Behavioral Perfect Equilibrium
in Bayesian Games

Elnaz Bajoori

No. 16/13

BATH ECONOMICS RESEARCH PAPERS

Department of Economics

Department of Economics
UNIVERSITY OF BATH

Behavioral Perfect Equilibrium in Bayesian Games

Elnaz Bajoori * János Flesch † Dries Vermeulen ‡

November 14, 2013

Abstract

We develop the notion of perfect equilibrium in Bayesian games where players may have infinitely many types and actions. We formulate and examine several definitions of perfection in this setting using behavior strategies. The differences between the definitions arise from using various notions of convergence such as uniform, pointwise and almost everywhere pointwise convergence on the type space.

Additionally, we illustrate the use of perfect equilibrium in the context of a second-price auction with incomplete information. In this auction, perfect equilibrium selects a unique equilibrium in a class of pure separating strategy profiles. Thus, the selected equilibrium is the unique pure strategy equilibrium in differentiable strategies that separates types. Moreover, when informational uncertainty vanishes, the selected equilibrium converges to the classical truthful dominant strategy equilibrium.

We also show that other, less intuitive, equilibria in which types are (partially) pooled are ruled out by our selection criterion. We further argue that standard selection criteria for second-price auctions, such as dominant strategy equilibrium or truthful bidding, have no bite in our incomplete information example. Bidders have no dominant strategies when information is incomplete, and the selected perfect separating equilibrium is not sincere.

JEL Codes. C72.

Keywords. Trembling hand perfect equilibrium, Bayesian game with infinite type spaces, Behavior strategy, Second-price auction with incomplete information.

* e.bajoori@bath.ac.uk. University of Bath, Dept. of Economics, Bath, UK.

† j.flesch@maastrichtuniversity.nl. Maastricht University, School of Business and Economics, Dept. of Quantitative Economics, Maastricht, The Netherlands.

‡ Corresponding author, d.vermeulen@maastrichtuniversity.nl. Maastricht University, School of Business and Economics, Dept. of Quantitative Economics, Maastricht, The Netherlands., P.O. Box 616, 6200 MD Maastricht, The Netherlands.

1 Introduction

For normal form games with finite action spaces, Nash [16] introduced the concept of Nash equilibrium and proved its existence. Since Nash equilibria are not always intuitive as a solution, several refinements of Nash equilibrium have been proposed in the literature. Among these, perfect equilibrium (Selten [21]) is one of the most commonly used refinement concepts.

Many applications of Nash equilibrium and its refinements are in the context of incomplete information games. Examples of games with incomplete information are signaling games (Kreps and Wilson [12]), principal-agent models (Laffont and Martimort [13]), and models of reputation such as the chain store paradox (Selten [22], Govindan [9]). In particular in auction design refinements of Nash equilibrium are used to reduce the number of equilibria and rule out the less intuitive or desirable equilibria. For example in the Vickrey auction, selection on dominant strategies singles out the truthful equilibrium, and rules out the (many) ex post equilibria in which partial pooling of types occurs. Also Jackson et al [10] employ perfect equilibrium in one of their examples to eliminate equilibria in which players may bid above their maximum possible valuation.

Our aim is to develop the refinement of perfect equilibrium in the context of such applications, in particular auction design. We focus on the game theoretic framework of Bayesian games, the standard tool to model incomplete information. We study the class of Bayesian games in which players, after the information phase in which they learn their type, play a one-shot game. We develop the notion of perfect equilibrium for such games, and in effect we propose three possible variations. We study the relations between these variations, and illustrate their use in the context of a second price auction.

When applying refinements to auctions and incomplete information games there are typically two problems to tackle. First, such games often feature discontinuities in the payoff functions. This issue is addressed in for example Reny [19] and Jackson et al. [10].

Second, incomplete information games, especially auctions, often have continuum type spaces and action spaces. Refinements for incomplete information games, such as for example sequential equilibrium (Kreps and Wilson [12]) and perfect Bayesian equilbrium (Fudenberg and Tirole [8]), are as a rule developed for finite games, and as such not immediately applicable in the context of auction theory and similar economic applications with incomplete information. An influential first attempt to generalize known equilibrium refinements to more general classes of games is the working paper by Myerson and Reny [15]. Our paper also contributes to the development of equilibrium refinement for incomplete information games, and generalizes the notion of perfect equilibrium to the class of Bayesian games [1].

[1]The result in Myerson and Reny [15] covers a large class of incomplete information games. Their solution

We take a conceptual view. A central solution concept in Bayesian games is that of Bayesian Nash Equilibrium (BNE), a direct extension of Nash equilibrium to games with incomplete information. A BNE is a profile of behavior strategies, one for each player, such that each player's strategy, given any type for this player, is an expected value maximizer given the strategies of his opponents, where expectations are taken over all possible types of the opponents [2]. So, the best response property that characterizes the equilibrium concept is required at the interim stage when the player already knows his own type [3].

In this paper we work with ex-interim probabilities, unless mentioned otherwise. Our goal is to define the notion of perfect BNE in Bayesian games, analyze its properties, and illustrate its predictive power in an elaborate example of a second price auction with incomplete information. Our work is motivated by the observation that, under incomplete information, standard selection criteria in auction design, such as dominant strategy equilibrium and truthful reporting, no longer have a bite. Also notions such as sequential equilibrium and perfect Bayesian equilibrium (Kreps and Wilson [12], Fudenberg and Tirole [8], Bonanno [6]) only have selective power in the context of extensive form games with multiple rounds. Thus in the setting of single round sealed bid auctions these refinements do not reduce the set of equilibria. According to the theory of refinements, the selection criterion that best suits our context is perfect equilibrium. This is the solution concept we develop for the class of Bayesian games.

We base our approach on completely mixed behavior strategies, i.e. behavior strategies which prescribe a completely mixed probability measure for each possible type. We first define perfection for general behavior strategy profiles, so not necessarily for BNEs yet. We define three versions of perfection, which differ in the kind of convergence over the set of types. Roughly speaking, a behavior strategy profile β is called perfect if there is a sequence $(\beta^k)_{k=1}^{\infty}$ of completely mixed behavior strategy profiles such that, for every player i, the distance between β_i^k and β_i and the distance between β_i^k and player i's best responses against β^k both converge to 0. We measure distance on the set of actions by means of the weak metric. Convergence on the set of types is expressed in one of the following three senses: uniformly for all types, pointwise for each type, and for almost every type. Consequently, we call β uniform-perfect, pointwise-perfect, or a.e.-pointwise-perfect. In general, such a profile β is not necessarily a BNE, since we did not impose strong conditions on the type and actions spaces and on the payoff functions in the Bayesian game.

concept assigns finitely additive probability measures, while our construction remains within the environment of countably additive probability measures.

[2] Behavior strategies are much in the spirit of Bayesian games. A behavior strategy of a player prescribes a probability measure on his set of actions, depending on the type of this player, that satisfies an additional measurability assumption. An alternative approach based on distributional strategies is investigated in Bajoori [4].

[3] A similar but weaker concept arises if one calculates ex-ante probabilities, and only requires the best response property for each player before he receives his own type. This approach is investigated for example by Reny [20] and by Milgrom and Weber [14].

If β is also a BNE, then we call β a uniform-perfect, pointwise-perfect, or a.e.-pointwise-perfect BNE, respectively.

In the first part of the paper we analyze each of these three notions of perfection in Bayesian games in detail, and we discuss the relations between them. We also pay attention to special cases, mainly when compactness is imposed on the type and action spaces or when the payoff functions satisfy a continuity property. We also briefly discuss finite Bayesian games, i.e. when there are only finitely many types and actions.

In the second part of the paper, we illustrate the use of perfection as a tool to select the more intuitive BNEs in Bayesian games. We do so by examining a sealed-bid second-price auction with two bidders. In this auction, the valuation function of each bidder does not only depend on his own type, but also on the type of his opponent. More precisely, for each $i = 1, 2$, bidder i's valuation is given by $v_i = 5 + t_i - \alpha t_j$, where $\alpha \in (0, 1)$ and $j \neq i$. Jackson et al. [10] and [11] considered the first price version of this auction with $\alpha = 4$ and proved that no BNE exists if each player has a positive probability to win in case of a tie.

This second-price auction admits multiple BNEs, but we find that perfection selects a BNE β that is unique in a certain class of separating strategy profiles. We also show that perfection rules out many BNEs in which pooling occurs among types. We highlight the subtleties involved in the choice of sequence of completely mixed behavior strategy profiles $(\beta^k)_{k=1}^{\infty}$, because the most straightforward candidate for this sequence, the uniform distribution, does not select a BNE.

We also show that standard complete information selection criteria such as dominant strategies and truthful reporting do not apply here. Sincere reporting is not an equilibrium, the selected BNE β is not sincere, and neither player has a dominant strategy in this auction. This emphasizes the necessity to employ more sophisticated refinement techniques such as perfect BNE to select among many BNEs in this auction.

As a final remark, we already observed that for example Reny [20] and Milgrom and Weber [14] take an ex ante approach. However, both from a conceptual and a computational view, the ex interim approach seems to be preferable. Conceptual, since perfect BNE is defined directly at the level of behavioral strategies. Computational, since perfect Bayesian Nash equilibrium is calculated at the level of the Bayesian game itself, not in the context of the induced strategic form game. Our computations in the second price auction emphasize these observations.

Related literature. Simon and Stinchcombe [23] define perfect equilibrium for strategic form games with compact action spaces. They discuss two essentially different appraoches. The first approach is a direct generalization of Selten's original definition, based on the notion of completely mixed strategies. The second approach to perfect equilibrium, by Simon and Stinchcombe referred to as the finitistic approach, uses the notion of an ε-perfect equilibrium in finite approximations of the original game. They show existence of these notions of perfect equilibrium, and investigate the properties of and relations between the various resulting solution concepts.

Bajoori et al. [3] examine the two approaches proposed in Simon and Stinchcombe [23] in further detail and provide an improved definition of the finitistic approach. Their results seem to imply a critique on the finitistic approach.

Jackson et al. [10] study games with incomplete information and discontinuous payoffs. In the incomplete information setting discontinuities often arise from indifferences between players' choices and the particular resolution of such indifferences in the description of the game. The paper shows in several examples that the resulting discontinuities may cause extreme behavior in equilibrium, or even non-existence of equilibrium. They show that the introduction of a communication phase before the start of the game may mitigate strategic effects, and restore existence of equilibrium with truthful reporting.

Reny [20] shows, under general conditions, the existence of a monotone pure-strategy equilibrium. The main innovations in the paper are that the result is shown to hold for a wide class of partially ordered spaces, and that best response sets only need to be join-closed (meaning that the join of two best responses is again a best response).

Myerson and Reny [15] develop the concept of sequential equilibrium for a very general class of multistage games with incomplete information. They prove existence of sequential equilibrium in terms of induced finitely additive conditional probability distributions.

Fudenberg and Tirole [8] define perfect Bayesian equilibrium and sequential equilibrium for finite Bayesian multi-period games, and show that these notions coincide when each player only has two types.

Our paper is structured as follows. First, we discuss some preliminary notions in Section 2 and present the model of Bayesian games in Section 3. Then, in Section 4, we define the concept of perfect BNE and analyze its properties. In Section 5, we apply our results to the above mentioned auction. The paper ends with an extensive appendix, which contains the proofs of several technical results that we use in earlier parts of the paper.

2 Preliminaries

A semi-ring $\mathcal{S}$ is a collection of subsets of a set X satisfying

1. $\emptyset \in \mathcal{S}$,
2. If $E, F \in \mathcal{S}$ then $E \cap F \in \mathcal{S}$,

3. If $E, F \in \mathcal{S}$ then there exists a countable collection $(G_i)_{i \in I}$ of pairwise disjoint sets in $\mathcal{S}$ such that $E \setminus F = \bigcup_{i \in I} G_i$.

A measure on a semi-ring $\mathcal{S}$ is a function $\mu : \mathcal{S} \to [0, \infty)$ such that

1. $\mu(\emptyset) = 0$,
2. (σ-additivity) For every countable collection $(A_i)_{i \in I}$ of pairwise disjoint sets in $\mathcal{S}$ with $\bigcup_{i \in I} A_i \in \mathcal{S}$ we have
$$\mu(\bigcup_{i \in I} A_i) = \sum_{i \in I} \mu(A_i).$$

A σ-field Σ is a collection of subsets of a set X satisfying

1. $X \in \Sigma$,
2. If $E \in \Sigma$ then $X \setminus E \in \Sigma$,
3. If $(E_i)_{i \in I}$ is a countable collection of sets in Σ then $\bigcup_{i \in I} E_i \in \Sigma$.

The pair (X, Σ) is called a measurable space.

A metric on a nonempty set X is a function $d : X \times X \to [0, \infty)$ such that for all $x, y, z \in X$ we have

1. $d(x, y) = 0 \Leftrightarrow x = y$ (identity of indiscernibles)
2. $d(x, y) = d(y, x)$ (symmetry)
3. $d(x, z) \leq d(x, y) + d(y, z)$ (triangle inequality).

The ordered pair (X, d) is called a metric space. For $X \subseteq \mathbb{R}$ we always consider the standard Euclidean metric given by $d(x, y) = |x - y|$. The distance $d(x, B)$ between a point $x \in X$ and a nonempty set $B \subseteq X$ is defined as

$$d(x, B) = \inf\{d(x, y) \mid y \in B\},$$

and if $B = \emptyset$, then $d(x, B) = \infty$. The ε-neighborhood of B is denoted by $B^\varepsilon = \{x \in X | \, d(x, B) < \varepsilon\}$. A set $U \subseteq X$ is called open if for every $x \in U$ there is an $\varepsilon > 0$ such that

$$\{y \in X \mid d(x, y) < \varepsilon\} \subseteq U.$$

A set $F \subseteq X$ is called closed if its complement $X \setminus F$ is open, and a set $C \subseteq X$ is called compact if, for every collection $\{U_\alpha \mid \alpha \in A\}$ of open sets such that $C \subseteq \cup_{\alpha \in A} U_\alpha$, it holds that A has a finite subset A' such that $C \subseteq \cup_{\alpha \in A'} U_\alpha$. Every compact set in a metric space is closed as well. The topology on X induced by metric d is the collection of all open sets, and the Borel σ-field Σ on X is the smallest σ-field that contains all open sets. A measure on (X, Σ) is a function $\mu : \Sigma \to [0, \infty)$ such that $\mu(\emptyset) = 0$ and μ is σ-additive. A measure μ is called a probability measure if $\mu(X) = 1$. For a metric space (X, d), the set of probability measures on the Borel σ-field on X is denoted by $\Delta(X)$. A probability measure $\mu \in \Delta(X)$ is completely mixed if $\mu(U) > 0$ for every nonempty open subset U of X. The weak (Prokhorov) metric ρ^w on $\Delta(X)$ is defined for every $\mu, \nu \in \Delta(X)$ by

$$\rho^w(\mu, \nu) = \inf\{\varepsilon > 0 \mid \forall B \in \Sigma : \mu(B) \leq \nu(B^\varepsilon) + \varepsilon \text{ and } \nu(B) \leq \mu(B^\varepsilon) + \varepsilon\}.$$

Let $\mu_n \in \Delta(X)$ for every $n \in \mathbb{N}$ and let $\mu \in \Delta(X)$. It is known that if the sequence μ_n converges to μ with respect to ρ^w, then $\int_X f(x)\mu_n(dx)$ converges to $\int_X f(x)\mu(dx)$ for every bounded and continuous function $f : X \to \mathbb{R}$. Conversely, if X is separable and $\int_X f(x)\mu_n(dx)$ converges to $\int_X f(x)\mu(dx)$ for every bounded and Lipschitz function $f : X \to \mathbb{R}$, then μ_n converges to μ with respect to ρ^w. Further, if X is compact, then so is $\Delta(X)$ with respect to ρ^w (cf. Prokhorov [18] and Parthasarathy [17]).

Let $(X_1, \Sigma_1), (X_2, \Sigma_2), \ldots, (X_n, \Sigma_n)$ be measurable spaces. The product σ-field $\otimes_{i=1}^n \Sigma_i$ on $\times_{i=1}^n X_i$ is defined to be the smallest σ-field that contains all the sets in $\times_{i=1}^n \Sigma_i$. Let μ_i be a measure on (X_i, Σ_i), for every $1 \leq i \leq n$, and μ be a measure on the semi-ring $\times_{i=1}^n \Sigma_i$ defined by

$$\mu(A_1 \times \ldots \times A_n) = \prod_{i=1}^n \mu_i(A_i),$$

for every $A_i \in \Sigma_i$, $1 \leq i \leq n$. The product measure μ^* is the extension of the measure μ to the measurable space $(\times_{i=1}^n X_i, \otimes_{i=1}^n \Sigma_i)$ via the formula

$$\mu^*(E) = \inf\left\{\sum_{k=1}^\infty \mu(\times_{i=1}^n A_i^k) \mid E \subseteq \bigcup_{k=1}^\infty (\times_{i=1}^n A_i^k),\ A_i^k \in \Sigma_i,\ 1 \leq i \leq n,\ \forall k \in \mathbb{N}\right\}. \tag{1}$$

According to Carathéodory's theorem, this extension is unique.

Let (X, Σ_1), (Y, Σ_2) be two measurable spaces. A function $f : X \longrightarrow Y$ is called measurable if $f^{-1}(B) \in \Sigma_1$ for every $B \in \Sigma_2$.

3 Bayesian Games

Definition 1 *A Bayesian game is a tuple* $\Gamma = (N, (T_i, d_{T_i})_{i \in N}, (A_i, d_{A_i})_{i \in N}, (\mu_i)_{i \in N}, (\Pi_i)_{i \in N})$ *where:*

1. $N = \{1, 2, .., n\}$ *is the set of players.*

2. *T_i is a nonempty set of player i's possible types with metric d_{T_i}. Let $\mathcal{T}_i$ denote the induced Borel σ-field on T_i, $T = \times_{i=1}^n T_i$, and $\mathcal{T} = \otimes_{i=1}^n \mathcal{T}_i$.*

3. *A_i is a nonempty set of player i's actions with metric d_{A_i}. Let $\mathcal{A}_i$ denote the induced Borel σ-field on A_i, $A = \times_{i=1}^n A_i$, and $\mathcal{A} = \otimes_{i=1}^n \mathcal{A}_i$.*

4. *μ_i is a probability measure on $(T_i, \mathcal{T}_i)$ for player i. Let $\mu = \times_{i=1}^n \mu_i$ be the product measure on $(T, \mathcal{T})$.*

5. *$\Pi_i : T \times A \to \mathbb{R}$ is player i's payoff function, bounded and measurable with respect to $\mathcal{T} \otimes \mathcal{A}$. Let $\Pi = (\Pi_i)_{i=1}^n$.*

The Bayesian game Γ is played as follows: First, nature draws a type $t_i \in T_i$ for each player i according to the probability measure μ_i. Each player i learns his own type t_i, but not the types of the other players. Then, each player i chooses an action $a_i \in A_i$, simultaneously and independently. Finally, depending on the types $t = (t_1, \ldots, t_n)$ and the chosen actions $a = (a_1, \ldots, a_n)$, each player i receives the payoff $\Pi_i(t, a)$.

Now, we discuss different classes of strategies for the players. We start with the simplest ones.

Definition 2 *A pure strategy for player i is a measurable function $p_i : T_i \to A_i$.*

Thus, a pure strategy prescribes one specific action depending on the player's type. Now we move on to the strategies which use some randomization for the choice of an action. In the finite version of the model, i.e., when T_i and A_i are finite, a mixed strategy is a probability measure on the set of pure strategies, whereas a behavior strategy prescribes, for each possible type, a probability measure on the set of available actions. Aumann [2] observed however that, for our infinite model, the above view of a mixed strategy leads to measure theoretic problems and does not provide an acceptable definition. Instead, a mixed strategy should be modeled by a random variable with values in the set of pure strategies, whose domain is a probability measure space that is used as the randomization device. This is the underlying idea of the definition of a mixed strategy for player i as a measurable function $\alpha_i : T_i \times [0,1] \to A_i$, where the uniform distribution is imposed on $[0,1]$. The interpretation of a mixed strategy α_i is that, after observing his own type t_i and drawing a randomization-variable s_i from $[0,1]$ according to uniform distribution, player i plays $\alpha_i(t_i, s_i)$. Note that, for every $s_i \in [0,1]$, the section function $\alpha_i(\cdot, s_i) : T_i \to A_i$ is a pure strategy. Behavior strategies are defined with similar considerations:

Definition 3 *(Milgrom and Weber) A behavior strategy for player i is a function $\beta_i : T_i \times \mathcal{A}_i \to [0,1]$ such that*

1. *the section function* $\beta_i(t_i, \cdot) : \mathcal{A}_i \to [0, 1]$ *is a probability measure for every* $t_i \in T_i$,

2. *the section function* $\beta_i(\cdot, B) : T_i \to [0, 1]$ *is measurable for every* $B \in \mathcal{A}_i$.

A behavior strategy β_i prescribes, depending on player i's type t_i, to choose an action according to the probability measure $\beta_i(t_i, \cdot)$. The second condition in the definition is included so that the ex-ante probability that player i's action falls into a set $B \in \mathcal{A}_i$ exists and is equal to $\int_{T_i} \beta_i(t_i, B)\ \mu_i(dt_i)$. We will usually define a behavior strategy by specifying the section function $\beta_i(t_i, \cdot)$ for every type $t_i \in T_i$. Behavior strategies are well suited for our purpose to define perfect equilibrium. Indeed, the probability measures $\beta_i(t_i, \cdot)$ are sufficient to describe player i's behavior. Moreover, as Aumann [2] showed, there is a many-to-one mapping from mixed to behavior strategies that preserves the players' expected payoffs, so mixed strategies would have no significant added value. For these reasons, we build our definitions on behavior strategies. From now on, by a strategy we will always mean a behavior strategy, unless mentioned otherwise.

Definition 4 *A strategy* β_i *for player* i *is called deterministic if, for every type* $t_i \in T_i$, *there is an action* $a_{i,t_i} \in A_i$ *such that* $\beta_i(t_i, \cdot)$ *is the Dirac measure on* a_{i,t_i}. *A strategy* β_i *is called completely mixed if the section function* $\beta_i(t_i, \cdot) : \mathcal{A}_i \to [0, 1]$ *is a completely mixed probability measure for every* $t_i \in T_i$. *The vector* $\beta = (\beta_1, \beta_2, \ldots, \beta_n)$, *where* β_i *is a strategy of player* i, *is called a strategy profile.*

Each pure strategy p_i naturally induces a deterministic strategy β_i, for which $\beta_i(t_i, \cdot)$ is the Dirac measure on $p_i(t_i)$ for every type $t_i \in T_i$.

For every strategy profile β and every player $i \in N$, we write $\beta_{-i} = (\beta_j)_{j \in N \setminus \{i\}}$ to denote the profile consisting of strategies of the players in $N \setminus \{i\}$. Further, we use $\mathbb{E}_{\tau_{-i}}(\Pi_i \mid t_i, \beta_i)$ to denote player i's expected payoff, given his type t_i and his strategy β_i, against a strategy profile τ_{-i}. Thus,

$$\mathbb{E}_{\tau_{-i}}(\Pi_i \mid t_i, \beta_i) = \int_{T_{-i}} \int_{A_{-i}} \int_{A_i} \Pi_i(t, a)\ \beta_i(t_i, da_i)\ \tau_{-i}(t_{-i}, da_{-i})\ \mu_{-i}(dt_{-i}).$$

In the expected payoff above, by integrals with respect to $\tau_{-i}(t_{-i}, da_{-i})$ and $\mu_{-i}(dt_{-i})$ we mean the iterated integrals with respect to $\tau_j(t_j, da_j)$ and $\mu_j(dt_j)$ for all $j \neq i$. Fubini's Theorem and Theorem 27 in the appendix guarantee the existence of the iterated integrals in the expression above, and also that the order of integration with respect to $\beta_i(t_i, da_i)$ and $\tau_j(t_j, da_j)$, $j \neq i$, is not relevant. In the special case where player i uses a deterministic strategy with corresponding pure strategy p_i, player i's expected payoff is denoted simply by $\mathbb{E}_{\tau_{-i}}(\Pi_i \mid t_i, p_i)$, and it is equal to

$$\mathbb{E}_{\tau_{-i}}(\Pi_i \mid t_i, p_i) = \int_{T_{-i}} \int_{A_{-i}} \Pi_i\left(t, (p_i(t_i), a_{-i})\right)\ \tau_{-i}(t_{-i}, da_{-i})\ \mu_{-i}(dt_{-i}).$$

Since, in the expected payoffs $\mathbb{E}_{\tau_{-i}}(\Pi_i \mid t_i, \beta_i)$ and $\mathbb{E}_{\tau_{-i}}(\Pi_i \mid t_i, p_i)$, it is irrelevant how player i chooses his actions for types other than t_i, we can naturally define $\mathbb{E}_{\tau_{-i}}(\Pi_i \mid t_i, \sigma_i)$ and $\mathbb{E}_{\tau_{-i}}(\Pi_i \mid t_i, a_i)$ for every probability measure σ_i on $(A_i, \mathcal{A}_i)$ and respectively for every action $a_i \in A_i$.

A probability measure σ_i on $(A_i, \mathcal{A}_i)$ is called a best response of player i for type $t_i \in T_i$ against a strategy profile τ_{-i}, if for every probability measure σ_i' on $(A_i, \mathcal{A}_i)$ we have

$$\mathbb{E}_{\tau_{-i}}(\Pi_i \mid t_i, \sigma_i) \geq \mathbb{E}_{\tau_{-i}}(\Pi_i \mid t_i, \sigma_i').$$

The set of such best responses is denoted by $BR_i(t_i, \tau_{-i})$. For a strategy profile τ, we will also use the notation $BR_i(t_i, \tau)$ instead of $BR_i(t_i, \tau_{-i})$.
A strategy β_i is called a best response of player i against a strategy profile τ_{-i}, if $\beta_i(t_i, \cdot) \in BR_i(t_i, \tau_{-i})$ for every $t_i \in T_i$. The set of such best responses is denoted by $BR_i(\tau_{-i})$. For a strategy profile τ, we will also use the notation $BR_i(\tau)$ instead of $BR_i(\tau_{-i})$. Note that all these best response sets can be empty, which is illustrated by the following simple example.

Example 5 Consider the following Bayesian game with only one player: $T_1 = \{t_1\}$, $A_1 = [0, 1]$, $\Pi_1(t_1, x) = x$ for every $x \in [0, 1)$, and $\Pi_1(t_1, 1) = 0$. In this game, the set of the best responses (that is, optimal strategies) of player 1 is empty. ♢

Now we define a central solution concept of Bayesian games, namely the concept of Bayesian Nash equilibrium.

Definition 6 *A strategy profile $\beta = (\beta_1, \beta_2, \ldots, \beta_n)$ is called a Bayesian Nash equilibrium (BNE), if β_i is a best response of player i against β, for every player i.*

There are Bayesian games which admit no BNE at all, for instance the game in Example 5. In many Bayesian games of economic interest, however, there exist multiple BNEs, and some of them are arguably more intuitive than others. As mentioned before, our goal is to develop the definition of perfection for BNEs, which can be a useful tool in such games to distinguish the more intuitive BNEs.

4 Behavioral Perfect Bayesian Nash Equilibrium

In this section we propose a number of possible definitions of perfect BNE for Bayesian games. We introduce three main definitions. In each of them, a different notion of convergence is used on the set of strategies. At the end of the section, we draw some conclusions in the context of finite games, i.e. when there are only finitely many types and actions.

4.1 Uniform-Perfect BNE

Definition 7 *A strategy profile $\beta = (\beta_1, \ldots, \beta_n)$ is called uniform-perfect, if there exists a sequence of profiles of completely mixed strategies $(\beta^k)_{k=1}^{\infty} = (\beta_1^k, \ldots, \beta_n^k)_{k=1}^{\infty}$ with the following properties for every player i:*

(1) $\lim_{k\to\infty} \sup_{t_i \in T_i} \rho^w(\beta_i^k(t_i, \cdot), \beta_i(t_i, \cdot)) = 0,$

(2) $\lim_{k\to\infty} \sup_{t_i \in T_i} \rho^w(\beta_i^k(t_i, \cdot), BR_i(t_i, \beta^k)) = 0.$

A strategy profile β is called a uniform-perfect BNE if β is a BNE and uniform-perfect.

Note that conditions (1) and (2) in the above definition require a uniform convergence of the corresponding distances over the set of possible types. In the special case when the game has complete information, i.e. each T_i is a singleton, our definition coincides with Simon and Stinchcombe's weak perfect equilibrium in [23].

A uniform-perfect strategy profile, and therefore a uniform-perfect BNE, does not always exist, which is for instance the case in Example 5, because the best response set in this game is always empty. Furthermore, even if a uniform-perfect strategy profile exists, it is not necessarily a BNE, which is illustrated by the following example.

Example 8 Consider the following Bayesian game with only one player: $T_1 = \{t_1\}$, $A_1 = [0, 1]$, $\Pi_1(t_1, x) = 1$ for every $x \in [0, 1)$, and $\Pi_1(t_1, 1) = 0$. Define strategies β_1 and β_1^k, for every $k \in \mathbb{N}$, as follows: $\beta_1(t_1, \cdot) = \delta_1(\cdot)$ and $\beta_1^k(t_1, \cdot) = (1 - \frac{1}{k})\delta_{1-\frac{1}{k}}(\cdot) + \frac{1}{k}\sigma(\cdot)$, for every $k \in \mathbb{N}$, where δ_x is the Dirac measure on action x and σ is the uniform distribution on the Borel sets of $[0, 1]$. Note that the set of the pure best responses (that is, optimal strategies) of player 1 consists of all actions in $[0, 1)$. So, the strategy β_1 and the sequence of completely mixed strategies $(\beta_1^k)_{k=1}^{\infty}$ satisfy conditions (1) and (2) of Definition 7. Therefore, β_1 is a uniform-perfect strategy profile, but it is clearly not a BNE. $\diamondsuit$

A uniform-perfect strategy profile may fail to exist, even if the type and action spaces are compact and the payoff functions are continuous, as is shown in Example 11. Nevertheless, it follows from Theorem 10 under fairly weak conditions that every uniform-perfect strategy profile is a BNE.

4.2 Pointwise-Perfect BNE

Definition 9 *A strategy profile $\beta = (\beta_1, \ldots, \beta_n)$ is called pointwise-perfect, if there exists a sequence of profiles of completely mixed strategies $(\beta^k)_{k=1}^{\infty} = (\beta_1^k, \ldots, \beta_n^k)_{k=1}^{\infty}$ with the following properties for every player i and each type $t_i \in T_i$:*

(1) $\lim_{k\to\infty} \rho^w(\beta_i^k(t_i, \cdot), \beta_i(t_i, \cdot)) = 0,$

(2) $\lim_{k\to\infty} \rho^w(\beta_i^k(t_i,\cdot), BR_i(t_i,\beta^k)) = 0.$

A strategy profile β is called a pointwise-perfect BNE if β is a BNE and pointwise-perfect.

Note that conditions (1) and (2) in the above definition require pointwise convergence of the corresponding distances over the set of possible types.

A pointwise-perfect strategy profile, and therefore a pointwise-perfect BNE, does not always exist, which is for instance the case in Example 5, because the best response set in this game is always empty. Furthermore, it follows from Example 8 that a pointwise-perfect strategy profile is not necessarily a BNE.

We do not know if a pointwise-perfect strategy profile always exists under the condition that the type and action spaces are compact and the payoff functions are continuous. Nevertheless, we have the following result.

Theorem 10 *In a Bayesian game where the action spaces are separable and $\Pi_i(t,\cdot)$ is continuous on A for every player i and every $t \in T$, every pointwise-perfect strategy profile is a BNE.*

Proof. Suppose that $\beta = (\beta_1, \ldots, \beta_n)$ is a pointwise-perfect strategy profile in such a Bayesian game. For β, take a sequence of strategy profiles $(\beta^k)_{k=1}^\infty$ as in Definition 9. Fix a player i and a type $t_i \in T_i$. Due to condition (2) in Definition 9, there exists a $K_i(t_i) \in \mathbb{N}$ such that $BR_i(t_i,\beta^k)$ is nonempty for every $k \geq K_i(t_i)$. Therefore, by condition (2) once more, there exists a sequence $(\sigma_{i,t_i}^k)_{k=1}^\infty$ of probability measures on $(A_i,\mathcal{A}_i)$ such that $\sigma_{i,t_i}^k \in BR_i(t_i,\beta^k)$ for every $k \geq K_i(t_i)$ and $\rho^w(\beta_i^k(t_i,\cdot),\sigma_{i,t_i}^k) \to 0$ as $k \to \infty$. By condition (1) and by the triangle inequality for ρ^w, this implies that $\rho^w(\beta_i(t_i,\cdot),\sigma_{i,t_i}^k) \to 0$ as $k \to \infty$.

For every $k \geq K_i(t_i)$ and every probability measure σ' on $(A_i,\mathcal{A}_i)$, we have due to $\sigma_{i,t_i}^k \in BR_i(t_i,\beta^k)$ that $\mathbb{E}_{\beta_{-i}^k}(\Pi_i \mid t_i,\sigma_i^k) \geq \mathbb{E}_{\beta_{-i}^k}(\Pi_i \mid t_i,\sigma')$, which means that

$$\int_{T_{-i}}\int_{A_{-i}}\int_{A_i} \Pi_i\left((t_i,t_{-i}),a\right)\ \sigma_{i,t_i}^k(da_i)\ \beta_{-i}^k(t_{-i},da_{-i})\ \mu_{-i}(dt_{-i}) \qquad (2)$$
$$\geq \int_{T_{-i}}\int_{A_{-i}}\int_{A_i} \Pi_i\left((t_i,t_{-i}),a\right)\ \sigma'(da_i)\ \beta_{-i}^k(t_{-i},da_{-i})\ \mu_{-i}(dt_{-i}).$$

According to Theorem 3.2 in [5], in a separable product space the weak convergence of the product measure is equivalent to the weak convergence of the marginal measures. Hence, because A is separable, the weak convergence of σ_{i,t_i}^k to $\beta_i(t_i,\cdot)$ and of $\beta_j^k(t_j,\cdot)$ to $\beta_j(t_j,\cdot)$ for every player $j \neq i$ and every type $t_j \in T_j$ implies that, for any $t_{-i} \in T_{-i}$, the product measure $\sigma_{i,t_i}^k \times \left(\times_{j\neq i}\beta_j^k(t_j,\cdot)\right)$ on $(A,\mathcal{A})$ weakly converges to the product measure $\beta_i(t_i,\cdot) \times (\times_{j\neq i}\beta_j(t_j,\cdot))$. Since the payoff function $\Pi_i((t_i,t_{-i}),\cdot)$ is continuous on A for every $t_{-i} \in T_{-i}$, by Fubini's theorem we have for every $t_{-i} \in T_{-i}$

$$\lim_{k\to\infty}\int_{A_{-i}}\int_{A_i} \Pi_i\left((t_i,t_{-i}),a\right)\ \sigma_{i,t_i}^k(da_i)\ \beta_{-i}^k(t_{-i},da_{-i})$$

$$= \int_{A_{-i}} \int_{A_i} \Pi_i\left((t_i, t_{-i}), a\right)\ \beta_i(t_i, da_i)\ \beta_{-i}(t_{-i}, da_{-i}),$$

and similarly

$$\lim_{k\to\infty} \int_{A_{-i}} \int_{A_i} \Pi_i\left((t_i, t_{-i}), a\right)\ \sigma'(da_i)\ \beta^k_{-i}(t_{-i}, da_{-i})$$

$$= \int_{A_{-i}} \int_{A_i} \Pi_i\left((t_i, t_{-i}), a\right)\ \sigma'(da_i)\ \beta_{-i}(t_{-i}, da_{-i}).$$

Thus, if we take the limit in inequality (2) when $k \to \infty$, by the dominated convergence theorem we obtain

$$\int_{T_{-i}} \int_{A_{-i}} \int_{A_i} \Pi_i\left(t, a\right)\ \beta_i(t_i, da_i)\ \beta_{-i}(t_{-i}, da_{-i})\ \mu_{-i}(dt_{-i})$$

$$\geq \int_{T_{-i}} \int_{A_{-i}} \int_{A_i} \Pi_i\left(t, a\right)\ \sigma'(da_i)\ \beta_{-i}(t_{-i}, da_{-i})\ \mu_{-i}(dt_{-i}),$$

which means $\mathbb{E}_{\beta_{-i}}(\Pi_i \mid t_i, \beta_i(t_i, \cdot)) \geq \mathbb{E}_{\beta_{-i}}(\Pi_i \mid t_i, \sigma')$. Hence, $\beta_i(t_i, \cdot) \in BR_i(t_i, \beta)$. Since player i and type t_i were chosen arbitrary, the strategy profile β is a BNE as claimed. ■

It is clear that every uniform-perfect BNE is also pointwise-perfect. Now, we provide a game which demonstrates that the converse is not always true. This game has a pointwise-perfect BNE, but it has no uniform-perfect BNE, in fact not even a uniform-perfect strategy profile, although the type and action spaces are compact and the payoff functions are continuous.

Example 11 Consider the following Bayesian game with two players: Player 1 has only one possible type, $T_1 = \{t_1\}$, whereas player 2's type space is $T_2 = \{t_2^1, t_2^2, \ldots, t_2^\infty\}$ in which t_2^∞ is the limit point of the sequence $(t_2^m)_{m=1}^\infty$. The probability measure μ_2 on $(T_2, \mathcal{T}_2)$ is arbitrary. The action spaces are $A_1 = \{U, D\}$ and $A_2 = \{L, R\}$. The payoff matrix when player 1 is the row-player and is given type t_1, and player 2 is the column-player and is given type t_2^m, for every $m \in \mathbb{N}$, is the following:

(t_1, t_2^m)	L	R
U	$0, -\frac{1}{m}$	$0, 0$
D	$1, \frac{1}{m^2}$	$1, 0$

and the payoff matrix when player 1 is given type t_1 and player 2 is given type t_2^∞ is:

(t_1, t_2^∞)	L	R
U	$0, 0$	$0, 0$
D	$1, 0$	$1, 0$

Observe that the type and action spaces in this game are compact, and the payoff functions are continuous.

Define $p_1(t_1) = D$ and $p_2(t_2^m) = L$, for every $m \in \mathbb{N} \cup \{\infty\}$. Let β_1 and β_2 be the deterministic strategies corresponding to p_1 and p_2 respectively. We claim that (β_1, β_2) is a pointwise-perfect BNE, but the game has no uniform-perfect strategy profile, so no uniform-perfect BNE either.

The proof that (β_1, β_2) is a pointwise-perfect BNE: Clearly, (β_1, β_2) is a BNE, so it remains to prove that (β_1, β_2) is pointwise-perfect. For every $k \in \mathbb{N}$, define two completely mixed strategies β_1^k and β_2^k by letting $\beta_1^k(t_1, \cdot) = (1 - \frac{1}{k})\delta_D(\cdot) + \frac{1}{k}\delta_U(\cdot)$ and $\beta_2^k(t_2^m, \cdot) = (1 - \frac{1}{k})\delta_L(\cdot) + \frac{1}{k}\delta_R(\cdot)$ for every $m \in \mathbb{N} \cup \{\infty\}$, where δ_x denotes the Dirac measure on x. The sequence $(\beta^k)_{k=1}^\infty = (\beta_1^k, \beta_2^k)_{k=1}^\infty$ clearly satisfies condition (1) of Definition 9, and it also satisfies condition (2) for player 1 and condition (2) for type t_2^∞ of player 2. It remains to verify condition (2) for an arbitrary type t_2^m for player 2 where $m \in \mathbb{N}$. Take such a type t_2^m. We have

$$\mathbb{E}_{\beta_1^k}(\Pi_2 \mid t_2^m, L) = (1 - \tfrac{1}{k}) \cdot \tfrac{1}{m^2} + \tfrac{1}{k} \cdot (-\tfrac{1}{m}),$$

which is strictly positive for large k. Hence, $BR_2(t_2^m, \beta^k) = \{\delta_L\}$ for large k, which implies that condition (2) of Definition 9 holds for type t_2^m. Thus, (β_1, β_2) is pointwise-perfect indeed.

The proof that this game has no uniform-perfect strategy profile: Suppose by way of contradiction that $\tau = (\tau_1, \tau_2)$ is a uniform-perfect strategy profile. Then, there is a sequence of completely mixed strategy profiles $(\tau^k)_{k=1}^\infty = (\tau_1^k, \tau_2^k)_{k=1}^\infty$ that satisfy conditions (1) and (2) of Definition 7 for τ. First notice that, by the triangle inequality for ρ^w, we have

$$\begin{aligned}\sup_{t_2 \in T_2} \rho^w(\tau_2(t_2, \cdot), BR_2(t_2, \tau^k)) &\leq \sup_{t_2 \in T_2} \left[\, \rho^w(\tau_2^k(t_2, \cdot), \tau_2(t_2, \cdot)) + \rho^w(\tau_2^k(t_2, \cdot), BR_2(t_2, \tau^k)) \,\right] \\ &\leq \sup_{t_2 \in T_2} \rho^w(\tau_2^k(t_2, \cdot), \tau_2(t_2, \cdot)) + \sup_{t_2 \in T_2} \rho^w(\tau_2^k(t_2, \cdot), BR_2(t_2, \tau^k)),\end{aligned}$$

and therefore conditions (1) and (2) imply

$$\lim_{k \to \infty} \sup_{t_2 \in T_2} \rho^w(\tau_2(t_2, \cdot), BR_2(t_2, \tau^k)) = 0. \tag{3}$$

For every $k \in \mathbb{N}$, the probability measure $\tau_1^k(t_1, \cdot)$ can be written in the form $\tau_1^k(t_1, \cdot) = (1 - \varepsilon_k)\delta_U(\cdot) + \varepsilon_k \delta_D(\cdot)$ with some $\varepsilon_k \in (0, 1)$. Due to condition (2) for player 1, we must have that $\varepsilon_k \to 1$ as $k \to \infty$.

For every $k \in \mathbb{N}$ and every type t_2^m for player 2, where $m \in \mathbb{N}$, we have

$$\mathbb{E}_{\tau_1^k}(\Pi_2 \mid t_2^m, L) = \varepsilon_k \cdot \tfrac{1}{m^2} + (1 - \varepsilon_k) \cdot (-\tfrac{1}{m}),$$

whereas $\mathbb{E}_{\tau_1^k}(\Pi_2 \mid t_2^m, R) = 0$. It has two consequences. First, for every t_2^m, where $m \in \mathbb{N}$, we have $BR_2(t_2^m, \tau^k) = \{\delta_L\}$ for large k, which in view of (3) yields $\tau_2(t_2^m, \cdot) = \delta_L(\cdot)$ for every t_2^m, where $m \in \mathbb{N}$. Second, for every $k \in \mathbb{N}$, we have $BR_2(t_2^m, \tau^k) = \{\delta_R\}$ for large m. In conclusion, for every $k \in \mathbb{N}$, if m is large, then $\tau_2(t_2^m, \cdot) = \delta_L(\cdot)$ and $BR_2(t_2^m, \tau^k) = \{\delta_R\}$. This is in contradiction with (3), so τ is not a uniform-perfect strategy profile. ♢

4.3 a.e.-Pointwise-Perfect BNE

Definition 12 *A strategy profile $\beta = (\beta_1, \ldots, \beta_n)$ is called a.e.-pointwise-perfect (where a.e. stands for almost everywhere), if for every player i there exists a set $S_i \in \mathcal{T}_i$ with $\mu_i(S_i) = 0$ and a sequence of profiles of completely mixed strategies $(\beta^k)_{k=1}^{\infty} = (\beta_1^k, \ldots, \beta_n^k)_{k=1}^{\infty}$ with the following properties for every player i and every type $t_i \in T_i \setminus S_i$:*

(1) $\lim_{k\to\infty} \rho^w(\beta_i^k(t_i, \cdot), \beta_i(t_i, \cdot)) = 0,$

(2) $\lim_{k\to\infty} \rho^w(\beta_i^k(t_i, \cdot), BR_i(t_i, \beta^k)) = 0.$

A strategy profile β is called an a.e.-pointwise-perfect BNE if β is a BNE and a.e.-pointwise-perfect.

An a.e.-pointwise-perfect strategy profile, and therefore an a.e.-pointwise-perfect BNE, does not always exist, which is for instance the case in Example 5, because the best response set in this game is always empty. Furthermore, it follows from Example 8 that an a.e.-pointwise-perfect strategy profile is not necessarily a BNE.

It is clear that every pointwise-perfect BNE is also a.e.-pointwise-perfect. The following example proves that the converse is not always true.

Example 13 Consider the following Bayesian game with two players: The type spaces are $T_1 = \{t_1\}$ and $T_2 = \{t_2^1, t_2^2\}$, and μ_2 is given by $\mu_2(t_2^1) = 1$ and $\mu_2(t_2^2) = 0$. The action spaces are $A_1 = \{U, D\}$ and $A_2 = \{L, R\}$. The payoff matrix is the following if player 1 is given type t_1 whereas player 2 is given type t_2^m for $m = 1, 2$:

(t_1, t_2^m)	L	R
U	$0, 1$	$0, 0$
D	$1, 0$	$1, 0$

Define $p_1(t_1) = D$, $p_2(t_2^1) = L$ and $p_2(t_2^2) = R$. Let β_1 and β_2 be the deterministic strategies corresponding to p_1 and p_2 respectively. It is clear that (β_1, β_2) is an a.e.-pointwise-perfect BNE, by choosing $S_1 = \emptyset$ and $S_2 = \{t_2^2\}$ in Definition 12. Yet, (β_1, β_2) is not pointwise-perfect, because player 2 chooses R with probability 1 if he receives type t_2^2.

We remark that a similar example can be made where $T_2 = [0, 1]$, μ_2 is the uniform distribution on T_2, and where the strategy profile is not pointwise-perfect only due to a single type, say type $t_2 = 1$. ♢

The following theorem presents conditions under which there is a strong connection between a.e.-pointwise and pointwise-perfection.

Proposition 14 *Let β be an a.e.-pointwise-perfect BNE in a Bayesian game with compact action spaces. Let S_i, for every player i, and $(\beta^k)_{k=1}^{\infty}$ be as in Definition 12 for β. Suppose that S_i is countable for every player i and $BR_i(t_i, \beta^k)$ is nonempty for every player i, every type $t_i \in S_i$ and every $k \in \mathbb{N}$. Then, there is a pointwise-perfect BNE $\hat{\beta}$ such that $\hat{\beta}_i(t_i, \cdot) = \beta_i(t_i, \cdot)$ holds for every player i and every type $t_i \in T_i \setminus S_i$.*

Proof. For such an a.e.-pointwise-perfect BNE β, we can construct a desired pointwise-perfect BNE $\hat{\beta}$ as follows. For every player $i \in N$, every type $t_i \in S_i$, and every $k \in \mathbb{N}$, by assumption of the proposition, we can choose a probability measure $\sigma^k_{i,t_i} \in BR_i(t_i, \beta^k)$. Now, define a profile of completely mixed strategies $\hat{\beta}^k$, for every $k \in \mathbb{N}$, by letting

$$\hat{\beta}^k_i(t_i, \cdot) = \begin{cases} \beta^k_i(t_i, \cdot) & \text{if } \ t_i \in T_i \setminus S_i \\ (1 - \frac{1}{k})\sigma^k_{i,t_i}(\cdot) + \frac{1}{k}\beta^k_i(t_i, \cdot) & \text{if } \ t_i \in S_i \end{cases}$$

for every player $i \in N$ and type $t_i \in T_i$.

In the remaining part of the proof, we use several fairly known results from analysis, which can all be found in [1]. By assumption, A_i is compact in the topology induced by the metric d_{A_i}, so the set $\Delta(A_i)$ of probability measures on $(A_i, \mathcal{A}_i)$ is also compact with respect to the topology induced by the weak metric. Consequently, the product space $\times_{i \in N} \times_{t_i \in S_i} \Delta(A_i)$ is compact in the product topology by Tychonoff's theorem, and because S_i is countable, this topology is even metrizable. Therefore, this topological space is sequentially compact, which assures the existence of a subsequence $(k_r)_{r=1}^{\infty}$ so that $\sigma^{k_r}_{i,t_i}$ converges to some $\sigma_{i,t_i} \in \Delta(A_i)$, for every player $i \in N$ and every $t_i \in T_i$, with respect to the weak metric. So, define a strategy for every player $i \in N$ by

$$\hat{\beta}_i(t_i, \cdot) = \begin{cases} \beta_i(t_i, \cdot) & \text{if } \ t_i \in T_i \setminus S_i \\ \sigma_{i,t_i}(\cdot) & \text{if } \ t_i \in S_i \end{cases}$$

for every type $t_i \in T_i$. Notice that the profile $\hat{\beta} = (\hat{\beta}_i)_{i \in N}$ is a pointwise-perfect equilibrium. Indeed, for the sequence $\hat{\beta}^{k_r}$, condition (1) of Definition 9 is obviously satisfied, and so is condition (2) because $BR_i(t_i, \hat{\beta}^{k_r}) = BR_i(t_i, \beta^{k_r})$ holds due to $\mu_j(S_j) = 0$ for all players j. ∎

The next example shows that the condition that S_i is countable is crucial in Proposition 14 if there are at least three players. We do not know if this condition is also crucial for games with only two players.

Example 15 Consider the following Bayesian game with three players. The type spaces are $T_1 = \{0, 1\}^{\mathbb{N}}$, $T_2 = \{t_2\}$ and $T_3 = \{t_3\}$. The metric d_{T_1} on T_1 is defined as follows: for $t_1, t'_1 \in T_1$, if $t_1 = t'_1$ then let $d_{T_1}(t_1, t'_1) = 0$, otherwise if m is the first coordinate in which t_1 and t'_1 differ, then let $d_{T_1}(t_1, t'_1) = 2^{-m}$. Notice that d_{T_1} induces the product topology on T_1. Further, let $\mu_1 = \delta_{(1,1,\ldots)}$, i.e. the Dirac measure on the type $(1, 1, \ldots)$. The action spaces are $A_1 = \{U, D\}$,

$A_2 = \{a_1, a_2, \ldots, a_\infty\}$, and $A_3 = \{b_1, b_2, \ldots, b_\infty\}$, in which $a_m = b_m = 1 - \frac{1}{m}$ for every $m \in \mathbb{N}$ and $a_\infty = b_\infty = 1$. For every type $t_1 \in \{0,1\}^{\mathbb{N}}$, let $f_m(t_1)$ be the m-th coordinate of the sequence t_1. When player 1 is given type $t_1 \in \{0,1\}^{\mathbb{N}} \setminus \{1, 1, \ldots\}$, the payoff of player 1 is independent of the action chosen by player 3, and it is given by

$t_1 \in \{0,1\}^{\mathbb{N}} \setminus \{(1,1,\ldots)\}$	a_m	a_∞
U	$f_m(t_1)$	0
D	$1 - f_m(t_1)$	0

whereas if player 1 is given type $t_1 = (1, 1, \ldots)$:

$t_1 = (1,1,\ldots)$	a_m	a_∞
U	0	0
D	0	0

The payoffs of players 2 and 3 are independent of the action chosen by player 1, and are given by:

	b_1	b_2	b_3	b_4	$\ldots$	b_∞
a_1	0,0	0,$\frac{1}{2}$	0,$\frac{1}{3}$	0,$\frac{1}{4}$	$\ldots$	0,0
a_2	$\frac{1}{2}$,0	0,0	0,$\frac{1}{3}$	0,$\frac{1}{4}$	$\ldots$	0,0
a_3	$\frac{1}{3}$,0	$\frac{1}{3}$,0	0,0	0,$\frac{1}{4}$	$\ldots$	0,0
a_4	$\frac{1}{4}$,0	$\frac{1}{4}$,0	$\frac{1}{4}$,0	0,0	$\ldots$	0,0
$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$
a_∞	0,0	0,0	0,0	0,0	$\ldots$	0,0

In general, if player 2 chooses action a_m and player 3 chooses action b_ℓ, with $m, \ell \in \mathbb{N}$, then player 2's payoff is $\frac{1}{m}$ if $m > \ell$, and 0 otherwise, and player 3's payoff is $\frac{1}{\ell}$ if $\ell > m$, and 0 otherwise.

We remark that the type and action spaces in this game are all compact. Since the payoff functions of players 2 and 3 are continuous, and player 1 has only finitely many actions, the best reply set $BR_i(t_i, \beta)$ is nonempty for every player i, every $t_i \in T_i$ and every strategy profile β. We claim: (1) there exists an a.e.-pointwise-perfect equilibrium in this game, but (2) this game admits no pointwise-perfect strategy profile.

An a.e.-pointwise-perfect BNE: First we construct an a.e.-pointwise-perfect BNE. Define $p_1(t_1) = U$ for every $t_1 \in T_1$, $p_2(t_2) = a_\infty$ and $p_3(t_3) = b_\infty$. Let $\widetilde{\beta}_i$ be the deterministic strategy for each player i corresponding to p_i. Now we prove that $\widetilde{\beta} = (\widetilde{\beta}_1, \widetilde{\beta}_2, \widetilde{\beta}_3)$ is an a.e.-pointwise-perfect BNE. Let $S_1 = T_1 \setminus \{(1,1,\ldots)\}$ and $S_2 = S_3 = \emptyset$. Further, take arbitrary completely mixed probability measures σ_2 and σ_3 on $(A_2, \mathcal{A}_2)$ and $(A_3, \mathcal{A}_3)$, respectively. For every $k \in \mathbb{N}$, define the completely mixed strategies

$$\beta_1^k(t_1, \cdot) = (1 - \tfrac{1}{k}) \cdot \delta_U(\cdot) + \tfrac{1}{k} \cdot \delta_D(\cdot) \quad \text{for every } t_1 \in T_1$$

$$\beta_2^k(t_2, \cdot) = (1 - \tfrac{1}{k^2}) \cdot \delta_{a_k}(\cdot) + \tfrac{1}{k^2} \cdot \sigma_2(\cdot)$$

$$\beta_3^k(t_3, \cdot) = (1 - \tfrac{1}{k^2}) \cdot \delta_{b_k}(\cdot) + \tfrac{1}{k^2} \cdot \sigma_3(\cdot).$$

If player 2 plays actions a_{k+1} against $\beta_3^k(t_3, \cdot)$, then he receives at least $(1 - \frac{1}{k^2})\frac{1}{k+1}$, whereas he receives at most $\frac{1}{k+2}$ by playing an action a_ℓ with $\ell > k+1$ and receives at most $\frac{1}{2k^2}$ by playing an action a_ℓ with $\ell < k+1$. Thus, for every $k \geq 3$, action a_{k+1} is player 2's (unique) best response to $\beta_3^k(t_3, \cdot)$, and for similar reasons, action b_{k+1} is player 3's (unique) best response to $\beta_2^k(t_2, \cdot)$. Based on this observation, one can check easily that S_1, S_2, S_3 and the sequence $\beta^k = (\beta_1^k, \beta_2^k, \beta_3^k)$ satisfy the conditions of Definition 12 for $\widetilde{\beta}$. Thus, $\widetilde{\beta}$ is an a.e.-pointwise-perfect BNE indeed.

No pointwise-perfect strategy profile: Now we claim that there is no pointwise-perfect strategy profile in this game. (Note that S_1 above is uncountable, so Proposition 14 does not apply.) Suppose by way of contradiction that $\beta = (\beta_1, \beta_2, \beta_3)$ is a pointwise-perfect strategy profile, with some sequence $(\beta^k)_{k=1}^\infty$ as required in Definition 9.

First we prove that $\beta_2(t_2, \cdot)$ and $\beta_3(t_3, \cdot)$ put probability 1 on action a_∞ and b_∞, respectively, i.e. $\beta_2(t_2, \{a_\infty\}) = \beta_3(t_3, \{b_\infty\}) = 1$. Suppose by way of contradiction that $\beta_2(t_2, \{a_\infty\}) < 1$; the proof is similar if $\beta_3(t_3, \{b_\infty\}) < 1$. Let $m^* = \min\{m \in \mathbb{N} \mid \beta_2(t_2, \{a_m\}) > 0\}$. Since player 3 can get a positive payoff against $\beta_2(t_2, \cdot)$, and $\beta_3(t_3, \cdot)$ is a best response to $\beta_2(t_2, \cdot)$, we must have $\beta_3(t_3, \{b_\infty\}) = 0$. So, also player 2 can get a positive payoff against $\beta_3(t_3, \cdot)$. Let $\ell^* = \min\{\ell \in \mathbb{N} | \beta_3(t_3, \{b_\ell\}) > 0\}$. Now, if $m^* \leq \ell^*$, then action a_{m^*} gives payoff zero to player 2, which is a contradiction as $\beta_2(t_2, \{a_{m^*}\}) > 0$. Similarly, if $\ell^* \leq m^*$, then action b_{ℓ^*} gives payoff zero to player 3, which is a contradiction as $\beta_3(t_3, \{b_{\ell^*}\}) > 0$. So, $\beta_2(t_2, \{a_\infty\}) = \beta_3(t_3, \{b_\infty\}) = 1$ must hold indeed.

For every $k \in \mathbb{N}$, let σ^k be the probability measure on $(A_2, \mathcal{A}_2)$ defined by $\sigma^k(a_\infty) = 0$ and

$$\sigma^k(a_m) = \frac{\beta_2^k(t_2, \{a_m\})}{1 - \beta_2^k(t_2, \{a_\infty\})}$$

for every $m \in \mathbb{N}$. So, $\sigma^k(a_m)$ equals the probability that action a_m is chosen with respect to $\beta_2^k(t_2, \cdot)$ conditioned on the event that a_∞ is not chosen. We now claim that for every $m \in \mathbb{N}$

$$\lim_{k \to \infty} \sigma^k(\{a_1, \ldots, a_m\}) = 0. \tag{4}$$

Suppose by way of contradiction that for some $m \in \mathbb{N}$ there exists a $Z > 0$ and a subsequence $(k_r)_{r=1}^\infty$ such that $\sigma^{k_r}(\{a_1, \ldots, a_m\}) \geq Z$ for every $r \in \mathbb{N}$. Notice that, for every $r \in \mathbb{N}$, action b_{m+1}

gives in expectation at least $\frac{Z}{m+1}$ to player 3 against σ^{k_r}, hence at least $(1 - \beta_2^{k_r}(t_2, \{a_\infty\})) \cdot \frac{Z}{m+1}$ against β^{k_r}. Now consider any action b_w with $w > \frac{m+1}{Z}$. Since player 3's highest payoff for action b_w is $\frac{1}{w}$, action b_w gives player 3 in expectation strictly less than $\frac{Z}{m+1}$ against σ^{k_r}, hence strictly less than $(1 - \beta_2^{k_r}(t_2, \{a_\infty\})) \cdot \frac{Z}{m+1}$ against β^{k_r}. Therefore, player 3's best responses are in $\{b_1, \ldots, b_{\lfloor \frac{m+1}{Z} \rfloor}\}$ against β^{k_r} for any $r \in \mathbb{N}$. Since player 3's best responses along this subsequence are not approaching b_∞ with respect to the weak metric, the sequence β^k cannot satisfy both conditions (1) and (2) of Definition 9. Thus, (4) holds as claimed.

Let $k_1 = 1$ and $m_1 \in \mathbb{N}$ be such that $\sigma^{k_1}(\{a_1, \ldots, a_{m_1}\}) > \frac{1}{2}$. Then, by (4), there exists a $k_2 > k_1$ and an $m_2 > m_1$ such that $\sigma^{k_2}(\{a_{m_1+1}, \ldots, a_{m_2}\}) > \frac{1}{2}$. By repeating this argument, we obtain in $\mathbb{N}$ two strictly increasing sequences $(m_r)_{r=1}^\infty$ and $(k_r)_{r=1}^\infty$ such that for every $r \in \mathbb{N}$

$$\sigma^{k_r}(\{a_m \,|\, m \in W_r\}) > \frac{1}{2}, \tag{5}$$

where $W_r = \{m_{r-1} + 1, \ldots, m_r\}$ and $m_0 = 0$. Note that the sets W_r, $r \in \mathbb{N}$, form a partition of $\mathbb{N}$. Now let t_1 be the type in $\{0,1\}^{\mathbb{N}}$ such that $f_m(t_1) = 1$ if $m \in W_r$ for an odd $r \in \mathbb{N}$ and $f_m(t_1) = 0$ if $m \in W_r$ for an even $r \in \mathbb{N}$. Take an arbitrary odd $r \in \mathbb{N}$. Then, $f_m(t_1) = 1$ for any $m \in W_r$, and since $\sigma^{k_r}(\{a_m \,|\, m \in W_r\}) > \frac{1}{2}$ due to (5), player 1's unique best response to β^{k_r} is action U. Similarly, when $r \in \mathbb{N}$ is even, player 1's unique best response to β^{k_r} is action D. So, the sequence β^k cannot satisfy both conditions (1) and (2) of Definition 9 for type t_1, which is a contradiction. Consequently, there is no pointwise-perfect strategy profile in this game. $\diamondsuit$

The following corollary follows easily from Proposition 14.

Corollary 16 *Consider a Bayesian game in which the type spaces are countable, the action spaces are compact, and the payoff functions are continuous. Then, for every a.e.-pointwise-perfect BNE β, there exists a pointwise-perfect BNE $\hat{\beta}$ such that, for every player i, we have $\hat{\beta}_i(t_i, \cdot) = \beta_i(t_i, \cdot)$ for μ_i-a.e. type $t_i \in T_i$.*

Now, we examine Bayesian games which satisfy the following two measurability conditions:

- Condition M1: For any player i and any strategies β_i^1, β_i^2 of player i, the map

  $$t_i \mapsto \rho^w(\beta_i^1(t_i, \cdot), \beta_i^2(t_i, \cdot))$$

  is measurable.

- Condition M2: For any player i, any strategy β_i of player i, and any strategy profile τ, the map

  $$t_i \mapsto \rho^w(\beta_i(t_i, \cdot), BR_i(t_i, \tau))$$

  is measurable.

Condition M1 is mild, because if the action spaces are σ-compact, then M1 is always satisfied (cf. Lemma 29). It is not clear if mild conditions are also sufficient to guarantee M2.

Now, we present a number of conditions that are all equivalent to a.e.-pointwise-perfection, under the additional assumptions of M1 and M2.

Proposition 17 *Consider a Bayesian game which satisfies conditions M1 and M2. Then, for every strategy profile $\beta = (\beta_1, \ldots, \beta_n)$, equivalent are:*

i. The strategy profile β is a.e.-pointwise-perfect.

ii. There exists a sequence of profiles of completely mixed strategies $(\beta^k)_{k=1}^{\infty} = (\beta_1^k, \ldots, \beta_n^k)_{k=1}^{\infty}$ such that for every player i and every $\varepsilon > 0$:

$$(ii.1) \quad \lim_{k\to\infty} \mu_i \left\{ t_i \in T_i \mid \rho^w(\beta_i^k(t_i,\cdot), \beta_i(t_i,\cdot)) \geq \varepsilon \right\} = 0,$$

$$(ii.2) \quad \lim_{k\to\infty} \mu_i \left\{ t_i \in T_i \mid \rho^w(\beta_i^k(t_i,\cdot), BR_i(t_i,\beta^k)) \geq \varepsilon \right\} = 0.$$

iii. There exists a sequence of profiles of completely mixed strategies $(\beta^k)_{k=1}^{\infty} = (\beta_1^k, \ldots, \beta_n^k)_{k=1}^{\infty}$ such that for every player i:

$$(iii.1) \quad \lim_{k\to\infty} \int_{T_i} \rho^w(\beta_i^k(t_i,\cdot), \beta_i(t_i,\cdot)) \, \mu_i(dt_i) = 0,$$

$$(iii.2) \quad \lim_{k\to\infty} \int_{T_i} \rho^w(\beta_i^k(t_i,\cdot), BR_i(t_i,\beta^k)) \, \mu_i(dt_i) = 0.$$

The intuition behind conditions (ii) and (iii) is the following. For every player i and every $k \in \mathbb{N}$, define a function $X_i^k : T_i \to \mathbb{R}$ by $X_i^k(t_i) = \rho^w\left(\beta_i^k(t_i,\cdot), \beta_i(t_i,\cdot)\right)$, and a function $Y_i^k : T_i \to \mathbb{R}$ by $Y_i^k(t_i) = \rho^w\left(\beta_i^k(t_i,\cdot), BR_i(t_i,\beta^k)\right)$. Due to the assumptions M1 and M2, these functions are measurable, and therefore they are random variables on the measurable space $(T_i, \mathcal{T}_i)$. Conditions (ii.1) and (ii.2) respectively require that both sequences X_i^k and Y_i^k converge in probability to zero (i.e. to the random variable that is zero everywhere). So, the probability that X_i^k and Y_i^k are far from zero becomes negligible for large k. Conditions (iii.1) and (iii.2) mean that both sequences X_i^k and Y_i^k converge in expectation to zero, with respect to the distribution μ_i. This makes intuitive sense because the type of player i is drawn from T_i according to μ_i.

Proof. The implication (i)→(ii) follows immediately from the fact that pointwise convergence almost everywhere implies convergence in probability. The implication (ii)→(iii) is also valid, because the sequences X_i^k and Y_i^k are uniformly bounded and hence if they converge in probability to zero, then they also converge to zero in expectation. So as to prove the implication (iii)→(i), we argue that both implications (iii)→(ii) and (ii)→(i) are valid. The implication (iii)→(ii) holds because convergence in expectation implies convergence in probability. Finally, the implication (ii)→(i) follows from the fact that convergence in probability implies the existence of a subsequence that pointwisely converges to the same limit almost everywhere. Hence, if the sequences X_i^k and Y_i^k converge in probability to zero, then there are subsequences $X_i^{k_m}$ and $Y_i^{k_m}$ which satisfy the conditions of Definition 12. ■

The equivalence of the conditions of Proposition 17 is remarkable in view of the fact that pointwise convergence almost everywhere, convergence in probability and convergence in expectation are not equivalent in general.

4.4 Perfection in Finite Bayesian Games

In this subsection, we discuss some direct consequences of our previous results on Bayesian games in which there are only finitely many types and actions. First of all, it is clear that the uniform and the pointwise approaches coincide. The a.e.-pointwise approach can however lead to different results, but only in the - perhaps less interesting - case when certain types occur with probability 0.

Theorem 18 *Every Bayesian game with finitely many types and actions admits a pointwise-perfect (or equivalently, uniform-perfect) BNE.*

Proof. We only provide a sketch of the proof. First, consider the corresponding game G' with ex-ante probabilities in which each player i has action set $A_i' = \times_{t_i \in T_i} A_i$ and wants to maximize his ex-ante expected payoff, i.e. his expected payoff before he learns his own type. The game G' is a game with complete information and with finite action spaces, so it has a perfect equilibrium σ in the classical sense with a corresponding sequence of completely mixed strategies $(\sigma^k)_{k=1}^{\infty}$. Let β be the (unique) strategy in the original game G such that, given any type $t_i \in T_i$, the strategies σ and β induce the same probability on every action $a_i \in A_i$, i.e. $\beta(t_i, a_i) = \sigma(A_i'[t_i, a_i])$, where $A_i'[t_i, a_i]$ is the set of all members of A_i' whose coordinate at position t_i is exactly a_i. Define β^k for every $k \in \mathbb{N}$ in a similar way with regard to σ^k. It is not difficult to see that β is an a.e.-pointwise-perfect BNE, with the sequence $(\beta^k)_{k=1}^{\infty}$ and with S_i being for every player i the set of all types $t_i \in T_i$ that occur with probability 0. The reason is that if a type $t_i \in T_i$ occurs with a positive probability, then this type is also taken into account by the ex-ante approach. Now, Proposition 14 completes the proof. ■

5 A second-price auction with incomplete information

We examine perfect equilibrium in a second price auction with incomplete information, in which bidders have no dominant strategies. In this auction we identify a symmetric BNE β, and argue that

β is a natural candidate solution for the auction from a conceptual and normative perspective. The BNE β separates the types, and it is differentiable–so that bids increase smoothly with type. We show that β is unique in this respect, there is no other strategy pair that features these properties. Moreover, as uncertainty vanishes, the BNE β converges to the classical truthful dominant strategy equilibrium in the Vickrey auction.

We show that perfect equilibrium is a useful tool to select the BNE β in this auction. In particular, we show that β is uniform-perfect. Next, we present a few classes of other BNEs for the auction. These BNEs are arguably less intuitive in the sense that types (at least partially) pool their bids. We show that perfection eliminates these equilibria, so that the BNE β is uniquely selected by perfection from a large class of BNEs. Finally we argue that the usual standard selection criteria sincere bidding and dominant strategy equilibrium, have no bite in this context, and do not single out the BNE β. As the proofs in this section are technical in nature, we only discuss the main results, and defer the formal proofs to the appendix.

Consider the following sealed-bid second price auction Γ^{α} for a single indivisible object, where α is a parameter in $(0,1)$. There are two bidders, whose respective types t_1 and t_2 are drawn independently from $T_1 = T_2 = [0,1]$ according to the uniform distribution. The valuations of the bidders are symmetric and are given by $v_1(t_1,t_2) = 5 + t_1 - \alpha t_2$ and $v_2(t_1,t_2) = 5 + t_2 - \alpha t_1$. The set of available bids is $A_1 = A_2 = [5-\alpha, 6]$. The tie-breaking rule can be arbitrary, and plays no role in our analysis.

5.1 A symmetric BNE

For each bidder i, define the pure strategy $B_i(t_i) = 5 + (1-\alpha)t_i$, and denote by β_i the corresponding deterministic strategy. In this subsection we argue that, from a conceptual and normative point of view, the strategy pair $\beta = (\beta_1, \beta_2)$ is a natural candidate solution for the auction.

Concretely, we have the following claims. The strategy pair β is a symmetric BNE, it separates the types, and it is differentiable–so that bids increase smoothly with type. It is also unique in this, there is no other strategy pair that features all these properties. Moreover, as uncertainty vanishes, so $\alpha \to 0$, the BNE β converges to the classical truthful dominant strategy equilibrium in the Vickrey auction.

We first observe that β is a BNE. The proof of this observation is deferred to the Appendix.

Proposition 19 *The symmetric deterministic strategy profile $\beta = (\beta_1, \beta_2)$ is a BNE.*

Next, we show that β is the unique BNE within the class of differentiable separating pure strategy pairs. Let $\mathfrak{F}$ be the class of deterministic strategies in which the corresponding pure strategies

$p_i : T_i \to A_i$ for bidders $i = 1, 2$ satisfy the following conditions:

(i) $p_1(0) = p_2(0)$ and $p_1(1) = p_2(1)$,

(ii) p_i is differentiable and $\frac{d}{dt_i} p_i(t_i) > 0$ for every $t_i \in T_i$.

Thus, $\mathfrak{F}$ is the class of pure and differentiable strategies for which the bid is strictly increasing in type. In particular, elements of $\mathfrak{F}$ separate the types of a player in the sense that the type can be inferred from the bid since no two types make the same bid [4]. We have the following proposition.

Proposition 20 *The BNE β is the only pure BNE in class $\mathfrak{F}$.*

Thus, the BNE β is the unique differentiable separating BNE. Moreover, as we already noted, when uncertainty vanishes (so, when $\alpha \to 0$) β converges to the truthful dominant strategy equilibrium in the Vickrey auction.

5.2 The selective power of perfect BNE

Thus, β is a natural candidate as a solution for the auction we study. We show that perfect BNE does select the differentiable separating BNE β, from among many equilibria in the auction. We first show that β is a perfect BNE.

Proposition 21 *The BNE β is uniform-perfect.*

We describe the idea of the proof. Let $C_M : [0, 1] \to [5 - \alpha, 6]$ be defined by $C_M(r) = 5 + (1 - \alpha)r$ for all $r \in [0, 1]$ and μ be the uniform distribution on $[5 - \alpha, 6]$. The first idea would be to look at the sequences of completely mixed strategies $\hat{\beta}_i^k$ given by $\hat{\beta}_i^k(t_i, \cdot) = (1 - \varepsilon_k)\delta_{C_M(t_i)}(\cdot) + \varepsilon_k \mu(\cdot)$ for both players $i = 1, 2$ and for every t_i, where δ is the Dirac measure and ε_k is a sequence in $(0, 1)$ converging to zero. One can verify that the sequence $\hat{\beta}_i^k$ satisfies condition (1) of Definition 7, and $\lim_{k\to\infty} \rho^w(\hat{\beta}_i^k(t_i, \cdot), BR_i(t_i, \hat{\beta}^k)) = 0$ for every $t_i \in (0, 1)$. However, the latter does not hold for $t_i = 0, 1$. More precisely, one can easily check that $BR_i(0, \hat{\beta}^k) = \{5 - \frac{1}{2}\alpha\}$ and $BR_i(1, \hat{\beta}^k) = \{6 - \frac{1}{2}\alpha\}$, for every k. However, by defining the completely mixed strategies in a more delicate way, we can prove that (β_1, β_2) is uniform-perfect. We introduce two more curves $C_T, C_L : [0, 1] \to [5 - \alpha, 6]$ by $C_T(r) = 6 - \alpha r^2$, and $C_L(r) = 5 - \alpha + \alpha(1 - r)^2$. These curves are depicted in the picture below.

[4] The first condition is mainly a normalization.

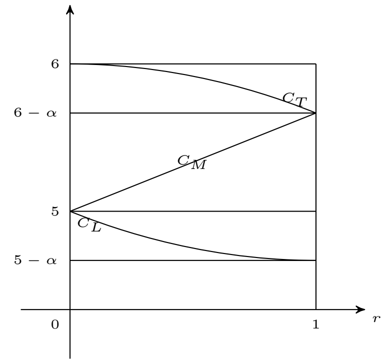

The curves C_T and C_L will be used to repair the above mentioned problem for types 0 and 1. For every $k \in \mathbb{N}$, let $\varepsilon_k = \frac{1}{k+3}$ and $\beta_1^k : T_1 \times \mathcal{A}_1 \to [0,1]$ be the completely mixed strategy for bidder 1 which is given for every $t_1 \in T_1 = [0,1]$ by

$$\beta_1^k(t_1, \cdot) = (1 - 2\varepsilon_k - \varepsilon_k^2)\ \delta_{C_M(t_1)}(\cdot) + \varepsilon_k\ \delta_{C_T(t_1)}(\cdot) + \varepsilon_k\ \delta_{C_L(t_1)}(\cdot) + \varepsilon_k^2\ \mu(\cdot).$$

The cumulative probability distribution with respect to $\beta_1^k(t_1, \cdot)$ is shown on $A_1 = [5 - \alpha, 6]$ in the picture below.

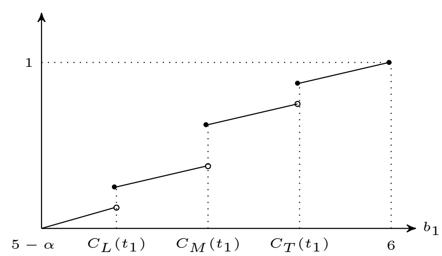

Similarly, define a completely mixed strategy β_2^k for bidder 2. In the Appendix we will show that the sequence of strategy profiles $(\beta^k)_{k=1}^{\infty} = (\beta_1^k, \beta_2^k)_{k=1}^{\infty}$ satisfies conditions 1 and 2 of Definition 7. Hence, the BNE β is uniform-perfect.

Next, we present a few classes of other BNEs for the auction. They are arguably less intuitive in the sense that types (at least partially) pool their bids. We show that perfection eliminates these equilibria. We start with Wolf and Sheep BNEs.

Proposition 22 *The deterministic strategy profile $\sigma = (\sigma_1, \sigma_2)$, given by $\sigma_1(t_1, \cdot) = \delta_6(\cdot)$ for every $t_1 \in T_1$ and $\sigma_2(t_2, \cdot) = \delta_{5-\alpha}(\cdot)$ for every $t_2 \in T_2$, is a BNE. However, σ is not pointwise-perfect, and hence not uniform-perfect either.*

It is worthwhile to note that the BNE introduced in Proposition 22 is just one of the many other BNEs of the same type. In all of this type of BNEs, the lowest bid of Wolf player is strictly larger than the highest bid of Sheep player.

The following proposition considers a deterministic BNE in which the corresponding pure strategies, as functions of the types, are not continuous. Also on this type of BNE we could construct many variations.

Proposition 23 *The deterministic strategy profile $\eta = (\eta_1, \eta_2)$ is a BNE, where η_1 and η_2 correspond to the pure strategies b_1 and b_2 given by*

$$b_1(t_1) = \begin{cases} 5 & \text{if } t_1 \in [0, x(\alpha)] \\ 6 & \text{if } t_1 \in (x(\alpha), 1] \end{cases} \quad \textit{and} \quad b_2(t_2) = \begin{cases} 5-\alpha & \text{if } t_2 \in [0, y(\alpha)] \\ 5.1 & \text{if } t_2 \in (y(\alpha), 1] \end{cases},$$

with $x(\alpha) = \frac{0.4+2\alpha}{4-\alpha^2}$ and $y(\alpha) = \frac{\alpha(0.2+\alpha)}{4-\alpha^2}$.
However, η is not pointwise-perfect, and hence not uniform-perfect either.

Thus, perfection selects the equilibrium β from among a large class of equilibria. Finally we argue that other standard selection criteria such as sincere bidding and dominance do not have a bite when uncertainty is present (that is, when $\alpha > 0$). In particular, these selection criteria do not single out the separating BNE β in this example.

Proposition 24 *There does not exist a symmetric, deterministic, and strictly increasing BNE in which bidders bid sincere given the opponent's bid function. In particular, the BNE β is not sincere, in the sense that for each bidder i, there are types t_i for which $B_i(t_i)$ does not equal his expected evaluation of the object given t_i.*

Proposition 25 *Neither bidder has a dominant strategy in the auction Γ^α in the following sense: neither bidder i has a strategy β_i such that for every type t_i, every strategy σ_i of bidder i and every strategy β_{-i} of bidder i's opponents we have*

$$\mathbb{E}_{\beta_{-i}}(\Pi_i \mid t_i, \beta_i) \geq \mathbb{E}_{\beta_{-i}}(\Pi_i \mid t_i, \sigma_i).$$

In fact in the Appendix we prove a somewhat stronger statement, namely that in the above statement we could replace "every type t_i" by "every type t_i in a subset of types with a strictly positive measure".

Hence, the usual standard selection criteria sincere bidding and dominant strategy equilibrium, have no bite in the context of Bayesian games, and do not single out the BNE β.

6 Appendix I: Semi-product measures

Definition 26 Let (X, Σ_1) and (Y, Σ_2) be two measurable spaces, μ be a measure on (X, Σ_1) and $\tau : X \times \Sigma_2 \to [0, \infty)$ be a bounded function with the following properties :

1. the section function $\tau(x, \cdot) : \Sigma_2 \to [0, \infty)$ is a measure on (Y, Σ_2) for every $x \in X$,

2. the section function $\tau(\cdot, B) : X \to [0, \infty)$ is a measurable function for every $B \in \Sigma_2$.

Let λ be a measure on the semi-ring $\Sigma_1 \times \Sigma_2$ defined by

$$\lambda(A \times B) = \int_A \tau(x, B)\, \mu(dx),$$

for every $A \times B \in \Sigma_1 \times \Sigma_2$. The unique extension of the measure λ by the formula (1) to the product σ-field $\Sigma_1 \otimes \Sigma_2$ is denoted by λ^* and is called the *semi-product measure* corresponding to τ and μ.

Note that λ is well-defined, because the section function $\tau(\cdot, B)$ is bounded and measurable for every $B \in \Sigma_2$. The following theorem states the integral respect to λ^* as iterated integrals respect to τ and μ. Dudley states a sketch of the proof in [7].

Theorem 27 *Let (X, Σ_1) and (Y, Σ_2) be two measurable spaces and let λ^* be the semi-product measure corresponding to τ and μ. Then, for every function $f : X \times Y \to \mathbb{R}$ bounded and measurable with respect to the σ-field $\Sigma_1 \otimes \Sigma_2$ we have*

$$\int_{X \times Y} f\, d\lambda^* = \int_X \int_Y f\, \tau(x, dy)\mu(dx).$$

In particular, $\lambda^*(E) = \int_X \tau(x, E_x)\mu(dx)$ *for every* $E \in \Sigma_1 \otimes \Sigma_2$ *where* $E_x = \{y \in Y \mid (x, y) \in E\}$.

Proof. Let $E \in \Sigma_1 \otimes \Sigma_2$. We first show that $\lambda^*(E) = \int_X \tau(x, E_x)\mu(dx)$. Let $\mathcal{N}$ be the collection of all sets $E \in \Sigma_1 \otimes \Sigma_2$ such that
(i). $x \mapsto \tau(x, E_x)$ is a measurable function, and
(ii). $\lambda^*(E) = \int_X \tau(x, E_x)\mu(dx)$.
We prove that $\mathcal{N}$ is a σ-field.

1. It is clear that $X \times Y \in \mathcal{N}$. Indeed, $x \mapsto \tau(x, Y)$ is measurable according to condition (2) of Definition 26, and $\lambda^*(X \times Y) = \lambda(X \times Y) = \int_X \tau(x, Y)\mu(dx)$.
2. We show that if $E \in \mathcal{N}$, then $E^c \in \mathcal{N}$:
(i). Since $x \mapsto \tau(x, E_x)$ is measurable and $\tau(x, Y) - \tau(x, E_x) = \tau(x, E_x^c)$, the map $x \mapsto \tau(x, E_x^c)$ is also measurable.

$$\begin{aligned}(ii).\ \lambda^*(E^c) = \lambda^*(X \times Y) - \lambda^*(E) &= \int_X \tau(x, Y)\, \mu(dx) - \int_X \tau(x, E_x)\mu(dx) \\ &= \int_X (\tau(x, Y) - \tau(x, E_x))\, \mu(dx) \\ &= \int_X \tau(x, Y \backslash E_x)\mu(dx) \\ &= \int_X \tau(x, E_x^c)\mu(dx).\end{aligned}$$

3. We show that if $(E^i)_{i=1}^{\infty}$ is a countable collection of sets in $\mathcal{N}$, then $\bigcup_{i=1}^{\infty} E^i \in \mathcal{N}$:

Without loss of generality we assume that $E_1, E_2, \ldots$ are pairwise disjoint, because otherwise we can consider the pairwise disjoint sets $E_1, E_2 \backslash E_1, E_3 \backslash (E_1 \cup E_2), \ldots$.

(i). Since $x \mapsto \tau(x, E^i_x)$ is measurable for every $i \in \mathbb{N}$, by the Monotone Convergence Theorem the map that assigns to x the real number $\lim_{n\to\infty} \sum_{i=1}^{n} \tau(x, E^i_x) = \sum_{i=1}^{\infty} \tau(x, E^i_x) = \tau(x, \bigcup_{i=1}^{\infty} E^i_x)$ is also measurable.

(ii). According to the Monotone Convergence Theorem we have

$$\begin{aligned}
\lambda^*(\bigcup_{i=1}^{\infty} E^i) = \sum_{i=1}^{\infty} \lambda^*(E^i) &= \lim_{n\to\infty} \sum_{i=1}^{n} \lambda^*(E^i) \\
&= \lim_{n\to\infty} \sum_{i=1}^{n} \int_X \tau(x, E^i_x)\ \mu(dx) \\
&= \int_X \lim_{n\to\infty} \sum_{i=1}^{n} \tau(x, E^i_x)\ \mu(dx) \\
&= \int_X \tau(x, \bigcup_{i=1}^{\infty} E^i_x)\ \mu(dx) \\
&= \int_X \tau(x, (\bigcup_{i=1}^{\infty} E^i)_x)\ \mu(dx).
\end{aligned}$$

Hence, $\mathcal{N}$ is a σ-field. Furthermore, it is clear that for every $A \in \Sigma_1$ and $B \in \Sigma_2$, $A \times B \in \mathcal{N}$. Since the product σ-field $\Sigma_1 \otimes \Sigma_2$ is the smallest σ-field that contains $\Sigma_1 \times \Sigma_2$, we have $\Sigma_1 \otimes \Sigma_2 = \mathcal{N}$. Therefore, $\lambda^*(E) = \int_X \tau(x, E_x)\ \mu(dx)$ for every set $E \in \Sigma_1 \otimes \Sigma_2$.

Now, we complete the proof of the theorem. We just proved that for every set $E \in \Sigma_1 \otimes \Sigma_2$ we have

$$\int_{X\times Y} \mathbf{1}_E\ d\lambda^* = \lambda^*(E) = \int_X \tau(x, E_x)\ \mu(dx) = \int_X \int_Y \mathbf{1}_{E_x}\ \tau(x, dy)\ \mu(dx).$$

Hence, it is clear that for any bounded and measurable step function ϕ on $X \times Y$ we have

$$\int_{X\times Y} \phi\ d\lambda^* = \int_X \int_Y \phi_x\ \tau(x, dy)\mu(dx) = \int_X \int_Y \phi\ \tau(x, dy)\mu(dx),$$

where $\phi_x(y) = \phi(x, y)$. Let $f : X \times Y \to \mathbb{R}$ be a bounded and measurable function with respect to the σ-field $\Sigma_1 \otimes \Sigma_2$. Then, there is a sequence of bounded and measurable step functions $\{\phi^n\}_{n=1}^{\infty}$

such that $\phi^n \uparrow f$ pointwisely. Thus, by the Monotone Convergence Theorem we have

$$\begin{aligned}
\int_{X\times Y} f\ d\lambda^* &= \int_{X\times Y} \lim_{n\to\infty} \phi^n\ d\lambda^* \\
&= \lim_{n\to\infty} \int_{X\times Y} \phi^n\ d\lambda^* \\
&= \lim_{n\to\infty} \int_X \int_Y \phi^n\ \tau(x,dy)\mu(dx) \\
&= \int_X \int_Y \lim_{n\to\infty} \phi^n\ \tau(x,dy)\mu(dx) \\
&= \int_X \int_Y f\ \tau(x,dy)\mu(dx).
\end{aligned}$$

■

7 Appendix II: Proofs for the auction in Section 5

Throughout the Appendix.2, we will use the notion $\mathbb{E}_{\tau_1}(\Pi_2 \mid t_2, b_2)$ for bidder 2's expected profit given his type t_2 and bid b_2 against a strategy τ_1 of bidder 1. If bidder 1 makes a bid $b_1 > b_2$, then bidder 2 does not win the object and has zero profit. Therefore, if bidding b_2 against τ_1 leads to a tie with zero ex-ante probability, then we have

$$\mathbb{E}_{\tau_1}(\Pi_2 \mid t_2, b_2) = \int_0^1 \int_{5-\alpha}^{b_2} (5 + t_2 - \alpha t_1 - b_1)\ \tau_1(t_1, db_1)\ dt_1. \tag{6}$$

Since the function $5 + t_2 - \alpha t_1 - b_1$ is bounded on $[0,1] \times [5-\alpha, 6]$ and measurable with respect to $\mathcal{T} \times \mathcal{A}$, then according to Theorem 27, $\mathbb{E}_{\tau_1}(\Pi_2 \mid t_2, b_2)$ is well-defined.

In the special case when τ_1 is a deterministic strategy, i.e. when $\tau_1(t_1, \cdot) = \delta_{p_1(t_1)}(\cdot)$ for a pure strategy p_1, we have

$$\int_{5-\alpha}^{b_2} (5 + t_2 - \alpha t_1 - b_1)\ \tau_1(t_1, db_1) = \begin{cases} 5 + t_2 - \alpha t_1 - p_1(t_1) & \text{if } p_1(t_1) \in [5-\alpha, b_2] \\ 0 & \text{if } p_1(t_1) \in (b_2, 6]. \end{cases}$$

Hence, if p_1 is increasing and $b_2 \in p_1(T_1)$, then we have

$$\mathbb{E}_{\tau_1}(\Pi_2 \mid t_2, b_2) = \int_0^{p_1^{-1}(b_2)} (5 + t_2 - \alpha t_1 - p_1(t_1))\, dt_1. \tag{7}$$

7.1 Proof of Proposition 19

We prove that the deterministic strategy profile $\beta = (\beta_1, \beta_2)$ in which β_i corresponds to the pure strategy $B_i(t_i) = 5 + (1-\alpha)t_i$ for each bidder i, is a BNE. For this purpose, we prove that β_2 is a best response against β_1. Then, due to symmetry, β_1 is also a best response against β_2, and the proof will be complete.

According to (7), if $b_2 \in [5, 6-\alpha]$ then we have

$$\begin{aligned}
\mathbb{E}_{\beta_1}(\Pi_2 \mid t_2, b_2) &= \int_0^{B_1^{-1}(b_2)} (5 + t_2 - \alpha t_1 - B_1(t_1))\, dt_1 \\
&= \int_0^{\frac{b_2-5}{1-\alpha}} (5 + t_2 - \alpha t_1 - 5 - (1-\alpha)t_1)\, dt_1 \\
&= \int_0^{\frac{b_2-5}{1-\alpha}} (t_2 - t_1)\, dt_1 \\
&= t_2 \left(\frac{b_2-5}{1-\alpha}\right) - \frac{1}{2}\left(\frac{b_2-5}{1-\alpha}\right)^2 .
\end{aligned}$$

By taking the first derivative with respect to b_2

$$\frac{d}{db_2}\mathbb{E}_{\beta_1}(\Pi_2 \mid t_2, b_2) = \frac{t_2}{1-\alpha} - \frac{1}{1-\alpha}\left(\frac{b_2-5}{1-\alpha}\right).$$

It is clear that

$$\frac{d}{db_2}\mathbb{E}_{\beta_1}(\Pi_2 \mid t_2, 5) \geq 0 \quad \text{and} \quad \frac{d}{db_2}\mathbb{E}_{\beta_1}(\Pi_2 \mid t_2, 6-\alpha) \leq 0,$$

which implies that the maximum of $\mathbb{E}_{\beta_1}(\Pi_2 \mid t_2, b_2)$ over $b_2 \in [5, 6-\alpha]$ is attained where $\frac{d}{db_2}\mathbb{E}_{\beta_1}(\Pi_2 \mid t_2, b_2) = 0$. This happens exactly when

$$b_2 = 5 + (1-\alpha)t_2 = B_2(t_2).$$

Notice that, for bidder 2, bidding less than 5 is never better than bidding 5 exactly, because all bids less than or equal to 5 win against β_1 with probability 0. Similarly, bidding more than $6-\alpha$ is never better than bidding $6-\alpha$ exactly, because all bids larger than or equal to $6-\alpha$ win against β_1 with probability 1. Therefore, β_2 is a best response to β_1 as claimed.

7.2 Proof of Proposition 20

Consider two pure strategies p_1 and p_2 that satisfy conditions (i), (ii) of Proposition 20, and the corresponding deterministic strategies τ_1 and τ_2. Suppose that (τ_1, τ_2) is a BNE. Our goal is to show that $p_1(t_1) = 5 + (1-\alpha)t_1$ and $p_2(t_2) = 5 + (1-\alpha)t_2$.

Notice that $p_1 : T_1 \to A_1$ and $p_2 : T_2 \to A_2$ are continuous and invertible, and therefore they have a continuous inverse. Thus, because $p_1(T_1) = p_2(T_2)$, the function $\hat{p} = p_1^{-1} \circ p_2$ is well defined and it is a continuous bijection from T_2 to T_1.

First we argue that $p_2(0) \leq 5$. Suppose by way of contradiction that $p_2(0) > 5$. According to (7) we have

$$\mathbb{E}_{\tau_1}(\Pi_2 \mid t_2, p_2(t_2)) = \int_0^{\hat{p}(t_2)} (5 + t_2 - \alpha t_1 - p_1(t_1))\; dt_1.$$

Notice that if $t_2 > 0$ then $\hat{p}(t_2) > 0$ due to conditions (i) and (ii), and moreover, for every type $t_2 < p_2(0) - 5$ we have

$$5 + t_2 - \alpha t_1 - p_1(t_1) \leq 5 + t_2 - p_1(t_1) \leq 5 + t_2 - p_1(0) = 5 + t_2 - p_2(0) < 0.$$

Therefore, $\mathbb{E}_{\tau_1}(\Pi_2 \mid t_2, p_2(t_2))$ is strictly negative. However, as $p_1(0) = p_2(0) > 5$, bidder 2 can get zero by bidding 5 for instance, which never wins. This is a contradiction with the assumption that τ_2 is a best response to τ_1, so $p_2(0) \leq 5$ holds indeed.

Now we prove that $p_2(1) \geq 6 - \alpha$. Write $c := 6 - \alpha - p_2(1)$ and suppose by way of contradiction that $c > 0$. According to (7), if we compare bids $p_2(1)$ and $p_2(t_2)$ for bidder 2 when his type is t_2, we obtain

$$\mathbb{E}_{\tau_1}(\Pi_2 \mid t_2, p_2(1)) - \mathbb{E}_{\tau_1}(\Pi_2 \mid t_2, p_2(t_2)) \quad = \quad \int_{\hat{p}(t_2)}^{1} (5 + t_2 - \alpha t_1 - p_1(t_1)) \; dt_1.$$

Notice that if $t_2 \in (1 - c, 1)$, then $\hat{p}(t_2) < 1$ due to conditions (i) and (ii), and moreover, for every $t_1 \in [0, 1]$ we have

$$\begin{aligned} 5 + t_2 - \alpha t_1 - p_1(t_1) \quad &\geq \quad 5 + t_2 - \alpha - p_1(1) \\ &= \quad 5 + t_2 - \alpha - p_2(1) \\ &= \quad 5 + t_2 - \alpha + c - 6 + \alpha \\ &= \quad -1 + t_2 + c \\ &> \quad 0. \end{aligned}$$

Therefore, for every $t_2 \in (1 - c, 1)$

$$\mathbb{E}_{\tau_1}(\Pi_2 \mid t_2, p_2(1)) - \mathbb{E}_{\tau_1}(\Pi_2 \mid t_2, p_2(t_2)) > 0,$$

which is in contradiction with the assumption that τ_2 is a best response to τ_1. Hence, $p_2(1) \geq 6 - \alpha$ holds indeed.

Now, we calculate bidder 2's best response bids against τ_1, given his type t_2. It is clear that, for any type t_2, bidding $p_2(0)$ is not worse than any bid in $[5 - \alpha, p_2(0)]$, because all these bids win with probability zero against τ_1. Similarly, bidding $p_2(1)$ is not worse than any bid in $[p_2(1), 6]$, because all these bids win with probability 1 against τ_1. It remains to determine the best bids in the interval $[p_2(0), p_2(1)]$, which we do by examining the derivative of $\mathbb{E}_{\tau_1}(\Pi_2 \mid t_2, b_2)$ with respect to b_2. By (7), we have for every bid $b_2 \in [p_2(0), p_2(1)]$ that

$$\begin{aligned}\mathbb{E}_{\tau_1}(\Pi_2 \mid t_2, b_2) &= \int_0^{p_1^{-1}(b_2)} (5 + t_2 - \alpha t_1 - p_1(t_1))\, dt_1 \\ &= \Big[(5+t_2)t_1 - \frac{\alpha}{2}t_1^2\Big]_0^{p_1^{-1}(b_2)} - \int_0^{p_1^{-1}(b_2)} p_1(t_1)\, dt_1 \\ &= (5+t_2)p_1^{-1}(b_2) - \frac{\alpha}{2}\left(p_1^{-1}(b_2)\right)^2 - \int_0^{p_1^{-1}(b_2)} p_1(t_1)\, dt_1.\end{aligned}$$

Therefore, the first derivative with respect to b_2 is

$$\begin{aligned}\frac{d}{db_2}\mathbb{E}_{\tau_1}(\Pi_2 \mid t_2, b_2) &= \frac{1}{p_1'\left(p_1^{-1}(b_2)\right)} \cdot \left(5 + t_2 - \alpha p_1^{-1}(b_2)\right) - \frac{1}{p_1'\left(p_1^{-1}(b_2)\right)} p_1(p_1^{-1}(b_2)) \\ &= \frac{1}{p_1'\left(p_1^{-1}(b_2)\right)} \cdot \left(5 + t_2 - \alpha p_1^{-1}(b_2) - b_2\right).\end{aligned}$$

Because $\hat{p}(0) = p_1^{-1}(p_2(0)) = p_1^{-1}(p_1(0)) = 0$ and $p_1'\left(p_1^{-1}(b_2)\right) > 0$ by conditions (i) and (ii), and because $p_2(0) \le 5$ we have for every $t_2 > 0$ that

$$\frac{d}{db_2}\mathbb{E}_{\tau_1}(\Pi_2 \mid t_2, p_2(0)) > 0. \tag{8}$$

Similarly, because $\hat{p}(1) = p_1^{-1}(p_2(1)) = p_1^{-1}(p_1(1)) = 1$ and $p_1'\left(p_1^{-1}(b_2)\right) > 0$ by conditions (i) and (ii), and because $p_2(1) \ge 6 - \alpha$ we have for every $t_2 < 1$ that

$$\frac{d}{db_2}\mathbb{E}_{\tau_1}(\Pi_2 \mid t_2, p_2(1)) < 0. \tag{9}$$

In view of (8) and (9), for all $t_2 \in (0,1)$, each best response bid b_2 in $[p_2(0), p_2(1)]$ has to satisfy

$$\frac{d}{db_2}\mathbb{E}_{\tau_1}(\Pi_2 \mid t_2, b_2) = 0,$$

which is equivalent to

$$b_2 = 5 + t_2 - \alpha p_1^{-1}(b_2).$$

Because τ_2 is a best response against τ_1, we must have for every $t_2 \in (0,1)$ that

$$p_2(t_2) = 5 + t_2 - \alpha p_1^{-1}(p_2(t_2)) = 5 + t_2 - \alpha\hat{p}(t_2).$$

By continuity of p_2 and $\hat{p}$, we obtain for all $t_2 \in [0,1]$ that

$$p_2(t_2) = 5 + t_2 - \alpha\hat{p}(t_2). \tag{10}$$

With a similar argument for bidder 1, we have for all $t_1 \in [0,1]$ that

$$p_1(t_1) = 5 + t_1 - \alpha\hat{p}^{-1}(t_1). \tag{11}$$

By equation (11) we have for all $t_2 \in [0,1]$ that

$$p_1\left(\hat{p}(t_2)\right) = 5 + \hat{p}(t_2) - \alpha\hat{p}^{-1}(\hat{p}(t_2)),$$

which yields $p_2(t_2) = 5 + \hat{p}(t_2) - \alpha t_2$. Thus, by equation (10) we have

$$5 + \hat{p}(t_2) - \alpha t_2 = 5 + t_2 - \alpha \hat{p}(t_2),$$

which implies $\hat{p}(t_2) = t_2$. Hence, from (10) we derive that

$$p_2(t_2) = 5 + (1 - \alpha)t_2,$$

and because of symmetry also that

$$p_1(t_1) = 5 + (1 - \alpha)t_1,$$

so the proof of Proposition 20 is complete.

7.3 Proof of Proposition 24

Suppose that (b_1, b_2) is symmetric and deterministic, $b_1 = b_2$ is strictly increasing, and b_1 is sincere given b_2. Since $b_1 = b_2$ by symmetry, and since $b_1 = b_2$ is strictly increasing, we have that

$$\mathbb{E}(t_2 \mid b_1(t_1) > b_2(t_2)) = \mathbb{E}(t_2 \mid t_1 > t_2) = \frac{1}{2}t_1.$$

So, since b_1 is sincere given b_2, using the above equality we have that

$$b_1(t_1) = \mathbb{E}(v_1(t_1, t_2) \mid b_1(t_1) > b_2(t_2)) = \mathbb{E}(5 + t_1 - \alpha t_2 \mid b_1(t_1) > b_2(t_2)) = 5 + t_1 - \frac{\alpha}{2}t_1.$$

So, $b_1(t_1) = 5 + t_1 - \frac{\alpha}{2}t_1$ and $b_2(t_2) = 5 + t_2 - \frac{\alpha}{2}t_2$. This implies that $(b_1, b_2) \in \mathfrak{F}$. So, since $b_1(t_1) > B_1(t_1)$ for every type $t_1 > 0$, Proposition 20 shows that (b_1, b_2) does not constitute a BNE.

It is worthwhile to mention that even if we look at the situation in which bidder 1 naively bids his ex ante expected valuation, the same conclusion holds. In that case bidder 1's expected evaluation of the object given his type t_1 is

$$\mathbb{E}(v_1(t_1, t_2) \mid t_1) = \mathbb{E}(5 + t_1 - \alpha t_2 \mid t_1) = 5 + t_1 - \frac{\alpha}{2},$$

and the same reasoning as before shows that also this does not yield a BNE.

7.4 Proof of Proposition 25

We prove that neither bidder has a dominant strategy. To this end we prove that every strategy σ_i of player $i = 1, 2$ is not dominant. Suppose σ_2 is a dominant strategy for bidder 2. Notice that if bidder 1 chooses strategy β_1 and if bidder 2 is given type $t_2 \in (0, 1)$, then bidding β_2 is bidder 2's unique best response (see the proof of Proposition 19). This implies that $\sigma_2 = \beta_2$. Now we show that the strategies β_2 is not in dominant strategies. To this end we show that β_2 is not a best response against all strategies of bidder 1. Suppose that bidder 1's strategy is given by $\lambda_1(t_1,\cdot) = \delta_{b_1(t_1)}(\cdot)$ for every $t_1 \in T_1$, where

$$b_1(t_1) = \begin{cases} \frac{11}{2} & \text{if } t_1 \in [0, \frac{1}{2}] \\ 6 & \text{if } t_1 \in (\frac{1}{2}, 1]. \end{cases}$$

We prove that if $t_2 = \frac{1+\alpha}{2}$, then β_2 is not a best response for bidder 2 against λ_1. Notice that bidder 2 gets zero by bidding $B_2(\frac{1+\alpha}{2})$, because it never wins against λ_1 due to $B_2(\frac{1+\alpha}{2}) = 5 + (1-\alpha)(\frac{1+\alpha}{2}) < \frac{11}{2}$. However, by (7), bidder 2's expected profit given type $t_2 = \frac{1+\alpha}{2}$ and $b_2 \in (\frac{11}{2}, 6)$ is

$$\mathbb{E}_{\lambda_1}(\Pi_2 \mid \frac{1+\alpha}{2}, b_2) = \int_0^{\frac{1}{2}} \left(5 + \frac{1+\alpha}{2} - \alpha t_1 - \frac{11}{2}\right) dt_1 = \frac{\alpha}{8} > 0.$$

This shows that if $t_2 = \frac{1+\alpha}{2}$ then bidding $b_2 \in (\frac{11}{2}, 6)$ is strictly better than bidding $B_2(\frac{1+\alpha}{2})$. Hence, β_2 is not in dominant strategies. This proves that bidder 2 does not have a dominant strategy at all, and by symmetry the same is true for bidder 1. The proof is complete.

7.5 Proof of Proposition 21

We prove that the BNE $\beta = (\beta_1, \beta_2)$ is uniform-perfect. For this purpose, define the curves $C_M, C_T, C_L : [0,1] \to [5-\alpha, 6]$ by

$$C_M(t_1) = 5 + (1-\alpha)t_1,$$

$$C_T(t_1) = 6 - \alpha t_1^2,$$

$$C_L(t_1) = 5 - \alpha + \alpha(1-t_1)^2.$$

Let μ be the uniform distribution on $A_1 = [5-\alpha, 6]$. For every $k \in \mathbb{N}$, let $\varepsilon_k = \frac{1}{k+3}$ and $\beta_1^k : T_1 \times \mathcal{A}_1 \to [0,1]$ be the completely mixed strategy for bidder 1 which is given for every $t_1 \in T_1 = [0,1]$ by

$$\beta_1^k(t_1,\cdot) = (1 - 2\varepsilon_k - \varepsilon_k^2)\ \delta_{C_M(t_1)}(\cdot) + \varepsilon_k\ \delta_{C_T(t_1)}(\cdot) + \varepsilon_k\ \delta_{C_L(t_1)}(\cdot) + \varepsilon_k^2\ \mu(\cdot).$$

Similarly, define a completely mixed strategy β_2^k for bidder 2. We show that the sequence of strategy profiles $(\beta^k)_{k=1}^{\infty} = (\beta_1^k, \beta_2^k)_{k=1}^{\infty}$ satisfies conditions 1 and 2 of Definition 7.
It is clear from the definition of $\beta_i(t_i,\cdot)$ and $\beta_i^k(t_i,\cdot)$, for each bidder i, that

$$\sup_{t_i \in T_i} \rho^w(\beta_i^k(t_i,\cdot), \beta_i(t_i,\cdot)) = 2\varepsilon_k + \varepsilon_k^2.$$

Hence, condition 1 of Definition 7 is satisfied. Condition 2 of Definition 7 follows from the claim below.

Claim: For every $\xi \in (0,\alpha)$, there is a $K_\xi \in \mathbb{N}$ such that for every $k > K_\xi$ we have:

1. For every $t_2 \in [0, \xi)$, the maximum of $\mathbb{E}_{\beta_1^k}(\Pi_2 \mid t_2, b_2)$ over $b_2 \in [5-\alpha, 6]$ can only be attained within $b_2 \in [5-\xi, 5+2\xi(1-\alpha)]$.

2. For every $t_2 \in [\xi, 1-\xi]$, the maximum of $\mathbb{E}_{\beta_1^k}(\Pi_2 \mid t_2, b_2)$ over $b_2 \in [5-\alpha, 6]$ can only be attained within $b_2 \in [C_M(t_2) - \xi(1-\alpha), C_M(t_2) + \xi(1-\alpha)]$.

3. For every $t_2 \in (1-\xi, 1]$ the maximum of $\mathbb{E}_{\beta_1^k}(\Pi_2 \mid t_2, b_2)$ over $b_2 \in [5-\alpha, 6]$ can only be attained within $b_2 \in [6-\alpha-2\xi(1-\alpha), 6-\alpha+\xi]$.

The above claim is illustrated in Figure 1. According to the claim, the gray area includes all best response bids for bidder 2 if $k > K_\xi$.

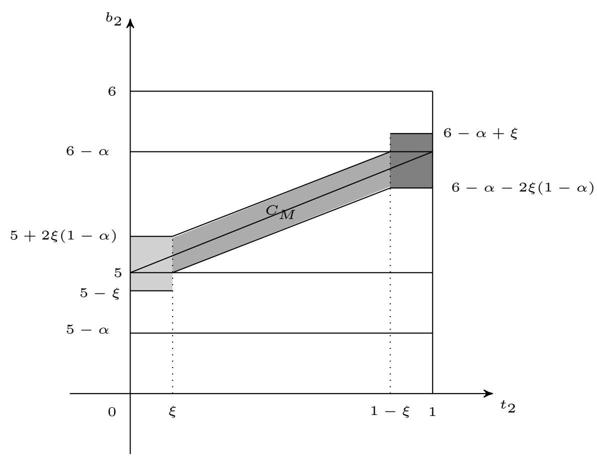

Figure 1: bidder 2's best response bids

Take an arbitrary $\xi \in (0, \min\{\alpha, \frac{1}{2}\})$. To prove this claim, it is enough to show that the following claims 1, 2 and 3 are valid. These claims corresponds to three cases, i.e., $6-\alpha \leq b_2 \leq 6$, $5-\alpha \leq b_2 \leq 5$ and $5 \leq b_2 \leq 6-\alpha$.

Claim 1: for the case $6-\alpha \leq b_2 \leq 6$
There is an $M_\xi^1 \in \mathbb{N}$ such that for every $k > M_\xi^1$ we have:

1. For every $t_2 \in [0, 1-\xi]$, the maximum of $\mathbb{E}_{\beta_1^k}(\Pi_2 \mid t_2, b_2)$ over $b_2 \in [6-\alpha, 6]$ is attained at $b_2 = 6-\alpha$.

2. For every $t_2 \in (1-\xi, 1]$, the maximum of $\mathbb{E}_{\beta_1^k}(\Pi_2 \mid t_2, b_2)$ over $b_2 \in [6-\alpha, 6]$ can only be attained within $b_2 \in [6-\alpha, 6-\alpha+\xi]$.

Claim 2: for the case $5-\alpha \leq b_2 \leq 5$
There is an $M_\xi^2 \in \mathbb{N}$ such that for every $k > M_\xi^2$ we have:

1. For every $t_2 \in [\xi, 1]$, the maximum of $\mathbb{E}_{\beta_1^k}(\Pi_2 \mid t_2, b_2)$ over $b_2 \in [5-\alpha, 5]$ is attained at $b_2 = 5$.

2. For every $t_2 \in [0, \xi)$, the maximum of $\mathbb{E}_{\beta_1^k}(\Pi_2 \mid t_2, b_2)$ over $b_2 \in [5-\alpha, 5]$ can only be attained within $b_2 \in [5-\xi, 5]$.

Claim 3: for case $5 \leq b_2 \leq 6-\alpha$
There is an $M_\xi^3 \in \mathbb{N}$ such that for every $k > M_\xi^3$ we have:

1. For every $t_2 \in [\xi, 1-\xi]$, the maximum of $\mathbb{E}_{\beta_1^k}(\Pi_2 \mid t_2, b_2)$ over $b_2 \in [5, 6-\alpha]$ can only be attained within $b_2 \in [C_M(t_2) - \xi(1-\alpha), C_M(t_2) + \xi(1-\alpha)]$.

2. For every $t_2 \in [0, \xi)$, the maximum of $\mathbb{E}_{\beta_1^k}(\Pi_2 \mid t_2, b_2)$ over $b_2 \in [5, 6-\alpha]$ can only be attained within $b_2 \in [5, 5+2\xi(1-\alpha)]$.

3. For every $t_2 \in (1-\xi, 1]$ the maximum of $\mathbb{E}_{\beta_1^k}(\Pi_2 \mid t_2, b_2)$ over $b_2 \in [5, 6-\alpha]$ can only be attained within $b_2 \in [6-\alpha-2\xi(1-\alpha), 6-\alpha]$.

It is important to note that Part (1) of Claim 1 shows that for every $t_2 \in [0, 1-\xi]$, the maximum of $\mathbb{E}_{\beta_1^k}(\Pi_2 \mid t_2, b_2)$ over $b_2 \in [6-\alpha, 6]$ is not more than the maximum over $b_2 \in [5, 6-\alpha]$, which is verified in Claim 3. Similarly, Part (1) of Claim 2 indicates that for every $t_2 \in [\xi, 1]$, the maximum of $\mathbb{E}_{\beta_1^k}(\Pi_2 \mid t_2, b_2)$ over $b_2 \in [5-\alpha, 5]$ is not more than the maximum over $b_2 \in [5, 6-\alpha]$.

Proof of Claim 1: We compute $\mathbb{E}_{\beta_1^k}(\Pi_2 \mid t_2, b_2)$ for every $b_2 \in [6-\alpha, 6]$. Note that $C_T^{-1}(b_2) \in [0,1]$ is the unique type t_1 for bidder 1 such that $C_T(t_1) = b_2$. By the definition of C_T, we have $C_T^{-1}(b_2) = \sqrt{\frac{6-b_2}{\alpha}}$. Therefore, by (6)

$$\begin{aligned}
\mathbb{E}_{\beta_1^k}(\Pi_2 \mid t_2, b_2) &= \int_0^1 \int_{5-\alpha}^{b_2} (5 + t_2 - \alpha t_1 - b_1)\ \beta_1^k(t_1, db_1)\, dt_1 \\
&= \int_0^{C_T^{-1}(b_2)} \int_{5-\alpha}^{b_2} (5 + t_2 - \alpha t_1 - b_1)\ \beta_1^k(t_1, db_1)\, dt_1 \qquad \text{I} \\
&+ \int_{C_T^{-1}(b_2)}^1 \int_{5-\alpha}^{b_2} (5 + t_2 - \alpha t_1 - b_1)\ \beta_1^k(t_1, db_1)\, dt_1. \qquad \text{II}
\end{aligned}$$

From the definition of $\beta_1^k(t_1, \cdot)$ we obtain

$$\begin{aligned}
\text{I} &= (1 - 2\varepsilon_k - \varepsilon_k^2) \int_0^{C_T^{-1}(b_2)} (5 + t_2 - \alpha t_1 - C_M(t_1))\ dt_1 \\
&+ \varepsilon_k \int_0^{C_T^{-1}(b_2)} (5 + t_2 - \alpha t_1 - C_L(t_1))\ dt_1 \\
&+ \varepsilon_k^2 \int_0^{C_T^{-1}(b_2)} \int_{5-\alpha}^{b_2} (5 + t_2 - \alpha t_1 - b_1)\, \mu(db_1)\, dt_1.
\end{aligned}$$

Similarly, we have

$$\begin{aligned}
\text{II} &= (1-2\varepsilon_k-\varepsilon_k^2)\int_{C_T^{-1}(b_2)}^{1}(5+t_2-\alpha t_1-C_M(t_1))\ dt_1\\
&+ \varepsilon_k\int_{C_T^{-1}(b_2)}^{1}(5+t_2-\alpha t_1-C_L(t_1))\ dt_1\\
&+ \varepsilon_k^2\int_{C_T^{-1}(b_2)}^{1}\int_{5-\alpha}^{b_2}(5+t_2-\alpha t_1-b_1)\,\mu(db_1)\ dt_1\\
&+ \varepsilon_k\int_{C_T^{-1}(b_2)}^{1}(5+t_2-\alpha t_1-C_T(t_1))\ dt_1.
\end{aligned}$$

Hence, by the definition of the curves C_M, C_T, and C_L we obtain

$$\begin{aligned}
\mathbb{E}_{\beta_1^k}(\Pi_2 \mid t_2,b_2)=\text{I}+\text{II} &= (1-2\varepsilon_k-\varepsilon_k^2)\int_0^1(5+t_2-\alpha t_1-5-(1-\alpha)t_1)\ dt_1\\
&+ \varepsilon_k\int_0^1\left(5+t_2-\alpha t_1-5+\alpha-\alpha(1-t_1)^2\right)\ dt_1\\
&+ \varepsilon_k^2\int_0^1\int_{5-\alpha}^{b_2}(5+t_2-\alpha t_1-b_1)\,\mu(db_1)\ dt_1\\
&+ \varepsilon_k\int_{C_T^{-1}(b_2)}^{1}\left(5+t_2-\alpha t_1-6+\alpha t_1^2\right)\ dt_1.
\end{aligned}$$

Thus,

$$\begin{aligned}
\mathbb{E}_{\beta_1^k}(\Pi_2 \mid t_2,b_2) &= (1-2\varepsilon_k-\varepsilon_k^2)(t_2-\frac{1}{2})+\varepsilon_k(t_2+\frac{\alpha}{6})\\
&+ \frac{\varepsilon_k^2}{1+\alpha}\int_{5-\alpha}^{b_2}(5+t_2-\frac{\alpha}{2}-b_1)\ db_1\\
&+ \varepsilon_k\int_{C_T^{-1}(b_2)}^{1}\left(-1+t_2-\alpha t_1+\alpha t_1^2\right)\ dt_1.
\end{aligned}$$

Note that for every $t_2\in[0,1]$

$$\begin{aligned}
\mathbb{E}_{\beta_1^k}(\Pi_2 \mid t_2,6-\alpha)-\mathbb{E}_{\beta_1^k}(\Pi_2 \mid t_2,6) &= -\frac{\varepsilon_k^2}{1+\alpha}\int_{6-\alpha}^{6}(5+t_2-\frac{\alpha}{2}-b_1)\ db_1\\
&- \varepsilon_k\int_0^1\left(-1+t_2-\alpha t_1+\alpha t_1^2\right)\ dt_1\\
&= \left(\frac{\alpha}{1+\alpha}\right)\varepsilon_k^2(1-t_2)+\varepsilon_k(1-t_2+\frac{\alpha}{6})>0,
\end{aligned}$$

and also that $\mathbb{E}_{\beta_1^k}(\Pi_2 \mid t_2,b_2)$ is continuous on the compact set $(t_2,b_2)\in[0,1]\times[5-\alpha,6]$. Therefore, there is a $y>0$ (without loss of generality we can assume that $y\in(0,\alpha-\xi)$, in order to make the analysis easier) such that for every $t_2\in[0,1]$ and $b_2\in(6-y,6]$

$$\mathbb{E}_{\beta_1^k}(\Pi_2 \mid t_2,6-\alpha)>\mathbb{E}_{\beta_1^k}(\Pi_2 \mid t_2,b_2).$$

This means that for every $t_2\in[0,1]$, bidding $6-\alpha$ is strictly better than any bid in the interval $(6-y,6]$. Hence, it suffices to prove Claim 1 over $b_2\in[6-\alpha,6-y]$. We do so by examining the first derivative of $\mathbb{E}_{\beta_1^k}(\Pi_2 \mid t_2, b_2)$ with respect to b_2, which equals to

$$\frac{d}{db_2}\, \mathbb{E}_{\beta_1^k}(\Pi_2 \mid t_2, b_2) \quad = \quad \varepsilon_k^2 \cdot q_1(t_2, b_2) + \varepsilon_k \cdot q_2(t_2, b_2),$$

where

$$\begin{aligned} q_1(t_2, b_2) &= \frac{1}{1+\alpha} \cdot (5 + t_2 - \frac{\alpha}{2} - b_2) \\ q_2(t_2, b_2) &= \frac{1}{2\alpha \cdot C_T^{-1}(b_2)} \cdot \left(-1 + t_2 - \alpha C_T^{-1}(b_2) + \alpha (C_T^{-1}(b_2))^2\right). \end{aligned}$$

Note that $q_2(t_2, b_2)$ is strictly negative for all $(t_2, b_2) \in [0,1] \times [6-\alpha, 6-y] \setminus \{(1, 6-\alpha)\}$, and $q_2(1, 6-\alpha) = 0$.

Because q_2 is strictly negative on the compact set $[0, 1-\xi] \times [6-\alpha, 6-y]$, it has a strictly negative upper-bound on $[0, 1-\xi] \times [6-\alpha, 6-y]$. Therefore, since ε_k^2 is relatively much smaller than ε_k for sufficiently large k, then $\frac{d}{db_2}\, \mathbb{E}_{\beta_1^k}(\Pi_2 \mid t_2, b_2) < 0$ for all $(t_2, b_2) \in [0, 1-\xi] \times [6-\alpha, 6-y]$. This proves that there exists an $N_\xi^1 \in \mathbb{N}$ such that for every $k > N_\xi^1$ we have $\frac{d}{db_2}\, \mathbb{E}_{\beta_1^k}(\Pi_2 \mid t_2, b_2) < 0$ for all $(t_2, b_2) \in [0, 1-\xi] \times [6-\alpha, 6-y]$. Hence, part (1) of Claim 1 holds for all $k > N_\xi^1$.

Because q_2 is strictly negative on the compact set $[1-\xi, 1] \times [6-\alpha+\xi, 6-y]$, we can apply a similar argument and find an $L_\xi^1 \in \mathbb{N}$ such that part (2) of Claim 1 holds for every $k > L_\xi^1$. Therefore, by choosing $M_\xi^1 = \max\{N_\xi^1, L_\xi^1\}$, Claim 1 is valid for every $k > M_\xi^1$.

Proof of Claim 2: We compute $\mathbb{E}_{\beta_1^k}(\Pi_2 \mid t_2, b_2)$ for every $b_2 \in [5-\alpha, 5]$. Note that type $C_L^{-1}(b_2) \in [0,1]$ is the unique type t_1 for bidder 1 such that $C_L(t_1) = b_2$. By the definition of C_L, we have $C_L^{-1}(b_2) = 1 - \sqrt{\frac{b_2-5}{\alpha} + 1}$. Hence, by (6)

$$\begin{aligned} \mathbb{E}_{\beta_1^k}(\Pi_2 \mid t_2, b_2) &= \int_0^1 \int_{5-\alpha}^{b_2} (5 + t_2 - \alpha t_1 - b_1)\ \beta_1^k(t_1, db_1)\, dt_1 \\ &= \int_0^{C_L^{-1}(b_2)} \int_{5-\alpha}^{b_2} (5 + t_2 - \alpha t_1 - b_1)\ \beta_1^k(t_1, db_1)\, dt_1 \qquad \text{III} \\ &+ \int_{C_L^{-1}(b_2)}^1 \int_{5-\alpha}^{b_2} (5 + t_2 - \alpha t_1 - b_1)\ \beta_1^k(t_1, db_1)\, dt_1. \qquad \text{IV} \end{aligned}$$

Therefore, we have

$$\text{III} \quad = \quad \varepsilon_k^2 \int_0^{C_L^{-1}(b_2)} \int_{5-\alpha}^{b_2} (5 + t_2 - \alpha t_1 - b_1)\, \mu(db_1)\, dt_1,$$

$$\begin{aligned} \text{IV} \quad &= \quad \varepsilon_k \int_{C_L^{-1}(b_2)}^1 (5 + t_2 - \alpha t_1 - C_L(t_1))\ dt_1 \\ &+ \quad \varepsilon_k^2 \int_{C_L^{-1}(b_2)}^1 \int_{5-\alpha}^{b_2} (5 + t_2 - \alpha t_1 - b_1)\ \mu(db_1)\, dt_1. \end{aligned}$$

Hence, by the definition of the curve C_L we obtain

$$\begin{aligned}
\mathbb{E}_{\beta_1^k}(\Pi_2 \mid t_2, b_2) = \text{ III} + \text{IV} \quad &= \quad \varepsilon_k^2 \int_0^1 \int_{5-\alpha}^{b_2} (5 + t_2 - \alpha t_1 - b_1)\ \mu(db_1)\ dt_1 \\
&+ \quad \varepsilon_k \int_{C_L^{-1}(b_2)}^1 \left(5 + t_2 - \alpha t_1 - 5 + \alpha - \alpha(1-t_1)^2\right)\ dt_1 \\
&= \quad \frac{\varepsilon_k^2}{1+\alpha} \int_{5-\alpha}^{b_2} (5 + t_2 - \frac{\alpha}{2} - b_1)\ db_1 \\
&+ \quad \varepsilon_k \int_{C_L^{-1}(b_2)}^1 (t_2 + \alpha t_1 - \alpha t_1^2)\ dt_1.
\end{aligned}$$

Note that for every $t_2 \in [0,1]$

$$\begin{aligned}
\mathbb{E}_{\beta_1^k}(\Pi_2 \mid t_2, 5) - \mathbb{E}_{\beta_1^k}(\Pi_2 \mid t_2, 5-\alpha) \quad &= \quad \frac{\varepsilon_k^2}{1+\alpha} \int_{5-\alpha}^{5} (5 + t_2 - \frac{\alpha}{2} - b_1)\ db_1 \\
&+ \quad \varepsilon_k \int_0^1 (t_2 + \alpha t_1 - \alpha t_1^2)\ dt_1 \\
&= \quad \left(\frac{\alpha}{1+\alpha}\right) \varepsilon_k^2 t_2 + \varepsilon_k (t_2 + \frac{\alpha}{6}) > 0,
\end{aligned}$$

and also that $\mathbb{E}_{\beta_1^k}(\Pi_2 \mid t_2, b_2)$ is continuous on the compact set $(t_2, b_2) \in [0,1] \times [5-\alpha, 5]$. Therefore, there is a $y > 0$ (similar to the previous case assume that $y \in (0, \alpha - \xi)$) such that for every $t_2 \in [0,1]$ and $b_2 \in [5-\alpha, 5-\alpha+y)$

$$\mathbb{E}_{\beta_1^k}(\Pi_2 \mid t_2, 5) > \mathbb{E}_{\beta_1^k}(\Pi_2 \mid t_2, b_2).$$

This means that for every $t_2 \in [0,1]$, bidding 5 is strictly better than any bid in the interval $[5-\alpha, 5-\alpha+y)$. Hence, it suffices to prove Claim 2 over the interval $[5-\alpha+y, 5]$. We do so by examining the first derivative of $\mathbb{E}_{\beta_1^k}(\Pi_2 \mid t_2, b_2)$ with respect to b_2, which equals to

$$\frac{d}{db_2} \mathbb{E}_{\beta_1^k}(\Pi_2 \mid t_2, b_2) \quad = \quad \varepsilon_k^2 p_1(t_2, b_2) + \varepsilon_k p_2(t_2, b_2),$$

where

$$\begin{aligned}
p_1(t_2, b_2) \quad &= \quad \frac{1}{1+\alpha} \cdot (5 + t_2 - \frac{\alpha}{2} - b_2), \\
p_2(t_2, b_2) \quad &= \quad \frac{1}{2\alpha \cdot (1 - C_L^{-1}(b_2))} \cdot (t_2 + \alpha C_L^{-1}(b_2) 0 - \alpha C_L^{-1}(b_2)^2).
\end{aligned}$$

Note that $p_2(t_2, b_2)$ is strictly positive for all $(t_2, b_2) \in [0,1] \times [5-\alpha+y, 5] \setminus \{(0,5)\}$, and $p_2(0,5) = 0$. With an argument similar to that of the previous case, we can find an M_ξ^2 such that for every $k > M_\xi^2$ Claim 2 is valid.

Proof of Claim 3: We compute $\mathbb{E}_{\beta_1^k}(\Pi_2 \mid t_2, b_2)$ for every $b_2 \in [5, 6-\alpha]$. Note that $C_M^{-1}(b_2) = \frac{b_2 - 5}{1-\alpha}$

for every $b_2 \in [5, 6-\alpha]$. With a similar argument as before we have

$$\begin{aligned}
\mathbb{E}_{\beta_1^k}(\Pi_2 \mid t_2, b_2) &= (1-2\varepsilon_k-\varepsilon_k^2)\int_0^{\frac{b_2-5}{1-\alpha}} (5+t_2-\alpha t_1 - C_M(t_1))\ dt_1 \\
&+ \varepsilon_k \int_0^1 (5+t_2-\alpha t_1 - C_L(t_1))\ dt_1 \\
&+ \varepsilon_k^2 \int_0^1 \int_{5-\alpha}^{b_2} (5+t_2-\alpha t_1 - b_1)\ \mu(db_1)\ dt_1.
\end{aligned}$$

Hence, by the definition of the curves C_M and C_L we obtain

$$\begin{aligned}
\mathbb{E}_{\beta_1^k}(\Pi_2 \mid t_2, b_2) &= (1-2\varepsilon_k-\varepsilon_k^2)\int_0^{\frac{b_2-5}{1-\alpha}} (5+t_2-\alpha t_1 - 5-(1-\alpha)t_1)\ dt_1 \\
&+ \varepsilon_k \int_0^1 \left(5+t_2-\alpha t_1 - 5+\alpha-\alpha(1-t_1)^2\right)\ dt_1 \\
&+ \varepsilon_k^2 \int_0^1 \int_{5-\alpha}^{b_2} (5+t_2-\alpha t_1 - b_1)\ \mu(db_1)\ dt_1.
\end{aligned}$$

Thus,

$$\begin{aligned}
\mathbb{E}_{\beta_1^k}(\Pi_2 \mid t_2, b_2) &= (1-2\varepsilon_k-\varepsilon_k^2)\int_0^{\frac{b_2-5}{1-\alpha}} (t_2 - t_1)\ dt_1 \\
&+ \varepsilon_k \int_0^1 \left(t_2+\alpha t_1 - \alpha t_1^2\right)\ dt_1 \\
&+ \frac{\varepsilon_k^2}{1+\alpha} \int_{5-\alpha}^{b_2} (5+t_2-\frac{\alpha}{2} - b_1)\ db_1 \\
&= (1-2\varepsilon_k-\varepsilon_k^2)\left(t_2\left(\frac{b_2-5}{1-\alpha}\right)-\frac{1}{2}\left(\frac{b_2-5}{1-\alpha}\right)^2\right) \\
&+ \varepsilon_k(t_2+\frac{\alpha}{6}) \\
&+ \frac{\varepsilon_k^2}{2(1+\alpha)}(b_2-5+\alpha)(5+2t_2-b_2).
\end{aligned}$$

So, the first derivative of $\mathbb{E}_{\beta_1^k}(\Pi_2 \mid t_2, b_2)$ with respect to b_2 is

$$\begin{aligned}
&\frac{d}{db_2}\mathbb{E}_{\beta_1^k}(\Pi_2 \mid t_2, b_2) \\
= \ &\frac{1-2\varepsilon_k-\varepsilon_k^2}{1-\alpha}\cdot\left(t_2-\frac{b_2-5}{1-\alpha}\right)+\frac{\varepsilon_k^2}{1+\alpha}\ (5+t_2-\frac{\alpha}{2}-b_2). \qquad (12)
\end{aligned}$$

First, we prove that for every $t_2 \in [\xi, 1-\xi]$ the maximum of $\mathbb{E}_{\beta_1^k}(\Pi_2 \mid t_2, b_2)$ cannot be attained on the boundary of interval $[5, 6-\alpha]$. Observe that

$$\frac{d}{db_2}\mathbb{E}_{\beta_1^k}(\Pi_2 \mid t_2, 5) = \frac{1-2\varepsilon_k-\varepsilon_k^2}{1-\alpha}\cdot t_2 + \frac{\varepsilon_k^2}{1+\alpha}\ (t_2-\frac{\alpha}{2}),$$

and

$$\frac{d}{db_2}\mathbb{E}_{\beta_1^k}(\Pi_2 \mid t_2, 6-\alpha) = \frac{1-2\varepsilon_k-\varepsilon_k^2}{1-\alpha}\cdot (t_2-1) + \frac{\varepsilon_k^2}{1+\alpha}\ (t_2-1+\frac{\alpha}{2}).$$

Hence, there is an $N_\xi^3 \in \mathbb{N}$ such that if $k > N_\xi^3$ then for every $t_2 \in [\xi, 1-\xi]$ we have

$$\frac{d}{db_2}\mathbb{E}_{\beta_1^k}(\Pi_2 \mid t_2, 5) > 0$$

and

$$\frac{d}{db_2}\mathbb{E}_{\beta_1^k}(\Pi_2 \mid t_2, 6-\alpha) < 0.$$

Therefore, by the continuity of $\mathbb{E}_{\beta_1^k}(\Pi_2 \mid t_2, b_2)$, we can conclude for every $t_2 \in [\xi, 1-\xi]$ that if b_2 maximizes $\mathbb{E}_{\beta_1^k}(\Pi_2 \mid t_2, b_2)$, then

$$\frac{d}{db_2}\mathbb{E}_{\beta_1^k}(\Pi_2 \mid t_2, b_2) = 0. \tag{13}$$

One can check with the help of (12) that the unique solution of (13), for every $t_2 \in [0,1]$, is

$$b_2(t_2) = \frac{(1-2\varepsilon_k-\varepsilon_k^2)(1+\alpha)}{v(\varepsilon_k)} \cdot (5+(1-\alpha)t_2) + \frac{\varepsilon_k^2(1-\alpha)^2}{v(\varepsilon_k)}(5+t_2-\frac{\alpha}{2}),$$

where

$$v(\varepsilon_k) = (1-2\varepsilon_k-\varepsilon_k^2)(1+\alpha) + \varepsilon_k^2(1-\alpha)^2.$$

One can verify that

$$|b_2(t_2) - C_M(t_2)| = \frac{\alpha\varepsilon_k^2(1-\alpha)^2}{v(\varepsilon_k)}\left|t_2 - \frac{1}{2}\right| \le \frac{\alpha\varepsilon_k^2(1-\alpha)^2}{2v(\varepsilon_k)}.$$

Therefore, there is an $L_\xi^3 \in \mathbb{N}$ such that for every $k > L_\xi^3$ and every $t_2 \in [0,1]$ we have

$$|b_2(t_2) - C_M(t_2)| \le \xi(1-\alpha).$$

By choosing $M_\xi^3 = \max\{N_\xi^3, L_\xi^3\}$, for every $k > M_\xi^3$, part (1) of Claim 3 is valid.

Since at $t_2 = 0$, for every $b_2 \in [5, 6-\alpha]$ we have

$$\frac{d}{db_2}\mathbb{E}_{\beta_1^k}(\Pi_2 \mid 0, b_2) = \frac{1-2\varepsilon_k-\varepsilon_k^2}{1-\alpha} \cdot \left(-\frac{b_2-5}{1-\alpha}\right) + \frac{\varepsilon_k^2}{1+\alpha}\,(5-\frac{\alpha}{2}-b_2) < 0,$$

the maximum at $t_2 = 0$ is obtained at $b_2 = 5$. Moreover, by part (1) of Claim 3 we have that if $k > \max\{N_\xi^3, L_\xi^3\}$, then for $t_2 = \xi$ the maximum is obtained within $b_2 \in [5, 5+2(1-\alpha)\xi]$. One can verify that in (12), the first derivative of unique solution $b_2(t_2)$ is strictly positive, then we can conclude that for every $t_2 \in [0, \xi)$ the maximum is obtained within $b_2 \in [5, 5+2(1-\alpha)\xi]$, for large k.

With a similar argument, since at $t_2 = 1$ for every $b_2 \in [5, 6-\alpha]$ we have

$$\frac{d}{db_2}\mathbb{E}_{\beta_1^k}(\Pi_2 \mid 1, b_2) = \frac{1-2\varepsilon_k-\varepsilon_k^2}{1-\alpha} \cdot \left(1-\frac{b_2-5}{1-\alpha}\right) + \frac{\varepsilon_k^2}{1+\alpha}\,(6-\frac{\alpha}{2}-b_2) > 0,$$

the maximum at $t_2 = 1$ is obtained in $b_2 = 6-\alpha$. Moreover, if $k > \max\{N_\xi^3, L_\xi^3\}$, then for $t_2 = 1-\xi$ the maximum is obtained within $b_2 \in [6-\alpha-2(1-\alpha)\xi, 6-\alpha]$. Similarly, since ${b_2}'(t_2)$ is strictly positive, we conclude that for every $t_2 \in (1-\xi, 1]$ the maximum is obtained within $b_2 \in [6-\alpha-2(1-\alpha)\xi, 6-\alpha]$ for large k. This complete the proof of Claim 3, for every $k > M_\xi^3$.

7.6 Proofs of Propositions 22 and 23

It is easy to check that strategy profiles σ and η are BNEs.

We now prove that the BNE $\sigma = (\sigma_1, \sigma_2)$, in which $\sigma_1(t_1, \cdot) = \delta_6(\cdot)$ and $\sigma_2(t_2, \cdot) = \delta_{5-\alpha}(\cdot)$, is not pointwise-perfect, consequently is not uniform-perfect. Let $(\sigma_1^k)_{k=1}^{\infty}$ be a sequence of completely mixed strategies for bidder 1 such that $\rho^w(\sigma_1^k, \sigma_1) \longrightarrow 0$ as $k \to \infty$. We show that bidder's 2 best response against σ_1^k does not converge to σ_2 when $k \to \infty$.
Suppose that bidder 1 plays σ_1^k, for some $k \in \mathbb{N}$. We prove that for $t_2 = 1$, the best response of bidder 2 is far from $\sigma_2(t_2, \cdot) = \delta_{5-\alpha}(\cdot)$. Note that bidder 2 gets always zero by choosing σ_2, because he never wins. So, it is enough to show that his expected profit is strictly positive by bidding $6-\alpha$, when $t_2 = 1$. We compute bidder 2's expected profit given type $t_2 = 1$ and bid $b_2 = 6 - \alpha$.

$$\mathbb{E}_{\sigma_1^k}(\Pi_2 \mid 1, 6-\alpha) \quad = \quad \int_0^1 \int_{5-\alpha}^{6-\alpha} (5 + 1 - \alpha t_1 - b_1)\, \sigma_1^k(t_1, db_1)\, dt_1.$$

Since σ_1^k is a completely mixed strategy we have $\sigma_1^k\,(t_1, (5-\alpha, 6-\alpha)) > 0$, therefore by using the fact that $b_1 \in (5-\alpha, 6-\alpha)$ we have

$$\begin{aligned} \mathbb{E}_{\sigma_1^k}(\Pi_2 \mid 1, 6-\alpha) \quad &> \quad \int_0^1 \int_{5-\alpha}^{6-\alpha} (5 + 1 - \alpha t_1 - 6 + \alpha)\, \sigma_1^k(t_1, db_1)\, dt_1 \\ &= \quad \alpha \int_0^1 \int_{5-\alpha}^{6-\alpha} (1 - t_1)\, \sigma_1^k(t_1, db_1)\, dt_1 \geq 0. \end{aligned}$$

This implies that $\mathbb{E}_{\sigma_1^k}(\Pi_2 \mid 1, 6-\alpha) > 0$, hence σ_2 is not pointwise-perfect.

With a similar argument, one can prove that the discontinuous BNE η is not pointwise-perfect, and hence not uniform-perfect either.

8 Appendix III: A measurability result for the weak distance

A metric space (X, d) is called σ-compact, if X is a countable union of compact subsets of X.
A strategy β_i for some player i is called simple if there exists a finite set $B \subseteq A_i$ such that $\beta_i(t_i, B) = 1$ for every $t_i \in T_i$.

Lemma 28 *Assume that the action space A_i is σ-compact for a player $i \in N$. Then, for every strategy β_i for player i, there exists a sequence $(\tau_i^k)_{k=1}^{\infty}$ of simple strategies such that*

$$\lim_{k\to\infty} \rho^w(\tau_i^k(t_i, \cdot), \beta_i(t_i, \cdot)) = 0 \tag{14}$$

for every type $t_i \in T_i$.

Proof.

Part 1: The definition of the sequence $(\tau_i^k)_{k=1}^\infty$: Since A_i is σ-compact, $A_i = \cup_{j=1}^\infty K_j$ for some compact sets $K_j \subseteq A_i$. Let $A_i^k = \cup_{j=1}^k K_j$ for every $k \in \mathbb{N}$. Note that A_i^k is compact for every $k \in \mathbb{N}$, the sequence $(A_i^k)_{k=1}^\infty$ is increasing and $A_i = \cup_{k=1}^\infty A_i^k$.

Take a $k \in \mathbb{N}$. For every $a \in A_i^k$, let $U^k(a) = \{a' \in A_i \mid d_{A_i}(a, a') < \frac{1}{k}\}$. Since A_i^k is compact, there exist actions $a_1^k, \ldots, a_{J^k}^k \in A_i^k$ such that

$$A_i^k \subseteq \bigcup_{j=1}^{J^k} U^k(a_j^k).$$

Define $B_1^k = U^k(a_1^k) \cap A_i^k$, and $B_j^k = \left[U^k(a_j^k) \cap A_i^k\right] \setminus \cup_{\ell=1}^{j-1} U^k(a_\ell^k)$ for every $j = 2, \ldots, J^k$. Without loss of generality we can assume, for every j, that B_j^k is non-empty; otherwise we could leave out a_j^k from the list $a_1^k, \ldots, a_{J^k}^k$. Thus, $\{B_1^k, \ldots, B_{J^k}^k\}$ forms a partition of A_i^k. Moreover, by construction, $B_j^k \in \mathcal{A}_i$ for every $j = 1, \ldots, J^k$. Take an arbitrary $b_j^k \in B_j^k$ for every $j = 1, \ldots, J^k$.

Now define the simple strategy τ_i^k for player i by

$$\tau_i^k(t_i, \cdot) = \beta_i(t_i, B_1^k) \cdot \delta_{b_1^k}(\cdot) + \ldots + \beta_i(t_i, B_{J^k}^k) \cdot \delta_{b_{J^k}^k}(\cdot) + (1 - \beta_i(t_i, A_i^k)) \cdot \delta_{b_1^k}(\cdot)$$

for every $t_i \in T_i$, where as usual, δ stands for the Dirac measure. Note that τ_i^k satisfies condition (2) of the definition of behavior strategies. The last term in the definition of $\tau_i^k(t_i, \cdot)$ will not play an important role, but it is needed so that $\tau_i^k(t_i, \cdot)$ is a probability measure.

Part 2: The proof that the sequence $(\tau_i^k)_{k=1}^\infty$ satisfies equality (14): Fix a type $t_i \in T_i$. Consider a bounded Lipschitz function $f : A_i \to \mathbb{R}$. Then, there exist $C_1, C_2 \geq 0$ such that $|f(a)| \leq C_1$ and $|f(a) - f(a')| \leq C_2 \cdot d_{A_i}(a, a')$ for every $a, a' \in A_i$. Let

$$D^k = \left| \int_{A_i} f(a)\ \tau_i^k(t_i, da) - \int_{A_i} f(a)\ \beta_i(t_i, da) \right|.$$

Since

$$\int_{A_i} f(a)\ \tau_i^k(t_i, da) = \sum_{j=1}^{J^k} f(b_j^k) \cdot \beta_i(t_i, B_j^k) + f(b_1^k) \cdot (1 - \beta_i(t_i, A_i^k))$$

and

$$\int_{A_i} f(a)\ \beta_i(t_i, da) = \sum_{j=1}^{J^k} \int_{B_j^k} f(a)\ \beta_i(t_i, da) + \int_{A_i \setminus A_i^k} f(a)\ \beta_i(t_i, da),$$

we have

$$\begin{aligned}
D^k &\leq \sum_{j=1}^{J^k} \left| f(b_j^k) \cdot \beta_i(t_i, B_j^k) - \int_{B_j^k} f(a)\ \beta_i(t_i, da) \right| + 2C_1 \cdot (1 - \beta_i(t_i, A_i^k)) \\
&\leq \sum_{j=1}^{J^k} \left| \int_{B_j^k} \left[f(b_j^k) - f(a)\right]\ \beta_i(t_i, da) \right| + 2C_1 \cdot (1 - \beta_i(t_i, A_i^k)) \\
&\leq \sum_{j=1}^{J^k} \sup_{a \in B_j^k} \left| f(b_j^k) - f(a) \right| \cdot \beta_i(t_i, B_j^k) + 2C_1 \cdot (1 - \beta_i(t_i, A_i^k)).
\end{aligned}$$

Since the diameter of B_j^k, for each $j = 1, \ldots, J^k$, is at most $\frac{2}{k}$, it follows that

$$\begin{aligned} D^k &\leq \sum_{j=1}^{J^k} \frac{2}{k} \cdot C_2 \cdot \beta_i(t_i, B_j^k) + 2C_1 \cdot (1 - \beta_i(t_i, A_i^k)) \\ &= \frac{2}{k} \cdot C_2 \cdot \beta_i(t_i, A_i^k) + 2C_1 \cdot (1 - \beta_i(t_i, A_i^k)). \end{aligned}$$

By the continuity of the measure, $\beta_i(t_i, A_i^k)$ converges to $\beta_i(t_i, A_i) = 1$ as $k \to \infty$, and hence D^k converges to zero as $k \to \infty$. As f was an arbitrary bounded Lipschitz function from A_i to $\mathbb{R}$, the proof of (14) is complete. □

Lemma 29 *Assume that the action space A_i is σ-compact for a player $i \in N$, and let β_i^1 and β_i^2 be two strategies for player i. Then, the function $f : T_i \to \mathbb{R}$ defined by $f(t_i) = \rho^w(\beta_i^1(t_i, \cdot), \beta_i^2(t_i, \cdot))$ is measurable.*

Proof. It suffices to prove that the set $\{t_i \in T_i \,|\, f(t_i) > r\}$ is measurable for every $r \in \mathbb{R}$. So, fix an arbitrary $r \in \mathbb{R}$.

Part 1: when β_i^1 and β_i^2 are simple strategies.

First assume that β_i^1 and β_i^2 are simple strategies. Then, there are finite sets $B^1, B^2 \subseteq A_i$ such that $\beta_i^1(t_i, B^1) = 1$ and $\beta_i^2(t_i, B^2) = 1$ for every $t_i \in T_i$. Let $B = B^1 \cup B^2$. We have

$$f(t_i) = \inf \left\{ \varepsilon > 0 \mid \beta_i^1(t_i, C) \leq \beta_i^2(t_i, C^\varepsilon) + \varepsilon \;\text{ and }\; \beta_i^2(t_i, C) \leq \beta_i^1(t_i, C^\varepsilon) + \varepsilon \;\; \forall C \in \mathcal{A}_i \right\}.$$

Notice that if $\beta_i^1(t_i, C) \leq \beta_i^2(t_i, C^\varepsilon) + \varepsilon$ for some $C \subseteq B$, then for any $\widetilde{C} \in \mathcal{A}_i$ satisfying $\widetilde{C} \cap B = C$, we obtain

$$\beta_i^1(t_i, \widetilde{C}) = \beta_i^1(t_i, C) \leq \beta_i^2(t_i, C^\varepsilon) + \varepsilon \leq \beta_i^2(t_i, \widetilde{C}^\varepsilon) + \varepsilon.$$

Similarly, if $\beta_i^2(t_i, C) \leq \beta_i^1(t_i, C^\varepsilon) + \varepsilon$ then

$$\beta_i^2(t_i, \widetilde{C}) \leq \beta_i^1(t_i, \widetilde{C}^\varepsilon) + \varepsilon.$$

Hence,

$$f(t_i) = \inf \left\{ \varepsilon > 0 \mid \beta_i^1(t_i, C) \leq \beta_i^2(t_i, C^\varepsilon) + \varepsilon \;\text{ and }\; \beta_i^2(t_i, C) \leq \beta_i^1(t_i, C^\varepsilon) + \varepsilon \;\; \forall C \subseteq B \right\}.$$

Notice that $f(t_i) > r$ holds if and only if there exists an $m \in \mathbb{N}$ and a set $C \subseteq B$ such that for $\varepsilon = r + \frac{1}{m}$ we have either $\beta_i^1(t_i, C) > \beta_i^2(t_i, C^\varepsilon) + \varepsilon$ or $\beta_i^2(t_i, C) > \beta_i^1(t_i, C^\varepsilon) + \varepsilon$. Indeed, the "only if"-part is immediate, whereas the "if"-part follows from the fact that $\beta_i^2(t_i, C^\varepsilon) + \varepsilon$ and $\beta_i^1(t_i, C^\varepsilon) + \varepsilon$ are increasing in ε.

For every $m \in \mathbb{N}$ and $C \subseteq B$, define

$$U_{m,C}^1 = \left\{ t_i \in T_i \,|\, \beta_i^1(t_i, C) > \beta_i^2(t_i, C^{r + \frac{1}{m}}) + r + \frac{1}{m} \right\}$$

$$U^2_{m,C} = \left\{ t_i \in T_i \,|\, \beta^2_i(t_i, C) > \beta^1_i(t_i, C^{r+\frac{1}{m}}) + r + \frac{1}{m} \right\}.$$

Since by condition (2) of the definition of behavior strategies, the section functions $\beta^1_i(\cdot, C) : T_i \to \mathbb{R}$ and $\beta^2_i(\cdot, C^{r+\frac{1}{m}}) : T_i \to \mathbb{R}$ are measurable, the sets $U^1_{m,C}$ are measurable. For a similar reason, the sets $U^2_{m,C}$ are measurable too. Now we have

$$\{t_i \in T_i \,|\, f(t_i) > r\} \;=\; \bigcup_{m=1}^{\infty} \bigcup_{C \subseteq B} \left[U^1_{m,C} \bigcup U^2_{m,C} \right].$$

Because the right hand side is a countable union of measurable sets, the set $\{t_i \in T_i \,|\, f(t_i) > r\}$ is measurable as well. So, the claim of the lemma holds for simple strategies.

Part 2: when β^1_i and β^2_i are arbitrary strategies.

Now we prove the lemma for arbitrary strategies β^1_i and β^2_i. By Lemma 28, there are two sequences of simple strategies $(\tau^{k,1}_i)^{\infty}_{k=1}$ and $(\tau^{k,2}_i)^{\infty}_{k=1}$ such that

$$\lim_{k\to\infty} \rho^w(\tau^{k,1}_i(t_i, \cdot), \beta^1_i(t_i, \cdot)) = 0, \quad \text{and} \quad \lim_{k\to\infty} \rho^w(\tau^{k,2}_i(t_i, \cdot), \beta^2_i(t_i, \cdot)) = 0$$

for every type $t_i \in T_i$. By the definition of f, we have

$$\{t_i \in T_i \,|\, f(t_i) > r\} = \{t_i \in T_i \,|\, \lim_{k\to\infty} \rho^w(\tau^{k,1}_i(t_i, \cdot), \tau^{k,2}_i(t_i, \cdot)) > r\}.$$

Notice that, for some $t_i \in T_i$, the inequality

$$\lim_{k\to\infty} \rho^w(\tau^{k,1}_i(t_i, \cdot), \tau^{k,2}_i(t_i, \cdot)) > r$$

holds if and only if there exists an $m \in \mathbb{N}$ such that $\rho^w(\tau^{k,1}_i(t_i, \cdot), \tau^{k,2}_i(t_i, \cdot)) > r + \frac{1}{m}$ holds for sufficiently large k. Therefore,

$$\{t_i \in T_i \,|\, f(t_i) > r\} = \bigcup_{m=1}^{\infty} \bigcup_{\ell=1}^{\infty} \bigcap_{k=\ell}^{\infty} \left\{ t_i \in T_i \,|\, \rho^w(\tau^{k,1}_i(t_i, \cdot), \tau^{k,2}_i(t_i, \cdot)) > r + \frac{1}{m} \right\}.$$

For each $k \in \mathbb{N}$, because the strategies $\tau^{k,1}_i$ and $\tau^{k,2}_i$ are simple, part 1 of the proof implies that set

$$\left\{ t_i \in T_i \,|\, \rho^w(\tau^{k,1}_i(t_i, \cdot), \tau^{k,2}_i(t_i, \cdot)) > r + \frac{1}{m} \right\}$$

is measurable. Therefore, the set $\{t_i \in T_i \,|\, f(t_i) > r\}$ is measurable as well. □

References

[1] Aliprantis CD, Border C (1999): Infinite Dimensional Analysis. Springer-Verlag.

[2] Aumann RJ (1964): Mixed and behavior strategies in infinite extensive games. *Advances in Game Theory*, Annals of Mathematical Studies 52, Princeton University Press, Princeton, N.J. 627–650.

[3] Bajoori E, Flesch J and Vermeulen D (2013): Perfect equilibrium in games with compact action spaces. *Games and Economic Behavior* 82, 490–502.

[4] Bajoori E (2011): Distributional perfect equilibrium in Bayesian games. Working paper.

[5] Billingsley P (1968): Convergence of probability measures. Wiley Series in Probability and Mathematical Statistics, John Wiley & Sons.

[6] Bonanno G (2013): AGM-consistency and perfect Bayesian equilibrium. Part I: definition and properties. *International Journal of Game Theory* 42, 567-592.

[7] Dudley RM (2002): Real analysis and probability. Cambridge University Press.

[8] Fudenberg D and Tirole J (1991): Perfect Bayesian Equilibrium and sequential equilibrium. *Journal of Economic Theory* 53, 236–260.

[9] Govindan S (1995): Stability and the chain store paradox. *Journal of Economic Theory* 66, 536–547.

[10] Jackson MO, Simon LK, Swinkels JM and Zame WR (2002): Communication and equilibrium in discontinuous games of incomplete information. *Econometrica* 70, 1711–1740.

[11] Jackson MO, Simon LK, Swinkels JM and Zame WR (2004): Corrigendum to "Communication and equilibrium in discontinuous games of incomplete information". *Econometrica* 72, 1927–1929.

[12] Kreps DM and Wilson R (1982): Sequential Equilibria. *Econometrica* 50, 863–894.

[13] Laffont JJ and Martimort D (2001): The theory of incentives: the principal-agent model. Princeton University Press.

[14] Milgrom PR and Weber RJ (1985): Distributional strategies for games with incomplete information. *Mathematics of Operations Research* 10, 619–632.

[15] Myerson RB and Reny PJ (2011): Sequential equilibria of multi-stage games with infinite sets of types and actions. *Working paper.*

[16] Nash JF (1951): Non-Cooperative Games. *The Annals of Mathematics 54* 2, 286–295.

[17] Parthasarathy KR (1967): *Probability Measures on Metric Spaces.* Academic Press, New York - London, 1967.

[18] Prokhorov YV (1956): Convergence of random processes and limit theorems in probability theory, *Theory of Probability and Its Applications* 1, 157–214.

[19] Reny, PJ (1999): On the existence of pure and mixed strategy Nash equilibria in discontinuous games. *Econometrica*, 67, 1029–1056.

[20] Reny PJ (2011): On the existence of monotone pure strategy equilibrium in Bayesian games. *Econometrica* 79, 499–553.

[21] Selten R (1975): Reexamination of the perfectness concept for equilibrium points in extensive games. *International Journal of Game Theory* 4, 25–55.

[22] Selten R (1978): The chain store paradox. *Theory and Decision* 9, 127–159.

[23] Simon LK and Stinchcombe MM (1995): Equilibrium refinement for infinite normal-form games. *Econometrica* 63, 1421–1443.