the Markov strategy:

$$w\left(q_t^{\sup}\right)=\begin{cases}\beta\Pi_B\left(q_t^{\sup}-\theta^{**}\right)+w^{**} & \textit{if } q_t^{\sup}>\theta^{**},\\ w^{**} & \textit{otherwise,}\end{cases}\tag{30}$$

where

$$\beta=-r+\sqrt{r^2+r},\tag{31}$$

and the buyer plays the Markov cut-off strategy:

$$k\left(w_t,q_t^{\sup};\theta\right)=\begin{cases}1 & \textit{if } \theta\geq\hat{\theta}\left(w_t,q_t^{\sup}\right),\\ 0 & \textit{otherwise,}\end{cases}\tag{32}$$

where the cut-off type is

$$\hat{\theta}\left(w_t,q_t^{\sup}\right)=\begin{cases}\frac{1+r}{(r+\beta)\Pi_B}\left(w_t-w^{**}\right)+\theta^{**} & \textit{if } q_t^{\sup}>\theta^{**},\\ \frac{1+r}{r\Pi_B}\left(w_t-w^{**}\right)+\theta^{**} & \textit{otherwise.}\end{cases}\tag{33}$$

The following corollary connects the MPBE characterized in Proposition 4 and the famous "Coase Conjecture".

Corollary 1 *Suppose $\Pi_S\to 0$, the MPBE characterized in Proposition 4 has the following properties:*
(i) The expected value for the seller, for $q_t^{\sup}\in[\theta^{**},1]$, is given by

$$W^{MPBE}\left(q_t^{\sup}\right)=\frac{\beta\Pi_B\left(q_t^{\sup}-\theta^{**}\right)^2}{2};\tag{34}$$

(ii) This MPBE implies the following features of the "Coase Conjecture":

$$\lim_{\Pi_S\to 0,r\to 0}w\left(q_0^{\sup}\right)=a,\ \lim_{\Pi_S\to 0,r\to 0}W^{MPBE}\left(q_0^{\sup}\right)=0.\tag{35}$$

As exemplified in part (ii) of Corollary 1, the "Coase Conjecture" features that as $r\to 0$, the price converges to the competitive level "in the twinkling of an eye" (Coase, 1972); also the buyer has almost the full bargaining power as $r\to 0$. The "Coase Conjecture", which focuses on the limiting case $r\to 0$, illustrates the more general Coase Problem—the limits on monopoly power created by monopolist's own future competition. At the heart of the Coase Problem is a commitment problem—the monopoly seller cannot commit her future self to not undercutting her current price offer if it is above the cost level a. Intuitively, for $r>0$, the payoff discounting on the buyer's behalf makes the good supplied in the future inferior to the good supplied currently, therefore allows the current price offer to contain a premium over the future price offer. This also sets in motion a price deflation process and causes a price *deflation expectation* whenever the price on offer is above the competitive level. For the limiting case of $r\to 0$, the physically identical goods supplied in the future and the current become almost perfect substitutes. Consequently the premium contained in the current price offer over the future offer is almost entirely eliminated. The lack of commitment on the seller's behalf then causes the initial price offer to almost immediately converge to the limit value of the price deflation process, which is the competitive price.

Although we derive the features of the "Coase Conjecture" from a particular MPBE, as is well known in the literature (Gul, Sonnenschein and Wilson, 1986), this MPBE turns out to be representative of all MPBE as far as the limiting results are concerned. The following theorem presents a more general version of the "Coase Conjecture", which is based on Theorem 3 of Gul, Sonnenschein and Wilson (1986). For the sake of brevity, we refrain from producing a formal proof of this theorem in this paper. Interested readers should consult their original proof.

Theorem 2 ("Coase Conjecture" (Gul, Sonnenschein and Wilson (1986))) *Suppose* $\Pi_S \to 0$ *and the seller is an expected utility maximizer endowed with a uniform prior probability measure with the probability density function* $\frac{dF(\theta)}{d\theta} = 1$ *for* $\theta \in [0,1]$. *For any* $\varepsilon > 0$ *and any Markov perfect Bayesian equilibrium (MPBE), there exists* $\bar{r} > 0$ *such that for any* $r \in (0, \bar{r})$ *the initial price offer in this MPBE is* $w_0 < a + \varepsilon$.

The "Coase Conjecture", however, only applies to the MPBE; it can be reversed when PBE is used as the solution concept instead of MPBE. The Folk Theorem presented below is based on Ausubel and Deneckere (1989a,b). These authors established that a large set of seller's expected payoffs—ranging from 0 ("the Coase Conjecture") to the "static monopoly profit"—can be obtained by a continuum of reputation equilibria (PBE) for $r \to 0$—which depend on both players playing some coordinated trigger strategies. We present the Folk Theorem first, followed by an intuitive explanation about how the trigger strategies work.

Theorem 3 (Reputation Equilibria – Folk Theorem) *Suppose* $\Pi_S \to 0$ *and the seller is an expected utility maximizer endowed with a uniform prior probability measure with the probability density function* $\frac{dF(\theta)}{d\theta} = 1$ *for* $\theta \in [0,1]$. *For* $r \to 0$, *any expected ex ante seller payoff* $W_0 \in \left(0, \frac{(\Pi_B - a)^2}{4\Pi_B}\right)$ *(where* $\frac{(\Pi_B - a)^2}{4\Pi_B}$ *is the "static monopoly profit") can be supported in a PBE.*

The trigger strategies underpinning this Folk Theorem—which are formally defined in the proof of the theorem—have the following feature: the seller's initial price offer w_0 is above the cost a and—conditional on that the game has not ended—the seller's strategy is to let the price offer deflate exponentially towards a, as is determined by following equation:

$$w_t - a = \left(\frac{1}{1 + \eta r}\right)^t (w_0 - a) \text{ for } t = 0, 1, \ldots; \tag{36}$$

where $\frac{1}{1+\eta r}$ is the contraction factor, which has the limiting value: $\lim_{(\eta r)\to 0} \frac{1}{1+\eta r} = 1$; $\eta > 0$ is a free strategy-specific parameter which affects the contraction factor $\frac{1}{1+\eta r}$: for given r, smaller η means slower price deflation. For $r \to 0$, the price deflation process is extremely slow. The trigger strategies of both players are contingent on a state variable $c_t \in \{1, 0\}$—$c_t = 1$ indicates the seller has a reputation for being a tough negotiator, and a credibility to sustain an extremely slow price deflation process. If the price mark-up $(w_t - a)$ ever deviates from this slow exponential contraction as described by (36) in period t, it will trigger $c_\tau = 0$ for all $\tau = t+1, t+2, \ldots,$—meaning the seller has forever lost her reputation—and it will induce both players to play the subsequent strategies of the MPBE described in Proposition 4—which, for $r \to 0$, make the price offer immediately to converge to a, and therefore serves as a punishment and deterrence for the seller to deviate. By way of creating and sustaining a payoff-enhancing reputation (or credibility), the seller can alleviate the commitment problem that is at the heart of the "Coase Conjecture". Interestingly, Ronald Coase (1972) not only conceived the "Coase Conjecture", but also famously pointed out how it could be reversed through contractual or institutional means in the durable good monopoly context. For example, the monopolist seller of the durable good could sign a contract with buyers promising to refund the price difference if she in the future sells the good at a lower price. This way the monopolist can commit her future self to not undercutting the current price offer, therefore induces the buyer to accept the initial "static monopoly price". Similarly, the institutional mechanism featured in a reputation equilibrium also commits the future self of the seller to not undercutting current price offer aggressively—i.e., to commit to deflating the price offer extremely slowly.

If Markov perfection is used as a refinement, the multiplicity of equilibrium outcomes can be eliminated for the limiting case of $r \to 0$, and the "Coase Conjecture" will be sustained. There is, however, no compelling argument to support this refinement. It eliminates all reputation equilibria, and it selects the MPBE only because the latter is simpler than the former. Given the fact that the seller's ex ante expected payoff is lower in the MPBE than any (interesting) PBE reputation equilibrium, a counter argument therefore can be made—for the seller there is a trade off between simplicity (cognitive economies) and loss of expected payoff and bargaining power; there is no obvious reason why the seller should be willing to coordinate with the buyer to play the MPBE, which is based on the prevalence of simplicity.

As an implication of Assumption 3, we can summarize the above arguments formally as the following:

Proposition 5 *For $r \to 0$, no MPBE in (the conventional Bayesian version of) the bargaining game (with uniform prior) is selected by the players collectively.*

Given that Markov perfection is an objectionable criterion in the context of (Bayesian) sequential bargaining, the refinement of PBE thereby requires some alternative criterion. Without a compelling criterion for equilibrium refinement emerging in this literature, bargaining indeterminacy remains unresolved theoretically. At least, it is fair to claim that the "Coase Conjecture" is reversed in the conventional Bayesian model of seller offer sequential bargaining with one-sided incomplete information. To understand why the "Coase Conjecture" is challenged, notice that the presumed uniform single prior allows the seller to identify the "static monopoly price", and also makes it plausible for her to commit to this "monopoly pricing" with the reputation mechanism. This strategic feasibility is underpinned by an epistemic commitment of the seller to the truth of the uniform single prior. Of course, in the extreme case of complete ignorance about the buyer's type, there is no rational ground for such an epistemic commitment of the seller to the truth of the uniform single prior, and the above analysis will not be suitable.[33] Instead, the application of our model of sequential bargaining under ambiguity and its MPOE prediction become a better fit.

7.2 "Coase Conjecture" 2.0

Interestingly, the "Coase Conjecture" finds a new incarnation in the unique MPOE in our model of sequential bargaining under ambiguity. Recall that $w_0 = w^{**} \to a$ as $\Pi_S \to 0$. Notably, this new incarnation does not rely on the limiting condition $r \to 0$. That is, for the new "Coase Conjecture" to hold, the good supplied by any future self of the seller does not have to be an almost perfect substitute of the good currently supplied.

[33] In the durable good monopoly setting that was analyzed by Coase (1972), the equivalent epistemic situation to the ambiguity of beliefs in our model would be a complete ignorance on the monopolist's behalf about the true demand function. Therefore, the monopolist would have a complete ignorance about what the true static monopoly price is.

In contrast to the conventional Bayesian model, our model with ambiguity about buyer types allows a convincing argument to justify the application of Markov perfection to refine the multiplicity of POE. Here the potential bargaining indeterminacy based on multiplicity of POE is caused by the large multiplicity of priors. Consequently, the epistemic justification for using Markov perfection as a refinement criterion in MPOE differs profoundly from that in MPBE. The criterion of Markov perfection selects the simplest (or the cognitively most economical) strategic interactions among multiplicity of equilibria. While the argument to justify the selection of MPBE among multiplicity of PBE is not compelling or not even appealing because it implicitly (and controversially) assumes that the seller generally favors simplicity in the trade off between cognitive economies and bargaining power. The same weakness does not affect the MPOE, because there is no discernible trade off between simplicity and payoff for the seller since her max-inf expected utility over all non-conditionally-weakly-dominated strategies are the same, namely zero.

To understand the last argument, note that in no POE does the offer price go below w^{**}, because such an offer gives the seller a negative infimum of expected payoff, which is inferior to the offer of w^{**}. Given this, two implications can be drawn: First, for buyer type $\theta = 0$, the expected seller payoff is zero for all POE, which must be the max-inf expected utility for the seller in all POE.

The second implication is that the unique MPOE gives the buyer the maximum payoff among all POE because the price w^{**} is the lowest among them. Therefore the buyer should have no incentive to oppose the unique MPOE in favor of any alternative POE ex ante.

As an implication of Assumption 3, we can summarize the above arguments formally as the following:

Proposition 6 *The unique MPOE in the ambiguity version of the bargaining game is selected by the players collectively.*

It is also worth mentioning that this new "Coase Conjecture" version 2.0 based on MPOE is robust to reversion that may be caused by a reputation equilibrium, because even if such a reputation equilibrium does exist as a POE, it is not an MPOE, and therefore it cannot survive the refinement based on Markov perfection.

7.3 Markov Perfection for Equilibrium Refinement? It Depends

According to Maskin and Tirole (2001), one practical reason for Markov perfection in applied game theory is that it "is often quite successful in eliminating or reducing a large multiplicity of equilibria in dynamic games, and thus in enhancing the predictive power of the model." Behind such practical virtue, there is also a philosophical appeal that "Markov strategies prescribe the *simplest form of behavior that is consistent with rationality*." For dynamic games with observable actions (i.e., perfect or almost perfect information), Maskin and Tirole (2001) demonstrate that this philosophical principle can be formulated precisely and consistently to give exact formal definitions of Markov strategy and Markov Perfect Equilibrium (MPE). "Informally, a Markov strategy depends only on payoff-relevant past events. More precisely, it is measurable with respect to the coarsest partition of histories for which, if all other players use measurable strategies, each player's decision-problem is also measurable."

Inspired by this deep insight, we require Markov strategy to "prescribe the *simplest form of behavior that is consistent with rationality,*" which is interpreted in the specific contexts of the models. For MPOE example analyzed in this paper, it is impossible to reduce the lists of state variables any further. One potential benefit of focusing the analysis on Markov equilibrium is to avoid multiplicity of equilibrium outcomes. This is achieved by the unique MPOE of our seller offer sequential bargaining game (see Theorem 1).

One of the motivations for favoring Markov strategy is the view that it can help avoid variables which are not directly payoff-relevant, but facilitate the "bootstrapping" property. According to Maskin and Tirole (2001), "[Markov] strategies depend on as few variables as possible; they involve no complex 'bootstrapping' in which each player conditions on a particular variable only because others do the same." It is, however, well known that this goal is not always desirable. Notably, the Markov perfect equilibrium of the infinitely repeated prisoners' dilemma game eliminates all cooperative (or collusive) equilibria, which Pareto dominate the MPE—a violation of Lemma 1. Also, the Markov perfect Bayesian equilibrium (MPBE) of the (seller offer) sequential bargaining game with one-sided incomplete information eliminates all the reputation equilibria, which enhance the bargaining power of the seller—a violation of Proposition 5. To say the least, there exists no compelling argument to justify the Markov equilibria in these contexts. In contrast, when we apply Markov perfection as a refinement criterion to the set of POE, we observe that the MPOE does not select any (weakly) Pareto dominated POE, or reduce bargaining power of any player. This is confirmed by the fact the selection of the MPOE does not reduce the uninformed seller's bargaining power while enhancing the informed buyer's bargaining power—a consistency with Assumption 3 and Proposition 6.

On the surface, the MPOE stands out among the multiplicity of POE because it is the simplest. Another contributing factor for the MPOE to be selected is that because of the severe epistemic disadvantage the seller has about the buyer's type, she lacks the incentive and willingness to experiment with higher prices in order to enhance her bargaining power. This fact about the seller's lack of incentive, which is known to the buyer, gives the buyer a re-enforced incentive to resist higher prices, on the basis of a strategically justified *deflationary expectation*. These incentive factors thus join force with cognitive economies to facilitate the selection of the MPOE, which underpins the "Coase Conjecture" 2.0.

8 Summary and Concluding Remarks

In this paper, we study a class of dynamic games of incomplete information in the context of sequential bargaining with one-sided incomplete information. A seller has the exclusive right to make offers to a buyer, whose valuation of the transaction is his private information and thus his type. If the offer is accepted, trade occurs; otherwise the seller can revise the offer in the next period. Payoffs for both players are discounted at a common positive rate. Unlike conventional Bayesian models of bargaining with asymmetric information, we do not assume that the seller has a unique prior distribution over the buyer's type $\theta \in [0, 1]$. Instead, we allow the seller to hold multiple prior probability measures over θ. We assume the seller has completely ignorance about value of θ in order to illustrate that the degree of incompleteness of information can significantly exceed what a conventional Bayesian model permits. This assumption is represented by the set of all plausible prior probability measures over $[0, 1]$. This gives rise to a process by which the seller iteratively selects beliefs over θ based on observational facts and a likelihood function that is derived from the common knowledge about the game and the concept of equilibrium. Following the works of Manski (2008) and GMMS (2010), the seller is assumed to maximize her infimum expected utility over non-weakly-dominated strategies. As a solution concept for the game, we propose an extension of perfect Bayesian equilibrium to a setting in which the uninformed seller has multiple priors. We call this extension the perfect objectivist equilibrium (POE). We show that the epistemic under-determination of probabilistic belief constitutes a novel cause for potential bargaining indeterminacy—i.e., multiplicity of POE—and call for its resolution through equilibrium refinement based on the criterion of Markov perfection. We fully characterize the Markov perfect objectivist equilibrium (MPOE) for the illustrative model, and establish the uniqueness of MPOE.

The refinement of perfect Bayesian equilibrium (PBE) in the bargaining context remains an unresolved problem partially because the refinement through Markov perfection is objectionable. We nevertheless argue that the refinement of POE through Markov perfection is better justifiable on an epistemic ground. The selection of Markov perfect Bayesian equilibrium (MPBE) among the multiplicity of PBE, including a continuum of reputation equilibria, entails the objectionable assumption that in the trade off by the uninformed seller between expected payoff and cognitive economies the latter prevails in general. The selection of MPOE among the multiplicity of POE does not entail such an assumption because the multiplicity of POE is caused by the ambiguity of probabilistic beliefs (or ignorance), which also makes the expected payoff for the uninformed seller undefined. Consequently this removes any meaningful trade off between expected payoff and cognitive economies and hence allows the latter to prevail. Intuitively, the severe ambiguity of probabilistic beliefs of the uninformed seller weakens her sense of the "surplus extraction" component in the bargaining game, strengthens her sense of the "surplus creation" component, and therefore transforms the bargaining game into a "coordination game" that incentivizes her to maximize the social surplus from trade. This result can be seen as a reincarnation and resurrection of the famous "Coase Conjecture"—in the face of severe ignorance of the buyer type or demand, an apparent monopoly results in a competitive outcome. Not only does this not require the discount rate to approach zero, it is also robust to reversion that may be caused by a reputation equilibrium. To distinguish from the original Coase Conjecture and the "Coase Conjecture" based on the conventional Bayesian models, we can call this new "Coase Conjecture" characterized in the current paper "Coase Conjecture" 2.0.

Our analysis has a striking implication for the post-crisis secular economic growth: After major financial or economic crises like the historical Great Depression or the recent Great Recession, when the intellectual and epistemic confidence of many entrepreneurs and investors are deeply shaken, their beliefs about future demands may well be characterized as more ambiguous in the sense which is the theme of the current paper. This implies that the Coase Problem—the limits on monopoly power created by monopolist's own future competition—probably should be taken more seriously in the post-crisis era.[34] After major financial or economic crises, it is likely that demand functions for novel durable goods are affected by uncertainty. Significantly, these will increase the ignorance and ambiguity on prospective entrepreneurs' behalf about the demand functions, then the "Coase Conjecture" 2.0 implies a shift of bargaining power from (uninformed) durable good monopolists to the buyers, which corrodes monopoly rents. This shift of bargaining power can manifest itself in buyers' *deflationary expectation* in the face of high prices. When the *deflationary expectation* becomes deeply rooted in buyers' strategic intentionality, the informationally seriously disadvantaged sellers can be easily persuaded to accept it and the associated bargaining power shift. Why does it matter? Well,

[34]Broadly speaking, taking the Coase Problem more seriously implies, for microeconomic analysis and industrial organization, some rethinking about analyzing market power in imperfectly competitive markets. For example, the conventional assumption that the true demand function is automatically and entirely known to the monopoly or oligopoly firm(s) seems to deserve qualification and maybe altering on epistemic ground. The gain from this change can be a better understanding of the relation between information (or a serious lack of information) about the demand function and market power.

if Schumpeter (1934) was right to claim that monopoly rent is the key to make innovation to pay for itself, and innovation is the engine of economic growth, then this "Coase Conjecture" 2.0 may deprive the post-crisis growth of an engine fueled by monopoly rents. It matters because this could be a contributing factor to *secular stagnation*, a notion which originated from Alvin Hansen (1939), and after being brought to prominence by Larry Summers recently, is becoming a serious concern for contemporary economists (Teulings and Baldwin, 2014).[35] To appreciate why this "Coase Conjecture" 2.0 may be relevant to the concern over secular stagnation, lets start with a quote from Summers (2014): " Perhaps Say's dubious law has a more legitimate corollary—'Lack of Demand creates Lack of Supply'." Then one can go a step further to suggest: "Lack of Knowledge about Demand Function creates Lack of Supply." The ignorance or ambiguity about the demand function, according to "Coase Conjecture" 2.0, limits a prospective innovator's power to extract monopoly rent—Schumpeterian entrepreneurial profit—therefore discouraging the supply of innovation in the first place.

Finally, we have two caveats: First, we believe the "Coase Conjectures" (old or new) are merely the poster child of the profound Coase Problem. Although the "Coase Conjectures" are extreme results based on some extreme assumptions, the profundity of the Coase Problem is unlikely to vanish even if these extreme assumptions are relaxed. Second, admitting the possibility of Secular Stagnation does not suggest that it is inevitable. For example, if the corrosion of private sector entrepreneurial profits by the Coase Problem is a key issue, then public subsidies for (private sector) entrepreneur-lead innovations can address the problem. Then a real challenge will be how public policies can achieve this goal intelligently under their own constraints of severe incompleteness of information and ambiguity.

References

[1] Anscombe, F. J. and R. J. Aumann (1963), "A Definition of Subjective Probability," *Annals of Mathematical Statistics*, 34 (1), pp. 199-205.

[2] Ausubel, Lawrence M. and Raymond J. Deneckere (1989a), "Reputation in Bargaining and Durable Goods Monopoly," *Econometrica*, 57 (3), pp. 511-531.

[3] Ausubel, Lawrence M. and Raymond J. Deneckere (1989b), "A Direct Mechanism Characterization of Sequential Bargaining with One-Sided Incomplete Information," *Journal of Economic Theory*, 48, pp. 18-46.

[4] Ausubel, Lawrence M. and Peter Cramton and Raymond J. Deneckere (2000), "Bargaining with Incomplete Information," *Handbook of Game Theory* (Aumann, Robert J. and Sergiu Hart, Ed.), Vol. 3, Ch. 50, Amsterdam: Elsevier Science B.V.

[35] According to the "Coase Conjecture" 2.0, although major financial or economic crises can exacerbate the Coase Problem, they do not have to be the exclusive cause of it. Other factors, such as, demographic changes—e.g., a decline of population growth rate—can also trigger demand changes and an epistemic deficiency about them. Consequently, "Coase Conjecture" 2.0 phenomenon can arise, and cause economic slow down. This in turn may contribute to deflationary pressure in the real sector of the economy and asset price bubble in financial markets, which result in financial instability and crises. This pattern seems to fit well empirically with the Great Depression, the Japanese Lost Decades, and, most recently, the Great Depression and Global Crisis.

[5] Bewley, Truman F. (2002), "Knightian Decision Theory. Part I," *Decisions in Economics and Finance*, 25, pp. 79-110. (The working paper version appeared in 1986.)

[6] Coase, R. H. (1972), "Durability and Monopoly" *Journal of Law and Economics*, 15 (1), pp. 143-149.

[7] de Finetti, Bruno (1937), "La Prévision: Ses Lois Logiques, Ses Sources Subjectives," *Annales de l'Institute Henri Poincaré*, 7, pp.1-68.

[8] Dudley, R. M. (2002), *Real Analysis and Probability*, Cambridge University Press, Cambridge.

[9] Edwards, A.W.F. (1992), *Likelihood* (Expanded Edition), Johns Hopkins University Press, London.

[10] Efron, Bradley (1998), "R. A. Fisher in the 21th Century," *Statistical Science*, 13 (2), pp. 95-114.

[11] Epstein, Larry G. and Martin Schneider (2003), "Recursive Multiple-Priors," *Journal of Economic Theory*, 113, pp. 1-31.

[12] Epstein, Larry G. and Martin Schneider (2007), "Learning under Ambiguity," *Review of Economic Studies*, 74, pp. 1275-1303.

[13] Fisher, R. A. (1925), *Statistical Methods for Research Worker*, Oliver and Boyd, Edinburgh.

[14] Fudenberg, Drew and David Levine and Jean Tirole (1985), "Infinite Horizon Models of Bargaining with One-sided Incomplete Information," in *Bargaining with Incomplete Information* (Alvin Roth, Ed.), London/New York, Cambridge University Press, pp. 77-98.

[15] Fudenberg, Drew and Jean Tirole (1991), *Game Theory*, MIT Press, Cambridge, Massachusetts.

[16] Gilboa, Itzhak and David Schmeidler (1989), "Maxmin Expected Utility with Non-unique Prior," *Journal of Mathematical Economics*, 18, pp. 141-153.

[17] Gilboa, Itzhak and Fabio Maccheroni and Massimo Marinacci and David Schmeidler (2010), "Objective and Subjective Rationality in a Multiple Prior Model," *Econometrica*, 78 (2), pp. 755-770.

[18] Gilboa, Itzhak and Massimo Marinacci (2011), "Ambiguity and the Bayesian Paradigm," Working Paper.

[19] Gul, F., H. Sonnenschein and R. Wilson (1986), "Foundations of Dynamic Monopoly and the Coase Conjecture," *Journal of Economic Theory*, 39, pp. 155-190.

[20] Hansen, Alvin H. (1939), "Economic Progress and Declining Population Growth," *American Economic Review*, 29, pp1-15.

[21] Harsanyi, John C. (1967), "Games with Incomplete Information Played by 'Bayesian' Players, I-III. Part I. The Basic Model," *Management Science*, 14 (3), Theory Series, pp. 159-182.

[22] Harsanyi, John C. (1968a), "Games with Incomplete Information Played by 'Bayesian' Players, I-III. Part II. Bayesian Equilibrium Points," *Management Science*, 14 (5), Theory Series, pp. 320-334.

[23] Harsanyi, John C. (1968b), "Games with Incomplete Information Played by 'Bayesian' Players, I-III. Part I. I-III. Part III. The Basic Probability Distribution of the Game," *Management Science*, 14 (7), Theory Series, pp. 486-502.

[24] Kennan, John and Robert Wilson (1993), "Bargaining with Private Information," *Journal of Economic Literature*, 31, pp. 45-104.

[25] Kreps, David M. and Robert Wilson (1982), "Sequential Equilibria," *Econometrica*, 50 (4), pp. 863-894.

[26] Manski, Charles F. (1995), *Identification Problems in Social Sciences*, Harvard University Press, Cambridge, Massachusetts.

[27] Manski, Charles F. (2008), "Partial Prescriptions for Decisions with Partial Knowledge," NBER Working Paper 14396.

[28] Maskin, Eric and John Riley (1984), "Monopoly with Incomplete Information," *RAND Journal of Economics*,15(2), pp. 171-196.

[29] Maskin, Eric and Jean Tirole (2001), "Markov Perfect Equilibrium I. Observable Actions," *Journal of Economic Theory*, 100, pp. 191-219.

[30] Ramsey, Frank P. (1926), "Truth and Probability," in *The Foundations of Mathematics and Other Logical Essays*, London, Kegan Paul, and New York, Harcourt, Brace and Co., 1931.

[31] Royden, H. L. (1988), *Real Analysis* (3rd Ed), Prentice Hall, Englewood Cliffs, New Jersey.

[32] Rubinstein, Ariel (1982), "Perfect Equilibrium in a Bargaining Model," *Econometrica*, 50 (1), pp. 97-109.

[33] Rubinstein, Ariel (1987), "A Sequential Strategic Theory of Bargaining," *Advances in Economic Theory: Fifth World Congress*, 5, edited by Truman Fassett Bewley, Chapter 5, Cambridge University Press, Cambridge, England.

[34] Savage, L. (1954), *The Foundations of Statistics*, New York: Wiley.

[35] Schumpeter, Joseph A. (1934), *The Theory of Economic Development, An Inquiry into Profits, Capital, Credit, Interest, and the Business Cycle*, Harvard University Press, Cambridge, Massachusetts.

[36] Sobel, Joel and Ichiro Takahashi (1983), "A Multistage Model of Bargaining," *Review of Economic Studies*, 50 (3), pp. 411-426.

[37] Stokey, Nacy L. and Robert E. Lucas with Edward C. Prescott (1989), *Recursive Methods in Economic Dynamics*, Harvard University Press, Cambridge, Massachusetts.

[38] Summers, Laurence H. (2014), "Reflections on the 'New Secular Stagnation Hypothesis'," *Secular Stagnation: Facts, Causes and Cures* (Teulings, Coen and Richard Baldwin, ed.), A VoxEU.org Book.

[39] Sutton, John (1986), "Non-Cooperative Bargaining Theory: An Introduction," *Review of Economic Studies*, 53 (5), pp. 709-724.

[40] Teulings, Coen and Richard Baldwin, ed., (2014), *Secular Stagnation: Facts, Causes and Cures*, A VoxEU.org eBook.

A Appendix

A.1 Proof of Proposition 1

For any (well defined) posterior $\mu_t(\cdot|h^t;\mu_0,s)\in\mathcal{M}_t(h^t;s)$, it follows from (11) that

$$\mu_t\left(\mathcal{I}\left(h^t;s\right)|h^t;\mu_0,s\right)=1.$$

Since $\mathcal{M}_0$ includes all plausible probability distributions over $\Sigma(\Theta)$—i.e., $\mathcal{M}_0$ is sufficiently "large" (inclusive)—we must have $\mu_t(\cdot|h^t;\mu_0,s)\in\mathcal{M}_0$. In conjunction, we must also have

$$\mu_t\left(\cdot|h^t;\mu_0,s\right)\in\left\{\mu_0\in\mathcal{M}_0|\mu_0\left(\mathcal{I}\left(h^t\right)\right)=1\right\},$$

which implies

$$\mathcal{M}_t\left(h^t;s\right)\subseteq\left\{\mu_0\in\mathcal{M}_0|\mu_0\left(\mathcal{I}\left(h^t\right)\right)=1\right\}. \tag{37}$$

Conversely, for $\mu\in\{\mu_0\in\mathcal{M}_0|\mu_0(\mathcal{I}(h^t;s))=1\}$, Bayes' Rule (11) implies

$$\mu_t\left(\mathcal{A}|h^t;\mu,s\right)=\mu\left(\mathcal{A}\cap\mathcal{I}\left(h^t;s\right)\right)=\mu\left(\mathcal{A}\right)-\mu\left(\mathcal{A}\cap\left(\Theta\backslash\mathcal{I}\left(h^t;s\right)\right)\right)=\mu\left(\mathcal{A}\right)$$

for all $\mathcal{A}\in\Sigma(\Theta)$, where the last equality follows from $\mu(\mathcal{A}\cap(\Theta\backslash\mathcal{I}(h^t;s)))\leq\mu(\Theta\backslash\mathcal{I}(h^t;s))=1-\mu(\mathcal{I}(h^t;s))=0$. Since the posterior equals the prior, we must have

$$\mu\in\mathcal{M}_t\left(h^t;s\right),$$

which implies

$$\left\{\mu_0\in\mathcal{M}_0|\mu_0\left(\mathcal{I}\left(h^t;s\right)\right)=1\right\}\subseteq\mathcal{M}_t\left(h^t;s\right). \tag{38}$$

It follows from (37) and (38) we must have

$$\mathcal{M}_t\left(h^t;s\right)=\left\{\mu_0\in\mathcal{M}_0|\mu_0\left(\mathcal{I}\left(h^t;s\right)\right)=1\right\}.$$

■

A.2 Proof of Proposition 2

This is a game of perfect information. To prove the strategy profile given by (19) and (20) constitutes a SPNE, it suffices to prove that each player's strategy is sequentially rational. Notice that given the seller's strategy $w(\theta)=\theta\Pi_B$, the maximal possible payoff the buyer can get in period t is 0 for all $t=0,1,\ldots$. This can already be achieved by the buyer's strategy

$$k\left(w_t;\theta\right)=\begin{cases}1 & \text{if } w_t\leq\theta\Pi_B,\\ 0 & \text{otherwise;}\end{cases}$$

if the seller deviates and offers $w_t\neq\theta\Pi_B$, it is optimal for the buyer to accept iff $w_t\leq\theta\Pi_B$, that is, iff the buyer has non-negative surplus from current trade and the payoff is at least as good as buying in the next period at price $\theta\Pi_B$, which gives the buyer a surplus of zero—this is exactly what the buyer strategy entails. Since buyer strategy in question is optimal for both $w_t=\theta\Pi_B$ and $w_t\neq\theta\Pi_B$, it must be sequentially rational. Similarly, given the buyer's strategy in question, the maximal possible payoff the seller can get in period t is

$$\theta\Pi_S + \theta\Pi_B - a = (\theta - \theta^{**})(\Pi_S + \Pi_B) > 0.$$

Note, this can already be achieved by the seller's strategy in question, which must also be sequentially rational.

What remains to be shown is that the strategy profile in question constitutes the unique SPNE for the given $\theta \in (\theta^{**}, 1]$. The key of the proof is to show that the full set of prices offered by any SPNE seller strategy in any subgame is a singleton—which only contains $w(\theta) = \theta\Pi_B$. Let $w^{\inf}$ denote the infimum of the full set of prices offered by the seller according some SPNE equilibrium strategy in some subgame. Since we have found a SPNE for this game and also know the prices offered in any subgame of any SPNE is bounded, so $w^{\inf}$—as the infimum of a non-empty bounded set—is well defined. Let $w(h^t;\theta)$ denote generally the offer made after any history h^t in any SPNE, then the definition of $w^{\inf}$ entails $w(h^t;\theta) \geq w^{\inf}$. Also, if the buyer rejects $w(h^t;\theta)$, the best price he can expect to get in the next period is no less than $w^{\inf}$—which implies the present value of payoff is no more than $\frac{\theta\Pi_B - w^{\inf}}{1+r}$, i.e., $\theta\Pi_B - w(h^t;\theta) \leq \frac{\theta\Pi_B - w^{\inf}}{1+r}$. Since this inequality holds generically, it must hold for the limit $w(h^t;\theta) \to w^{\inf}$—that is, $\theta\Pi_B - w^{\inf} \leq \frac{\theta\Pi_B - w^{\inf}}{1+r}$—which implies $w^{\inf} \geq \theta\Pi_B$ (for $r > 0$) and hence $w(h^t;\theta) \geq \theta\Pi_B$. As a result in any SPNE, we must have $k(h^t;\theta) = 1$ if $w_t < \theta\Pi_B$—that is, any (negative deviation from the infimum) $w_t < \theta\Pi_B$ would be accepted immediately. Since any offer $w_t > \theta\Pi_B$ implies strictly negative payoff for the buyer it must be rejected. Therefore the highest price that can be accepted is $\theta\Pi_B$. So far we have established that if $w(h^t;\theta)$ is a SPNE strategy of the seller, which is acceptable by the buyer's strategy in a SPNE, then we must have $w(h^t;\theta) = \theta\Pi_B$—which gives the seller a surplus of $(\theta - \theta^{**})(\Pi_S + \Pi_B) > 0$.

We are still to prove that in any SPNE, the offer $w(h^t;\theta)$ must be acceptable by the buyer's strategy in that SPNE—implying that $w(h^t;\theta) = \theta\Pi_B$. Otherwise, we must suppose that $w(h^t;\theta) > \theta\Pi_B$ (finitely) for some SPNE after some history h^t. Since $w(h^t;\theta)$ will be rejected, the maximal discounted payoff the seller can get in future trade is $\frac{(\theta-\theta^{**})(\Pi_S+\Pi_B)}{1+r}$, which is strictly smaller than $[(\theta - \theta^{**})(\Pi_S + \Pi_B) - \varepsilon]$ for some $\varepsilon \in \left(0, \frac{r(\theta-\theta^{**})(\Pi_S+\Pi_B)}{1+r}\right)$—the payoff the seller can get if she offers $w_t = \theta\Pi_B - \varepsilon$. Notice that w_t—which is below the infimum—is acceptable by the buyer's strategy in all SPNE. So we derive a contradiction from the supposition that $w(h^t;\theta) > \theta\Pi_B$ (finitely)—the contradiction implies the negation of the supposition—we must have $w(h^t;\theta) = \theta\Pi_B$, which is identical to the strategy described by (19) in Proposition 2**.**

Last we prove that in any SPNE the buyer's strategy must entail $k(h^t;\theta) = 1$ if $w_t = \theta\Pi_B$. Otherwise, we have to suppose that there exists a SPNE such that the seller offers $w(h^t;\theta) = \theta\Pi_B$ and buyer's strategy is $k(h^t;\theta) = 0$ if $w_t = \theta\Pi_B$—this strategy profile obviously results in a rejection of $w(h^t;\theta)$ in period t. After that, the seller's discounted payoff can not exceed $\frac{(\theta-\theta^{**})(\Pi_S+\Pi_B)}{1+r}$. Instead, if the seller deviates to the price $w_t = \theta\Pi_B - \varepsilon$ for some $\varepsilon \in \left(0, \frac{r(\theta(\Pi_S+\Pi_B)-a)}{1+r}\right)$, then the payoff for the seller will be $[(\theta - \theta^{**})(\Pi_S + \Pi_B) - \varepsilon] > \frac{(\theta-\theta^{**})(\Pi_S+\Pi_B)}{1+r}$—which is a profitable deviation for the seller. This contradicts the supposition. Consequently, the buyer's strategy must must be $k(h^t;\theta) = 1$ if $w_t = \theta\Pi_B$. Overall, the buyer's strategy must be $k(h^t;\theta) = 1$ iff $w_t \leq \theta\Pi_B$, which is identical to strategy described by (20) in Proposition 2**.**

It is straightforward to verify that the unique SPNE outcome is that trade occurs on the equilibrium path. Since the social surplus is maximized, it is impossible to make either player better off without making the other one worse off. Hence, Pareto efficiency is obtained.■

A.3 Proof of Theorem 1

To see that the buyer's cut-off strategy $s_B^{**}(h^t)$—as specified by (23)—is sequentially rational, note that given the seller's constant strategy w^{**}—as specified in (21))—the maximal possible payoff (without discounting) for the buyer is

$$\begin{cases} (\theta - \theta^{**})\Pi_B > 0 & \text{for } \theta \in (\theta^{**}, 1], \\ 0 & \text{for } \theta \in [0, \theta^{**}], \end{cases}$$

which can already be achieved by the buyer strategy $s_B^{**}(h^t)$; if the seller deviates and offers $w_t \neq w^{**}$, it is optimal for the buyer to accept iff

$$\theta\Pi_B - w_t \geq \max\left\{0, \frac{\theta\Pi_B - w^{**}}{1+r}\right\},$$

that is, iff the buyer has non-negative surplus from current trade and the payoff is at least as good as buying in the next period at price w^{**}—this is exactly what the buyer strategy $s_B^{**}(h^t)$ entails. Since $s_B^{**}(h^t)$ is optimal for both $w_t = w^{**}$ and $w_t \neq w^{**}$, it must be sequentially rational.

Conversely, we will show that the constant seller strategy $s_S^{**}(h^t) = w^{**}$ is sequentially rational—i.e., given the buyer's cut-off strategy $s_B^{**}(h^t)$, it passes the max-inf criterion[36] and is not weakly dominated conditional on any plausible set of posteriors $\mathcal{M}_t(h^t; s^{**})$.

First, notice that the cut-off type for $w_t = w^{**}$ is $\hat{\theta}(w^{**}) = \theta^{**}$, who brings payoff $\theta^{**}\Pi_S + w^{**} - a = 0$ to the seller, which is the infimum of seller payoff for the constant strategy $s_S^{**}(h^t) = w^{**}$. For any alternative seller strategy $w(h^t)$, consider buyer type $\theta \in [0, \theta^{**})$ and notice that the seller's surplus from trading with type $\theta \in [0, \theta^{**})$ is strictly negative since the set $\{w_t - a + \theta\Pi_S | -\Pi_S \leq w_t \leq \theta\Pi_B, \theta \in [0, \theta^{**})\} = [-\Pi_S - a, 0)$ is an interval in the negative region. Given the buyer strategy $s_B^{**}(h^t)$, a buyer never makes a loss from trade, therefore the infimum of the seller's payoff when dealing with type $\theta \in [0, \theta^{**})$ is zero at best. That means the constant strategy w^{**} satisfies the max-inf criterion.

Second, we prove that constant strategy w^{**} is not weakly dominated conditional on any set of posteriors $\mathcal{M}_t(h_t; s^{**})$. Suppose the opposite—that is, there exists an alternative (pure) strategy $w(h^t)$ and a history $h^{t'}$ and the associated set of posteriors $\mathcal{M}_{t'}\left(h^{t'}; s^{**}\right)$ such that the continuation strategy of $w(h^t)$ following history $h^{t'}$ weakly dominates w^{**} conditional on $\mathcal{M}_{t'}\left(h^{t'}; s^{**}\right)$.

The seller's belief updating following history $h^{t'}$ can be exhaustively categorized into the following two cases: (1) $\mathcal{M}_{t'}\left(h^{t'}; s^{**}\right) = \mathcal{M}_0$, which occurs if $t' = 0$ or $\mathcal{I}\left(h^{t'}; s^{**}\right) = \varnothing$, or (2) $\mathcal{M}_{t'}\left(h^{t'}; s^{**}\right) = \{\mu \in \mathcal{M}_0 | \mu([0, \theta^{**})) = 1\}$, which occurs if $\mathcal{I}\left(h^{t'}; s^{**}\right) = [0, \theta^{**}) \neq \varnothing$—in detail, for $t' = 1, 2, \ldots$, the offers are $w_\tau = w^{**}$ for $\tau = 0, 1, \ldots, t'$ and are always rejected, then the seller can infer that $\theta \in [0, \theta^{**})$.

In case (2), evidently w^{**} is optimal since—given the buyer strategy $s_B^{**}(h^t)$—it will not induce any type $\theta \in [0, \theta^{**})$ to buy (by lowering the price). So it is trivial to show that w^{**} cannot be weakly dominated by $w(h^t)$ conditional on case (2). Therefore one has to suppose that w^{**} is weakly dominated by $w(h^t)$ conditional on case (1) $\mathcal{M}_{t'}\left(h^{t'}; s^{**}\right) = \mathcal{M}_0$.

Obviously, if $w\left(h^{t'}\right) = w^{**}$, then a buyer of all types $\theta \in [\theta^{**}, 1]$ will buy at price w^{**} and all types $\theta' \in [0, \theta^{**})$ will not buy at price w^{**}. Since it is suboptimal for the seller to induce any type $\theta' \in [0, \theta^{**})$ to buy, it is impossible for $w(h^t)$ to weakly dominate w^{**} if $w\left(h^{t'}\right) \leq w^{**}$, we have to infer that $w\left(h^{t'}\right) > w^{**}$.

It follows that all types $\theta \in \left[\theta^{**}, \frac{w\left(h^{t'}\right)}{\Pi_B}\right)$ will not buy at price $w\left(h^{t'}\right)$ because of negative payoff. If there exists some $\theta' \in \left(\theta^{**}, \frac{w\left(h^{t'}\right)}{\Pi_B}\right)$ such that the subsequent strategy of $w(h^t)$ fails to eventually induce θ' to buy, then $w(h^t)$ is inferior to w^{**} conditional on the Dirac measure $\delta_{\theta'} \in \mathcal{M}_{t'}\left(h^{t'}\right) = \mathcal{M}_0$, therefore the former does not weakly dominate the latter. Consequently, in order for $w(h^t)$ to weakly dominate w^{**} conditional on $\mathcal{M}_{t'}\left(h^{t'}; s^{**}\right) = \mathcal{M}_0$, for any type $\theta_n := \theta^{**} + \frac{\frac{w\left(h^{t'}\right)}{\Pi_B} - \theta^{**}}{n+1} \in \left(\theta^{**}, \frac{w\left(h^{t'}\right)}{\Pi_B}\right)$, $n = 1, 2, \ldots$, there must exist a period $t_n \geq t' + 1$ such that the offer $w(h^{t_n})$ induces type θ_n to buy in period t_n. Necessarily

[36] For brevity and without loss of clarity, in the remainder of this paper we write "max-inf criterion" in place of "max-inf expected utility criterion".

$w\left(h^{t_n}\right) \leq \theta_n \Pi_B$. Notice that θ_n is decreasing in n. Since the buyer strategy $s_B^{**}\left(h^t\right)$ is a cut-off strategy, an offer that is accepted according to $s_B^{**}\left(h^t\right)$ by type θ_n must also be accepted by type $\theta_{n'} > \theta_n$. Then it is obvious that $t_n \geq t_{n'}$ if $\theta_n < \theta_{n'}$—or, equivalently, $n > n'$—implying that the lower type θ_n must not buy the good earlier than the higher type $\theta_{n'}$.

For $w\left(h^t\right)$ not to be inferior to w^{**} conditional on Dirac measure θ_n it must be at least equally profitable for the seller to sell to type θ_n at price $w\left(h^{t_n}\right)$ in period t_n as to sell at price w^{**} in period t'. As an implication the following inequalities must hold:

$$w^{**} + \theta_n \Pi_S - a \leq \frac{w\left(h^{t_n}\right) + \theta_n \Pi_S - a}{(1+r)^{t_n - t'}} \leq \frac{\theta_n \Pi_B + \theta_n \Pi_S - a}{(1+r)^{t_n - t'}},$$

where the second inequality is implied by the necessary condition: $w\left(h^{t_n}\right) \leq \theta_n \Pi_B$—for type θ_n to accept $w\left(h^{t_n}\right)$. It follows that

$$(1+r)^{t_n - t'} \leq \frac{\theta_n \Pi_B + \theta_n \Pi_S - a}{w^{**} + \theta_n \Pi_S - a} = \frac{\Pi_B + \Pi_S}{\Pi_S}.$$

Since $r > 0$, we must have

$$t_n - t' \leq \frac{\ln\left(\frac{\Pi_B + \Pi_S}{\Pi_S}\right)}{\ln(1+r)} < \infty \text{ for all } n = 1, 2, \ldots, \tag{39}$$

that is, for all $n = 1, 2, \ldots$, $(t_n - t')$ is uniformly bounded from above. Without loss of generality, let $\bar{T} := \max\left\{t' + i | i \leq \frac{\ln \frac{\Pi_S + \Pi_B}{\Pi_S}}{\ln(1+r)} \text{ for all } i = 0, 1, \ldots,\right\}$ be the deadline to conclude trade with all types $\theta_n = \theta_1, \theta_2, \ldots$. It must follow that there exists some period $T \in \left\{t', \ldots, \bar{T}\right\}$ such that $w\left(h^T\right) \leq w^{**}$ for some h^T and $w\left(h^\tau\right) > w^{**}$ for all h^τ, $\tau = t', \ldots, T-1$. Notice that T is defined as the earliest period in which the price will reach or cross w^{**} from above according to strategy $w\left(h^t\right)$ if the game has not ended. Let $\underline{w} := \min\left\{w\left(h^\tau\right) | \tau = t', \ldots, T-1, \forall \text{ plausible } h^\tau\right\}$ be the lowest price that strategy $w\left(h^\tau\right)$ reaches before period T—which is the minimum of a non-empty finite set, and must exist. Therefore all types θ_n such that $\theta_n \in \left(\theta^{**}, \frac{\underline{w}}{\Pi_B}\right)$ buy at price $w\left(h^T\right) \leq w^{**}$ in period $T > t'$—in response to the deviation strategy $w\left(h^t\right)$ and according to buyer strategy $s_B^{**}\left(h^t\right)$—which is based on the expectation that the future prices will be constantly w^{**}; while they buy at price w^{**} in period t' in response to the constant strategy w^{**}. Clearly, for each type $\theta_n \in \left(\theta^{**}, \frac{\underline{w}}{\Pi_B}\right)$, the strategy $w\left(h^t\right)$ is inferior to the constant strategy w^{**} conditional on the Dirac measure δ_{θ_n}—because the former sells at no higher price $\left(w\left(h^T\right) \leq w^{**}\right)$ to type θ_n and with a delay—which contradicts the presupposition that the former weakly dominates the latter conditional on $\mathcal{M}_{t'}\left(h^{t'}\right)$. In conclusion, the constant strategy w^{**} is not weakly dominated conditional on any $\mathcal{M}_t\left(h^t\right)$.

The Markov perfection property of strategy profile s^{**} is confirmed by the fact that the list of state variables—i.e., $\varnothing$ for the seller, and $\{\theta, w_t\}$ for the buyer—cannot be reduced further.

We now derive the uniqueness of MPOE. Notice that to be an MPOE, the seller's strategy must be constant—let it be denoted by w. If $0 \leq w < w^{**}$, then it incentivizes type $\theta \in \left[\frac{w}{\Pi_B}, \theta^{**}\right)$ to buy and leaves the seller with negative payoff and fails the max-inf criterion; if $w^{**} < w \leq \Pi_B$, the buyer's sequentially rational strategy must be

$$k\left(w_t; \theta\right) = \begin{cases} 1 & \text{if } \theta \Pi_B - w_t \geq \max\left\{0, \frac{1}{1+r}\left(\theta \Pi_B - w\right)\right\}, \\ 0 & \text{otherwise.} \end{cases}$$

Given this as the buyer's strategy, the constant strategy w cannot be sequentially rational—because following a rejection of w, the seller will infer that the true buyer type is $\theta \in \left[0, \frac{w}{\Pi_B}\right)$ and given this belief the constant strategy w^{**} weakly dominates w. To see why? Note that the price offer w means no trade ex post; with the price the offer w^{**}, non-negative seller payoff is guaranteed, and the seller gets positive surplus from trade conditional on priors δ_θ such that $\theta \in \left(\theta^{**}, \frac{w}{\Pi_B}\right)$. Consequently, for MPOE we must have $w = w^{**}$ and the buyer's strategy as given by $s_B^{**}\left(h^t\right)$, and hence the uniqueness of the MPOE.

The rest of the proof is straightforward.■

A.4 An Example of Weakly Dominated Strategy which Satisfies the max-inf Expected Utility Criterion

Consider a one-period model with the (artificial) buyer type space: $[0,1)$. Suppose the buyer plays the following strategy:

$$k\left(w;\theta\right) = \begin{cases} 1 & \text{if } \theta \geq \max\left\{\frac{w}{\Pi_B}, \hat{\theta}\left(w\right)\right\}, \\ 0 & \text{otherwise,} \end{cases}$$

where w is the price on offer; $\hat{\theta}\left(w\right) = \frac{(1+r)w - w^{**}}{r\Pi_B}$ is the cut-off type—for our purpose of looking at the seller choice under ambiguity, it does not matter whether the buyer's strategy is optimal or not. We argue that the following strategy of the seller—which is to offer the price $\hat{w} = \Pi_B$—is a weakly dominated strategy which satisfies the max-inf criterion. To see this, first notice that for all types $\theta \in [0, \theta^{**})$, the seller's surplus from trade is negative since the set $\{w - a + \theta\Pi_S | -\Pi_S \leq w \leq \theta\Pi_B, \theta \in [0, \theta^{**})\} = [-\Pi_S - a, 0)$. Therefore the infimum of seller surplus for every possible strategy $w \in [0, \Pi_B]$ is zero at best. Second, notice that under strategy $\hat{w} = \Pi_B$ there will be no trade, so the seller surplus is always zero—which is the maximum infimum. Third, we argue that strategy $\hat{w} = \Pi_B$ is weakly dominated by the strategy $w^{**} := \theta^{**}\Pi_B$, conditional on the ambiguous beliefs $\{\tilde{\mu} \in \mathcal{M}_0 | \tilde{\mu}\left([0,1)\right) = 1\}$.

The argument is as follows: For all type $\theta \in [0, \theta^{**}]$ both strategies give the same surplus for the seller—which is zero and also the maximum infimum of seller surplus. For all $\theta \in (\theta^{**}, 1)$, strategy w^{**} gives strictly positive and hence larger seller surplus. Therefore, the strategy $\hat{w} = \Pi_B$ gives strictly lower expected value of seller surplus than strategy w^{**} for all beliefs $\mu' \in \{\mu \in \mathcal{M}_0 | \mu\left([0,1)\right) = 1, \mu\left((\theta^{**}, 1)\right) > 0\}$; while they are equivalent for all beliefs $\nu' \in \{\nu \in \mathcal{M}_0 | \nu\left([0,1)\right) = 1, \nu\left((\theta^{**}, 1)\right) = 0\}$. Furthermore, for all beliefs $\tilde{\mu}' \in \{\tilde{\mu} \in \mathcal{M}_0 | \tilde{\mu}\left([0,1)\right) = 1\}$ strategy w^{**} gives the seller either equal or higher surplus in comparison with strategy $\hat{w} = \Pi_B$. Thus, strategy $\hat{w} = \Pi_B$ is weakly dominated by strategy w^{**} for the ambiguous beliefs $\{\tilde{\mu} \in \mathcal{M}_0 | \tilde{\mu}\left([0,1)\right) = 1\}$.

A.5 Proof of Proposition 3

The proof of existence will be by construction of a POE for each $w_0 \in [w^{**}, \bar{w}]$. For the limiting case: $w_0 = w^{**}$, Theorem 1 has already established that the MPOE $\left(s^{**}, \mathcal{M}\left(s^{**}\right)\right)$—which has initial price offer $w_0 = w^{**}$—constitutes a POE. So we only need to deal with the case $w_0 \in (w^{**}, \bar{w}]$.

Denote the POE to be constructed by $\left(s, \mathcal{M}\left(s\right)\right)$. Let the seller's strategy in s be

$$s_S\left(h^t\right) = w\left(t\right) = \begin{cases} w_0 & \text{for } t = 0, \\ w^{**} & \text{otherwise.} \end{cases} \tag{40}$$

Let $\hat{\theta}\left(w_t\right)$ be the cut-off type of buyer who is indifferent between buying at price $w_t \geq w^{**}$ in the current period t and buying at price w^{**} in the next period. Notice that it suffices to only compare the current and next period payoffs because the price is expected to be (constant) w^{**} in all future periods—so buying in any further future period than the next is inferior to buying in the next period because of the delay for all types $\theta \in (\theta^{**}, 1]$. The indifference condition is given by

$$\Pi_B\hat{\theta}\left(w_t\right) - w_t = \frac{\Pi_B\hat{\theta}\left(w_t\right) - w^{**}}{1+r}, \tag{41}$$

which has the solution

$$\hat{\theta}(w_t) = \frac{(1+r)\,w_t - w^{**}}{r\Pi_B}. \tag{42}$$

It is easy to verify the following monotonic relation between buyer surplus and type—conditional on $\theta \geq \frac{w_t}{\Pi_B}$:

$$\Pi_B\theta - w_t \gtreqless \frac{\Pi_B\theta - w^{**}}{1+r} \text{ iff } \theta \gtreqless \hat{\theta}(w_t),$$

which implies that it is a best response to the seller's strategy $w(t)$ for the buyer to play the following cut-off strategy:

$$k(w_t;\theta) = \begin{cases} 1 & \text{if } \theta \geq \max\left\{\frac{w_t}{\Pi_B}, \hat{\theta}(w_t)\right\}, \\ 0 & \text{otherwise.} \end{cases} \tag{43}$$

To prove that the seller's strategy $w(t)$ is sequentially rational, we first show that given the buyer's strategy $k(w_t;\theta)$, the seller's strategy $w(t)$ passes the max-inf criterion. To see this, notice that

$$w(t) \geq w^{**} \text{ for all } t = 0, 1, \ldots,$$

$$\hat{\theta}(w_t) = \begin{cases} \frac{(1+r)w_t - w^{**}}{r\Pi_B} > \theta^{**} & \text{for } w_t > w^{**}, \\ \theta^{**} & \text{for } w_t = w^{**}. \end{cases}$$

It follows that given the buyer's strategy, the infimum of payoff for the seller is 0 under strategy $w(t)$—since it will be rejected by all types $\theta \in [0, \theta^{**})$—yielding payoff 0 as the lowest. Faced by possible types $\theta \in [0, \theta^{**})$ no alternative strategies can yield positive payoff for the seller, therefore the infimum of their payoffs cannot be positive. Thus, strategy $w(t)$ passes the max-inf criterion.

Next we establish that the seller's strategy $w(t)$ is not weakly dominated conditional on any beliefs $\mathcal{M}_t(h^t;s)$. Suppose the opposite—that is, there exist some t', history $h^{t'}$ and beliefs $\mathcal{M}_{t'}\left(h^{t'};s\right)$ and an alternative seller strategy $w(h^t)$ such that $w(h^t)$ weakly dominates $w(t)$ conditional on $\mathcal{M}_t\left(h^{t'};s\right)$.

The seller's belief updating following history $h^{t'}$ can be exhaustively categorized into the following four cases: (1) $\mathcal{I}\left(h^{t'};s\right) = [0,1]$ with $t' = 0$, or (2) $\mathcal{I}\left(h^{t'};s\right) = \left[0, \hat{\theta}(w_0)\right)$ with $t^{'} = 1$ and $\mathcal{I}\left(h^{t'};s\right) \neq \varnothing$, or (3) $\mathcal{I}\left(h^{t'};s\right) = [0, \theta^{**})$ with $t^{'} \geq 2$ and $\mathcal{I}\left(h^{t'};s\right) \neq \varnothing$, or (4) $\mathcal{I}\left(h^{t'};s\right) = \varnothing$ with $t^{'} \geq 1$. Notice that cases (2) and (3) occur after the seller has (always) played according to strategy $w(t)$ up to time $t' \geq 1$ and no offer has been accepted. In these cases

$$\mathcal{M}_{t'}\left(h^{t'};s\right) = \left\{\mu \in \mathcal{M}_0 | \mu\left(\left[0, \hat{\theta}(w(t'-1))\right)\right) = 1\right\}.$$

Case (4) occurs after the sellers knows by time t' that the observed history $h^{t'}$ contradicts the equilibrium path of s. In this case, we have $t^{'} \geq 1$, $\mathcal{I}\left(h^{t'};s\right) = \varnothing$ and $\mathcal{M}_{t'}\left(h^{t'};s\right) = \mathcal{M}_0$.

For cases (2), (3) and (4), arguments that are similar to the proof of Theorem 1 can be applied here to show it is logically impossible for the subsequent strategies of $w(t)$—which become the constant strategy w^{**} from $t^{'} \geq 1$—to be weakly dominated by (the subsequent strategies of) $w(h^t)$. So what remains to be dealt with is case (1), which is at the beginning of the game $t' = 0$ and $\mathcal{M}_{t'}\left(h^{t'};s\right) = \mathcal{M}_0$.

For case (1), the initial offer of strategy $w(t)$ is $w_0 > w^{**}$, which is uniquely optimal if the buyer is of type $\hat{\theta}(w_0) \in (\theta^{**}, 1]$. Notice that from (42) it follows that $\hat{\theta}(w_0) = \frac{(1+r)w_0 - w^{**}}{r\Pi_B}$. To satisfy the condition $\hat{\theta}(w_0) \in (\theta^{**}, 1]$, we must have $w^{**} < w_0 \leq w^{**} + \frac{r(\Pi_B - w^{**})}{1+r} := \bar{w}$, which is satisfied by assumption. In order for $w(h^t)$ to be at least as good as $w(t)$ for type $\hat{\theta}(w_0)$, we must have $w\left(h^0\right) = w_0$.

Subsequently, if $w(h^t) = w^{**}$ for $t = 1$, then all types $\theta \in [\theta^{**}, 1]$ will have bought the good within two periods and all the types $\theta \in [0, \theta^{**})$ will not and should never be incentivized to buy in the interest of seller—and in this case it is impossible for $w(h^t)$ to weakly dominate $w(t)$ since the latter does not incentivize any type $\theta \in [0, \theta^{**})$ to buy. If $w(h^1) < w^{**}$, the strategy will have negative payoff for the seller facing buyer type $\left[\hat{\theta}\left(w\left(h^1\right)\right), \theta^{**}\right)$, therefore cannot weakly dominate $w(t)$—note, $w(t)$ has the payoff of 0 for the seller facing buyer type $\left[\hat{\theta}\left(w\left(h^1\right)\right), \theta^{**}\right)$. Then we have to infer that $w(h^1) > w^{**}$ for $t = 1$.

It follows that all types $\theta \in \left(\theta^{**}, \min\left(\frac{w(h^0)}{\Pi_B}, \frac{w(h^1)}{\Pi_B}\right)\right)$ will not have bought the good according to strategy $s_B(h^t)$ in response to offers $w(h^t)$ within two periods—while they will in response to strategy $w(t)$. If there exists some $\theta' \in \left(\theta^{**}, \min\left(\frac{w(h^0)}{\Pi_B}, \frac{w(h^1)}{\Pi_B}\right)\right)$ such that the subsequent strategy of $w(h^t)$ fails to eventually incentivize θ' to buy, then $w(h^t)$ is inferior to $w(t)$ conditional on the Dirac measure $\delta_{\theta'} \in \mathcal{M}_{t'}\left(h^{t'}; s\right) = \mathcal{M}_0$—notice $t' = 0$ for case (1)—and therefore $w(h^t)$ does not weakly dominate $w(t)$. Consequently, in order for $w(h^t)$ to weakly dominate $w(t)$ conditional on $\mathcal{M}_{t'}\left(h^{t'}\right) = \mathcal{M}_0$, for any type $\theta_n := \theta^{**} + \frac{\min\left(\frac{w(h^0)}{\Pi_B}, \frac{w(h^1)}{\Pi_B}\right) - \theta^{**}}{n+1} \in \left(\theta^{**}, \min\left(\frac{w(h^0)}{\Pi_B}, \frac{w(h^1)}{\Pi_B}\right)\right)$, $n = 1, 2, \ldots$, there must exist a period $t_n \geq 1$ such that $w(h^{t_n})$ incentivizes type θ_n to buy in period t_n. Then arguments that are similar to the proof of Theorem 1 can be applied here to show that there exists an $\underline{w} \in \left(w^{**}, \min\left(\frac{w(h^0)}{\Pi_B}, \frac{w(h^1)}{\Pi_B}\right)\right)$ and $1 < T' < \infty$ such that all types θ_n such that $\theta_n \in \left(\theta^{**}, \frac{\underline{w}}{\Pi_B}\right)$ buy at price $w\left(h^{T'}\right) \leq w^{**}$ in period $T' > 1$ in response to the deviation strategy $w(h^t)$; while they buy at price w^{**} in period 1 in response to the strategy $w(t)$. Clearly, for each type θ_n such that $\theta_n \in \left(\theta^{**}, \frac{\underline{w}}{\Pi_B}\right)$, the strategy $w(h^t)$ is inferior to the strategy $w(t)$ conditional on the Dirac measure δ_{θ_n} because the former sells at no higher price to type θ_n and with a delay—which contradicts the presupposition that the former weakly dominates the latter conditional on $\mathcal{M}_{t'}\left(h^{t'}; s\right)$. In conclusion, the strategy $w(t)$ is not weakly dominated conditional on any $\mathcal{M}_t(h^t; s)$.■

A.6 Proof of Proposition 4

Using the envelope theorem, we can derive $W'(q_t^{\sup})$—evaluated under equation: $\frac{r\Pi_B q_{t+1}^{\sup} + w\left(q_{t+1}^{\sup}\right)}{1+r} = w(q_t^{\sup})$—as:

$$W'(q_t^{\sup}) = f(q_t^{\sup})\left[w(q_t^{\sup}) - a\right],$$

where $f(\theta)$ is the probability density function. The first-order condition (FOC) for the maximization problem—evaluated under equation: $\frac{r\Pi_B q_{t+1}^{\sup} + w\left(q_{t+1}^{\sup}\right)}{1+r} = w(q_t^{\sup})$ and substituting $W'\left(q_{t+1}^{\sup}\right)$—is

$$\left[F(q_t^{\sup}) - F\left(q_{t+1}^{\sup}\right)\right]\frac{r\Pi_B + w'\left(q_{t+1}^{\sup}\right)}{1+r} - f\left(q_{t+1}^{\sup}\right)\left\{w(q_t^{\sup}) - a - \frac{1}{1+r}\left[w\left(q_{t+1}^{\sup}\right) - a\right]\right\} = 0. \tag{44}$$

Since the seller's prior is a uniform distribution, we have $F(\theta) = \theta$ and $f(\theta) = 1$, the FOC gives

$$\left(q_t^{\sup} - q_{t+1}^{\sup}\right)\frac{r\Pi_B + w'\left(q_{t+1}^{\sup}\right)}{1+r} - \left\{w(q_t^{\sup}) - a - \frac{1}{1+r}\left[w\left(q_{t+1}^{\sup}\right) - a\right]\right\} = 0. \tag{45}$$

We consider a plausible MPBE such that for all $q_t^{\sup} \in [\theta^{**}, 1]$ the function w has the following linear form:

$$w(q_t^{\sup}) = \beta\Pi_B(q_t^{\sup} - \theta^{**}) + a, \tag{46}$$

where $\beta \in [0, 1)$ is a parameter to be determined. It follows from equations (27) and (46) the cut-off function

$\hat{\theta}$ takes the following linear form:

$$\hat{\theta}(w_t) = \frac{1+r}{(r+\beta)\Pi_B}(w_t - a) + \theta^{**}. \tag{47}$$

Given the value of state variable $q_t^{\sup}$, if the seller follows her putative MPBE strategy in question we have $w_t = w(q_t^{\sup})$. Conditional on this, equations (46), (47) and $q_{t+1}^{\sup} = \hat{\theta}(w_t)$ imply that

$$\frac{q_{t+1}^{\sup} - \theta^{**}}{q_t^{\sup} - \theta^{**}} = \frac{(1+r)\beta}{r+\beta}. \tag{48}$$

Then from (45), (46) and (48) we can derive for $q_t^{\sup} \geq \theta^{**}$ the following quadratic equation of β:

$$\beta^2 + 2r\beta - r = 0.$$

which has the following unique positive solution:

$$\beta = -r + \sqrt{r^2 + r}. \tag{49}$$

We thus have the seller's strategy $w(q_t^{\sup})$ as specified by (30) and (31). It is easy to verify that the buyer's best response to this strategy of the seller is characterized by (32) and (33).

From equations (46), (48) and the fact that $w^{**} \to a$, it follows that for $q_t^{\sup} \in [\theta^{**}, 1]$ and $\tau = t, t+1, \ldots,$

$$w(q_t^{\sup}) - a = \beta\Pi_B(q_t^{\sup} - \theta^{**}),$$

$$q_\tau^{\sup} - \theta^{**} = \left[\frac{(1+r)\beta}{r+\beta}\right]^{\tau-t}(q_t^{\sup} - \theta^{**}),$$

and

$$w(q_\tau^{\sup}) - a = \left[\frac{(1+r)\beta}{r+\beta}\right]^{\tau-t}[w(q_t^{\sup}) - a].$$

As a result, the expected value function of the seller, for $q_t^{\sup} \in [\theta^{**}, 1]$, is given by:

$$\begin{aligned}
W(q_t^{\sup}) &= \textstyle\sum_{\tau=t}^{\infty} \frac{1}{(1+r)^{\tau-t}}\left[F(q_\tau^{\sup}) - F\left(q_{\tau+1}^{\sup}\right)\right][w(q_\tau^{\sup}) - a] \\
&= \textstyle\sum_{\tau=t}^{\infty} \frac{1}{(1+r)^{\tau-t}}\left(q_\tau^{\sup} - q_{\tau+1}^{\sup}\right)[w(q_\tau^{\sup}) - a] \\
&= \textstyle\sum_{\tau=t}^{\infty} \frac{\left[1-\frac{(1+r)\beta}{r+\beta}\right](q_\tau^{\sup}-\theta^{**})}{(1+r)^{\tau-t}}\beta\Pi_B(q_\tau^{\sup} - \theta^{**}) \\
&= \beta\Pi_B\left[1 - \frac{(1+r)\beta}{r+\beta}\right](q_t^{\sup} - \theta^{**})^2 \textstyle\sum_{\tau=t}^{\infty} \frac{\left(\left(\frac{(1+r)\beta}{r+\beta}\right)^2\right)^{\tau-t}}{(1+r)^{\tau-t}} \\
&= \frac{\beta\Pi_B\left[1-\frac{(1+r)\beta}{r+\beta}\right]}{1-\frac{\left(\frac{(1+r)\beta}{r+\beta}\right)^2}{1+r}}(q_t^{\sup} - \theta^{**})^2 \\
&= \frac{\beta\Pi_B\left(r+\beta-\beta r-\beta^2\right)(q_t^{\sup}-\theta^{**})^2}{r+2\beta-\beta^2} \\
&= \frac{\beta\Pi_B(q_t^{\sup}-\theta^{**})^2}{2} \text{ (using the fact: } \beta^2 = r - 2r\beta)
\end{aligned} \tag{50}$$

It then can be shown that, for $q_{t+1}^{\sup} \in [\theta^{**}, q_t^{\sup}]$, the second derivative of the objective function (of maximization) w.r.t. $q_{t+1}^{\sup}$ is

$$\frac{\partial\left\{\left[F(q_t^{\sup}) - F(q_{t+1}^{\sup})\right]\frac{r\Pi_B + w'\left(q_{t+1}^{\sup}\right)}{1+r} - f\left(q_{t+1}^{\sup}\right)\left[\frac{r\Pi_B q_{t+1}^{\sup} + w\left(q_{t+1}^{\sup}\right)}{1+r} - a\right] + \frac{1}{1+r}W'\left(q_{t+1}^{\sup}\right)\right\}}{\partial q_{t+1}^{\sup}} = -\frac{(2r+\beta)\Pi_B}{1+r} < 0,$$

that is, the objective function is concave in $q_{t+1}^{\sup}$. This verifies the hypothesis that given $W\left(q_{t+1}^{\sup}\right)$, the linear strategy $w\left(q_t^{\sup}\right)$—or, equivalently, $q_{t+1}^{\sup}=\frac{(1+r)\beta}{r+\beta}\left(q_t^{\sup}-\theta^{**}\right)+\theta^{**}$—solves the maximization problem embedded in the Bellman equation (29). Consequently, function W—as specified by (50)—must also solve the Bellman equation (29). By Blackwell's theorem and the contraction mapping (fixed-point) theorem, the function W specified by (50) is the unique solution to the Bellman equation (29). Since we have $q_t^{\sup}\to\theta^{**}$ as $t\to\infty$, we must have $W\left(q_t^{\sup}\right)\to 0$ as $t\to\infty$ and $\lim_{t\to\infty}\frac{1}{(1+r)^t}W\left(q_t^{\sup}\right)=0$. The last condition implies that the Principle of Optimality[37] applies. Therefore the strategy w as characterized by (30) and (31) must maximize the expected seller payoff given the belief—the uniform probability distribution over the support $[0,q_t^{\sup}]$—and the buyer's strategy as specified by (26) and (47). Then, by definition, w must also be sequentially rational.■

A.7 Proof of Corollary 1

Part (i): This result comes from the expected value function of the seller, (50), which has already been derived in the proof of Proposition **4**.

Part (ii): As a preliminary, we derive the following limiting result:

$$\lim_{r\to 0}\beta=\lim_{r\to 0}\left(-r+\sqrt{r^2+r}\right)=0, \tag{51}$$

Consequently, we have, for $q_t^{\sup}\in[\theta^{**},1]$ and $t=0,1,\ldots,$

$$\lim_{r\to 0}w\left(q_t^{\sup}\right)=\lim_{r\to 0}\left[\beta\Pi_B\left(q_t^{\sup}-\theta^{**}\right)+w^{**}\right]=w^{**}\to a \text{ as } \Pi_S\to 0$$

and

$$\lim_{r\to 0}W^{MPBE}\left(q_t^{\sup}\right)=\lim_{r\to 0}\frac{\beta\Pi_B\left(q_t^{\sup}-\theta^{**}\right)^2}{2}=0. \tag{52}$$

Because $q_0^{\sup}=1\in[\theta^{**},1]$, we have

$$\lim_{\Pi_S\to 0,r\to 0}w\left(q_0^{\sup}\right)=a,\lim_{\Pi_S\to 0,r\to 0}W^{MPBE}\left(q_0^{\sup}\right)=0.$$

■

A.8 Proof of Theorem 3

Consider that the buyer plays the following (trigger) cut-off strategy:

$$k\left(t,w_t,c_t,q_t^{\sup};\theta\right)=\begin{cases}1 & \text{if } \theta\geq\hat{\theta}\left(t,w_t,c_t,q_t^{\sup}\right),\\ 0 & \text{otherwise,}\end{cases} \tag{53}$$

where $\hat{\theta}\left(t,w_t,c_t,q_t^{\sup}\right)$ is the cut-off type—who, condition on c_t, is indifferent between buying in period t at price w_t and buying in the next period—$\hat{\theta}\left(t,w_t,c_t,q_t^{\sup}\right)$ is such that

$$\hat{\theta}\left(t,w_t,c_t,q_t^{\sup}\right)\Pi_B-w_t=\frac{\hat{\theta}\left(t,w_t,c_t,q_t^{\sup}\right)\Pi_B-w\left(t+1,c_{t+1},\hat{\theta}\left(t,w_t,c_t,q_t^{\sup}\right)\right)}{1+r}, \tag{54}$$

[37] For a textbook treatment of the Principle of Optimality and the conditions for it to be valid, see Stokey and Lucas (1989, chapter 4).

where function w is the seller's (trigger) strategy, defined by

$$w\left(t, c_t, q_t^{\sup}\right) = \begin{cases} \left(\frac{1}{1+\eta r}\right)^t (w_0 - a) + a & \text{if } c_t = 1, \\ \beta\left(\Pi_B q_t^{\sup} - a\right) + a & \text{if } c_t = 0, q_t^{\sup} > \theta^{**}, \\ a & \text{if } c_t = 0, q_t^{\sup} \leq \theta^{**}, \end{cases} \tag{55}$$

where the value of parameter β is specified by (31); state variable c_t is the reputation (or credibility) indicator—as is determined by the following law of motion:

$$\begin{aligned} & c_{t+1} = c_t 1_{w_t = w\left(t, c_t, q_t^{\sup}\right)} \text{ for } t = 0, 1, \ldots, \\ & c_0 = 1, \end{aligned} \tag{56}$$

where state variable $q_t^{\sup}$ is an estimate of the supremum of the updated buyer's type space—which (conditional on the game has not ended) is defined by the following law of motion:

$$\begin{aligned} & q_{t+1}^{\sup} = \min\left\{q_t^{\sup}, \hat{\theta}\left(t, w_t, c_t, q_t^{\sup}\right)\right\} \text{ for } t = 0, 1, \ldots, \\ & q_0^{\sup} = 1. \end{aligned} \tag{57}$$

The indicator function $1_{w_t = w\left(t, c_t, q_t^{\sup}\right)}$ signals—by assigning value 0—a deviation of price w_t from the strategy value $w\left(t, c_t, q_t^{\sup}\right)$. The first instance of such a deviation triggers a (self-fulfilling) switch of the seller's strategy from the "tough" phase to the "soft" phase, together with a (self-fulfilling) switch of the buyer's expectation about and response to it. In the "tough" phase ($c_t = 1$), the strategies of both players depend on time t instead of state variable $q_t^{\sup}$; in the "soft" phase ($c_t = 0$) the strategies depend on state variable $q_t^{\sup}$ instead of time t.

On the (tentative PBE) equilibrium path, we have $c_t = 1$ for $t = 0, 1 \ldots,$ and

$$w_t - a = \left(\frac{1}{1+\eta r}\right)^t (w_0 - a), \tag{58}$$

that is, the price mark-up $(w_t - a)$ deflates exponentially over time with contraction factor $\frac{1}{1+\eta r} \in (0, 1)$. Notice that $\lim_{r \to 0} \frac{1}{1+\eta r} = 1$, that is, for $r \to 0$, $(w_t - a)$ contracts extremely slowly. The economic meaning of this is that the seller commits to keeping the offer price at almost the initial level w_0 within any finite period of time.

From (54), (58) and (57) we can derive the following linear relations between $\left(q_{t+1}^{\sup} - \frac{a}{\Pi_B}\right)$, $\left(\hat{\theta}\left(t, w_t, 1, q_t^{\sup}\right) - \frac{a}{\Pi_B}\right)$ and $(w_t - a)$:

$$q_{t+1}^{\sup} - \frac{a}{\Pi_B} = \hat{\theta}\left(t, w_t, 1, q_t^{\sup}\right) - \frac{a}{\Pi_B} = \frac{1 + r - \frac{1}{1+\eta r}}{r \Pi_B} (w_t - a). \tag{59}$$

At time t, the expected value for the seller is (in general) given by

$$\begin{aligned} & W\left(t, 1, q_t^{\sup}\right) = \sum_{\tau=0}^{\infty} \frac{1}{(1+r)^{\tau-t}} \left[F\left(q_\tau^{\sup}\right) - F\left(q_{\tau+1}^{\sup}\right)\right] (w_\tau - a) \\ & = \sum_{\tau=t}^{\infty} \frac{1}{(1+r)^{\tau-t}} \left(q_\tau^{\sup} - q_{\tau+1}^{\sup}\right) (w_\tau - a) \\ & = \sum_{\tau=t}^{\infty} \frac{\left(\left(q_\tau^{\sup} - \frac{a}{\Pi_B}\right) - \left(q_{\tau+1}^{\sup} - \frac{a}{\Pi_B}\right)\right)}{(1+r)^{\tau-t}} (w_\tau - a). \end{aligned}$$

For non-initial periods: $t = 1, 2, \ldots$

$$\begin{aligned}
&W\left(t,1,q_t^{\sup}\right) = \sum_{\tau=t}^{\infty} \frac{\left(\frac{1+r-\frac{1}{1+\eta r}}{r\Pi_B}\left(w_{\tau-1}-a\right)-\frac{1+r-\frac{1}{1+\eta r}}{r\Pi_B}\left(w_{\tau}-a\right)\right)}{(1+r)^{\tau-t}}\left(w_\tau - a\right)\\
&= \frac{\left(1+r-\frac{1}{1+\eta r}\right)\left(\frac{1}{\frac{1}{1+\eta r}}-1\right)}{r\Pi_B}\sum_{\tau=t}^{\infty}\frac{\left(\left(\frac{1}{1+\eta r}\right)^{\tau-t}\right)^2}{(1+r)^{\tau-t}}\left(w_t-a\right)^2\\
&= \frac{\left(1+r-\frac{1}{1+\eta r}\right)\left(\frac{1}{\frac{1}{1+\eta r}}-1\right)}{r\Pi_B\left(1-\frac{\left(\frac{1}{1+\eta r}\right)^2}{1+r}\right)}\left(\frac{1}{1+\eta r}\right)^{2t}\left(w_0-a\right)^2\\
&\to \frac{(1+\eta)\eta}{\Pi_B(1+2\eta)}\left(w_0-a\right)^2 \ \ (\text{as } r\to 0).
\end{aligned}$$

For the initial period $t = 0$

$$\begin{aligned}
&W\left(0,1,q_0^{\sup}\right) = \sum_{\tau=0}^{\infty}\frac{\left(\left(q_\tau^{\sup}-\frac{a}{\Pi_B}\right)-\left(q_{\tau+1}^{\sup}-\frac{a}{\Pi_B}\right)\right)}{(1+r)^{\tau-t}}\left(w_\tau-a\right)\\
&= \left(\left(q_0^{\sup}-\frac{a}{\Pi_B}\right)-\left(q_1^{\sup}-\frac{a}{\Pi_B}\right)\right)\left(w_0-a\right)+W\left(1,1,q_1^{\sup}\right)\\
&= \left(\left(1-\frac{a}{\Pi_B}\right)-\frac{1+r-\frac{1}{1+\eta r}}{r\Pi_B}\left(w_0-a\right)\right)\left(w_0-a\right)+W\left(1,1,q_1^{\sup}\right)\\
&\to \left(\left(1-\frac{a}{\Pi_B}\right)-\frac{1+\eta}{\Pi_B}\left(w_0-a\right)\right)\left(w_0-a\right)+\frac{(1+\eta)\eta}{\Pi_B(1+2\eta)}\left(w_0-a\right)^2 \ \ (\text{as } r\to 0)\\
&= \left(1-\frac{a}{\Pi_B}\right)\left(w_0-a\right)-\frac{(1+\eta)^2}{(1+2\eta)\Pi_B}\left(w_0-a\right)^2 \ \ (\text{as } r\to 0).
\end{aligned}$$

Given η and for $r \to 0$, the value of w_0 that maximizes $W\left(0,1,q_0^{\sup}\right)$ is:

$$w_0 = \frac{(1+2\eta)\left(\Pi_B-a\right)}{2\left(1+\eta\right)^2}+a. \tag{60}$$

We impose the above relation between free parameters w_0 and η—notice that w_0 is now a decreasing function of parameter η, which has the following limiting values:

$$w_0 \to \frac{\Pi_B+a}{2} \ (\text{"static monopoly price"}) \text{ as } \eta\to 0, \tag{61}$$

$$w_0 \to a \text{ (competitive price) as } \eta \to \infty. \tag{62}$$

Consequently, the initial period (expected) value for the seller is

$$W\left(0,1,q_0^{\sup}\right) = \frac{(1+2\eta)\left(\Pi_B-a\right)^2}{4\left(1+\eta\right)^2\Pi_B}, \tag{63}$$

which is a decreasing function of parameter η, with the following limiting values:

$$W\left(0,1,q_0^{\sup}\right) \to \frac{\left(\Pi_B-a\right)^2}{4\Pi_B} \ (\text{"static monopoly profit"}) \text{ as } \eta\to 0, \tag{64}$$

$$W\left(0,1,q_0^{\sup}\right) \to 0 \text{ as } \eta\to\infty. \tag{65}$$

The non-initial period (expected) value, for $r \to 0$, $t = 1, 2, \ldots$, is

$$W\left(t,1,q_t^{\sup}\right) = \frac{\eta\left(1+2\eta\right)\left(\Pi_B-a\right)^2}{4\left(1+\eta\right)^3\Pi_B},$$

which is an inverted-U shaped function of parameter η.

Overall, for $0 < \eta < \infty$ and $r \to 0$, we have $W(t, 1, q_t^{\sup}) > 0$ (finitely) for all $t = 0, 1, \ldots$. Intuitively, this finitely positive value $W(t, 1, q_t^{\sup})$ deters the seller from initiating a deviation from the strategy specified by (58)—because such a deviation will trigger $c_\tau = 0$ for all $\tau = t+1, t+2, \ldots$, and a switch of strategies by both players—the seller's strategy will become

$$w(t, 0, q_t^{\sup}) = \begin{cases} \beta(\Pi_B q_t^{\sup} - a) + a & \text{if } q_t^{\sup} > \theta^{**}, \\ a & \text{otherwise;} \end{cases} \tag{66}$$

and the buyer's strategy will become

$$k(t, w_t, 0, q_t^{\sup}; \theta) = \begin{cases} 1 & \text{if } \theta \geq \hat{\theta}(t, w_t, 0, q_t^{\sup}), \\ 0 & \text{otherwise,} \end{cases} \tag{67}$$

where the cut-off type is

$$\hat{\theta}(t, w_t, 0, q_t^{\sup}) = \begin{cases} \frac{1+r}{(r+\beta)\Pi_B}(w_t - a) + \theta^{**} & \text{if } q_t^{\sup} > \theta^{**}, \\ \frac{1+r}{r\Pi_B}(w_t - a) + \theta^{**} & \text{otherwise.} \end{cases} \tag{68}$$

In comparison with (30), (32) and (33) it becomes obvious that the above strategies are identical to the MPBE strategies specified in Proposition 4 (noticing that $w^{**} \to a$ and $\theta^{**} \to \frac{a}{\Pi_B}$ as $r \to 0$). Therefore these strategies of the two players are mutual best responses—they are both sequentially rational conditional on $c_t = 0$. From (34) and (52) we know the seller's expected value becomes

$$W(t, 0, q_t^{\sup}) = W^{MPBE}(q_t^{\sup}) = \frac{\beta \Pi_B \left(q_t^{\sup} - \frac{a}{\Pi_B}\right)^2}{2} \to 0 \text{ as } r \to 0.$$

To complete the proof that it is not profitable for the seller to initiate a deviation from the trigger strategy, we need to show that the expected payoff from deviation is no more than the expected value $W(t, 1, q_t^{\sup}) > 0$ (finitely). Suppose at time t we have $c_t = 1$ and $q_t^{\sup} > \theta^{**}$, and the seller chooses $w_t < w(t, 1, q_t^{\sup})$, which will trigger $c_\tau = 0$ for all $\tau = t+1, t+2, \ldots$. Then it will be expected by the buyer that next period price offer will be $w\left(t+1, 0, q_{t+1}^{\sup}\right) = \beta\left(\Pi_B q_{t+1}^{\sup} - a\right) + a$, where $q_{t+1}^{\sup} = \min\left\{q_t^{\sup}, \hat{\theta}(t, w_t, 0, q_t^{\sup})\right\}$ will be the cut-off type of buyer who will be indifferent between buying currently at w_t and buying next period at $w\left(t+1, 0, q_{t+1}^{\sup}\right)$. We have the following equation:

$$\hat{\theta}(t, w_t, 0, q_t^{\sup}) \Pi_B - w_t = \frac{\hat{\theta}(t, w_t, 0, q_t^{\sup}) \Pi_B - \left[\beta\left(\Pi_B \hat{\theta}(t, w_t, 0, q_t^{\sup}) - a\right) + a\right]}{1+r},$$

which has the solution:

$$\hat{\theta}(t, w_t, 0, q_t^{\sup}) = \frac{(1+r) w_t - (1-\beta) a}{(r+\beta) \Pi_B} \to \infty \text{ as } r \to 0 \text{ for } w_t > a \text{ (finitely).}$$

The implication is that for $r \to 0$, the only way for $w_t \geq a$ to be acceptable by any type of buyer is $w_t \to a$. Consequently, the current period payoff is $(w_t - a)\left[F(q_t^{\sup}) - F\left(q_{t+1}^{\sup}\right)\right] \to 0$, and the continuation value is no more than $\frac{W^{MPBE}\left(q_{t+1}^{\sup}\right)}{1+r} \to 0$; while the expected value of playing the trigger strategy is $W(t, 1, q_t^{\sup}) > 0$ (finitely) for finite $\eta > 0$. That is, the deviation from the trigger strategy cannot be profitable for $r \to 0$. As a result, the trigger strategies constitute a PBE. For $r \to 0$, as η ranges from 0 to ∞, $W(0, 1, q_0^{\sup})$ ranges from $\frac{(\Pi_B - a)^2}{4\Pi_B}$ ("static monopoly profit") to 0.■