



Ambiguity and price competition

R. R. Routledge¹ · R. A. Edwards²

Published online: 4 November 2019
© The Author(s) 2019

Abstract

There are few models of price competition in a homogeneous-good market which permit general asymmetries of information amongst the sellers. This work studies a price game with discontinuous payoffs in which both costs and market demand are ex ante uncertain. The sellers evaluate uncertain profits with maximin expected utilities exhibiting ambiguity aversion. The buyers in the market are permitted to split between sellers tying at the minimum price in arbitrary ways which may be deterministic or random. The role of the primitives in determining equilibrium prices in the market is analyzed in detail.

Keywords Ambiguity · Maximin utilities · Asymmetric information · Game theory

1 Introduction

In a classic Bertrand price competition, goods are homogeneous and sellers select their price with complete information regarding the structure of market demand, their production costs, and the production technologies of their competitors. The lowest priced seller secures all forthcoming demand or an equal share if multiple sellers tie. This environment is far from realistic. Demand and costs are typically uncertain at the ex ante stage and sellers will possess different information about the state of the world. Sharing rules may also be different from equal sharing at price ties. Can we guarantee equilibrium existence in such a market with discontinuous payoffs, arbitrary sharing

University of Nottingham, United Kingdom. Email: Robert.Edwards@Nottingham.ac.uk The authors would like to thank the participants of the 2018 OLIGO Workshop (University of Piraeus), the attendees of the Liverpool economic research seminar, and their colleagues, especially Christian Bach, Dominique Demouglin, Olga Gorelkina, Yiquan Gu, Erin Hengel, Sanna Laksa, Eunyoung Moon, Harvey Upton and Yunchou Wu for their encouragement and comments during the writing of this paper.

✉ R. R. Routledge
R.R.Routledge@liverpool.ac.uk

¹ UoL Management School, UoL, Chatham Street, Liverpool L69 7ZH, UK

² University of Nottingham, Nottingham NG7 2RD, UK

at price ties, and incomplete, asymmetric information across sellers? Surprisingly, this fundamental issue remains an open question.

One might wonder why it is worth studying price competition in the context of ambiguity? The first point to emphasise is that there are few papers which study price competition in the context of *any* form of uncertainty.¹ The reason for this is probably technical: Bertrand competition in a homogeneous-good market is a game with discontinuous payoffs and there are few existence results for Nash equilibria in such contexts. The second point is that, in the context of a one-shot game, assuming standard uncertainty (that a seller can assign probabilities to all relevant states) is a strong and often unrealistic assumption. Therefore, in this context, permitting ambiguity is more plausible (and our model does contain the special case in which sellers can assign probabilities to all states of the world). Our work is particularly interesting because, by introducing ambiguity into the standard Bertrand price game, we are able to provide straightforward pure strategy equilibrium existence results, despite the game having discontinuous payoffs.

In the model, we study there are three primitives: the market demand, the sellers' cost functions and the sharing rule at price ties. Uncertainty is permitted to affect both the market demand and the cost functions (and could be extended to the sharing rule, but we omit this to keep the notation as simple as possible). In this context, it is plausible that a seller knows their own costs (which we assume) but is uncertain about other sellers' costs and the market demand. Indeed, there is a large literature studying market competition under uncertainty when sellers know the relevant probabilities (Vives 1999, Ch.8). In a one-shot game, though, a seller has no previous history upon which to base the probabilities of different states of the world. Consequently, a seller is unlikely to be able to assign probabilities to different market states with any certainty or confidence.² In this context, ambiguity aversion is a more appropriate decision rule to use.

We consider sellers that face ambiguity regarding exogenously specified state-contingent demand and cost functions and information is asymmetric across sellers. An information partition containing multiple states captures states of the world between which the seller cannot distinguish. In this incomplete information game, ambiguity averse sellers simultaneously and independently choose prices to maximise their maximin expected utilities (MEU) as in Gilboa and Schmeidler (1989). Maximin preferences have gained significant traction in recent years as an alternative to the increasingly refuted Bayesian paradigm.³

Our objective is to provide new conditions under which pure strategy price equilibria exist, whilst permitting discontinuities in the demand function, asymmetric information across sellers and arbitrary price-tie sharing rules. The prices chosen by sellers must also be measurable with respect to their private information, which reflects the way in which sellers' information partitions constrains their choices of prices. In this

¹ Notable exceptions are Spulber (1995) and Wambach (1999). Although in both these papers, the uncertainty is ex ante symmetric. Our model permits general asymmetries of information.

² See, amongst others, the discussion in de Castro and Yannelis (2011) and He and Yannelis (2015a).

³ See for instance: Allingham (2002, Ch.3), Correia-Vioglio et al. (2013), Correia-da-Silva and Hervés-Beloso (2009, 2012); de Castro and Yannelis 2011; He and Yannelis (2015a, 2017a,b); Pulford and Colman (2007).

context, we are able to provide direct and constructive existence results for a pure strategy price equilibrium in the incomplete information game.⁴

We show that two relatively mild conditions are sufficient to guarantee a continuum of pure strategy price equilibria, which lend to a wide range of applications. Firstly, there exist at least two sellers of each information type. Secondly, cost functions are not too different. This strictly generalises previous results that hinged critically on equal or winner-takes-all sharing rules.⁵ The first condition overcomes the well-known ‘open-set’ problem, where the best response of the seller may not be well defined.⁶ Requiring that no seller has a unique information partition, combined with similar costs, ensures that we are able to construct equilibria in which sellers with symmetric information choose the same price. Therefore, a seller cannot profitably increase their price as they will be tied with at least one rival. This condition is also reasonably unrestrictive and will be satisfied if we consider the case where the set of sellers is replicated any number of times (as in Debreu and Scarf 1963).

Of particular interest in our model is the relationship between the ex post equilibrium sets and pure strategy Nash equilibria in the ex ante price game. It is well known that in complete information price games with sellers who have strictly convex costs there often exists a continuum of pure strategy Nash equilibria (Dastidar 1995; Vives 1999, p.122). In this work, we have been able to find a connection between these ex post equilibrium sets and the Nash equilibria of the ex ante game: with maximin expected utilities, one can construct ex ante Nash equilibria by carefully selecting equilibria from the ex post equilibrium sets. A similar point has been noted by He and Yannelis (2017a) for discontinuous games with asymmetric information satisfying the Reny (1999) conditions. However, unlike them, we still require that the prices posted in the marketplace be measurable with respect to the private information of the sellers.

Many departures from the classical Bertrand setup have been extensively explored in the literature. Recent developments include generalising existence results to permit discontinuities in demand and cost functions (Baye and Morgan 2002), a wider range of production technologies (Saporiti and Coloma 2010; Baye and Kovenock 2008), arbitrary sharing rules at price ties (Bagh 2010; Hoernig 2007), and incomplete and (possibly) differential information amongst sellers. This paper makes contributions that span all these active directions of research.

Section 2 presents the Bertrand game we study. Section 3 outlines existence conditions for pure strategy equilibrium. Section 4 provides illustrative examples which should give the reader a clear understanding of the results, and their wide applicability. Section 5 discusses related literature on price competition and ambiguity. Finally, we present some concluding remarks about directions for future research in this area.

⁴ Being able to give direct and constructive existence results regarding pure equilibria is unusual when one departs from the classical complete information game of Bertrand competition. Recent papers have shown that finding an equilibrium point in more complex setting can be difficult, and often equilibria only exist in mixed strategies. See, for a recent example of a behavioural Bertrand price game, Edwards (2019).

⁵ Winner-takes-all sharing specifies that one of the tied sellers at the minimum price is selected randomly to serve the entire market.

⁶ See Vives (1999, p.123), and Blume (2003), for a discussion of this problem.

2 The Bertrand Game

The model consists of a finite set of sellers $N = \{1, \dots, n\}$, $n \geq 2$, who are producing a single perfectly homogeneous good. The uncertainty will be modelled by a finite set $\Omega = \{\omega_1, \dots, \omega_m\}$ which is the set of possible **states of the world**. There is a probability distribution, μ , over the set Ω which describes the probability of each state occurring. It shall be assumed that $\mu(\omega) > 0$ for every $\omega \in \Omega$ so no state of the world is redundant. Each seller has a state-contingent **cost function** given by $C_i : \mathfrak{R}_+ \times \Omega \rightarrow \mathfrak{R}_+$. The following conditions are imposed upon the cost functions.

Assumption 2.1 *For every $i \in N$, and every $\omega \in \Omega$, $C_i(\cdot, \omega)$ is strictly convex, strictly increasing and satisfies $C_i(0, \omega) = 0$.*

There is a state-contingent **market demand function** for the homogeneous good given by $D : \mathfrak{R}_+ \times \Omega \rightarrow \mathfrak{R}_+$. The following conditions are imposed upon the demand function.

Assumption 2.2 *For every $\omega \in \Omega$, there exist positive finite real numbers $\bar{x}(\omega)$ and $\bar{y}(\omega)$ such that $D(x, \omega) = 0$ for all $x \geq \bar{x}(\omega)$, and $D(0, \omega) = \bar{y}(\omega)$. The function $D(\cdot, \omega)$ is strictly decreasing on $(0, \bar{x}(\omega))$.*

Note that no continuity, or any smoothness conditions, is imposed upon the demand function for the good. As is standard in asymmetric information models, the **private information** of seller i is modelled by an algebra Σ_i on the set Ω . Let P_i be the partition of the set Ω associated with Σ_i (the minimal elements in Σ_i for which no strict subset is also contained in Σ_i). Let $\Sigma = \bigcap_{i \in N} \Sigma_i$ and $P(\Sigma)$ be the partition of Ω associated with the coarsest algebra, Σ , generated by the intersection of the sellers' algebras. Whenever two states of the world are in the same element in the partition P_i , it is taken that seller i is unable to distinguish between those two states. A function $f : \Omega \rightarrow \mathfrak{R}_+$ will be called P_i -measurable if, whenever $\omega_p \in E$ and $\omega_q \in E$ for some event $E \in P_i$, then $f(\omega_p) = f(\omega_q)$. Facing these information restrictions, the strategy set of seller i in the Bertrand game is

$$L_i = \{f : \Omega \rightarrow \mathfrak{R}_+ : f \text{ is } P_i\text{-measurable}\}.$$

Let $L = \times_{i \in N} L_i$ be the joint strategy set. It shall be taken as given that a seller knows their cost function before posting their price in the marketplace. To avoid a seller being able to infer more information about the state of the world than that given by their partition, the following condition is imposed upon the cost functions.

Assumption 2.3 *For every $i \in N$, the function $C_i(\cdot, \omega)$ is P_i -measurable.*

In the classical Bertrand game sellers post prices in the market with a commitment to meet all the market demand forthcoming.⁷ As a result, consumers will buy from the sellers posting the minimum price in the market. At this point, it becomes necessary

⁷ An excellent and succinct summary of the differences between Bertrand competition and Bertrand-Edgeworth competition (which does permit rationing of the market demand) is Vives (1999, Ch.5).

to specify what happens when different sellers tie at the minimum price. How is the market demand split between them? The most commonly used rule in the literature is the “equal sharing” rule which assumes that sellers tying at the minimum price receive an equal share of the market demand. However, this is clearly restrictive and recent research has indicated that the existence and nature of equilibria in Bertrand price games are highly sensitive to the choice of the sharing rule at minimum price ties.⁸ Given this problem, we shall consider general classes of random and deterministic sharing rules. Let the set S_i be given by

$$S_i = \{S \subseteq N : i \in S \text{ and } |S| \geq 2\}.$$

The set S_i is all the possible coalitions of sellers which i could tie with at the minimum price.

2.1 The sharing rules

A **deterministic sharing rule** is a set of functions $\{g_{i,S}\}_{i \in N, S \in S_i}$ satisfying $g_{i,S} : \mathfrak{R}_+ \rightarrow (0, 1)$ with $S \in S_i$ and $\sum_{i \in S} g_{i,S}(x) = 1$ for every $x \in \mathfrak{R}_+$. In words, a sharing rule for seller i with coalition $S \in S_i$ is a function which maps the price tied at into the share of the market demand received. Given a deterministic sharing rule, the shared profit when seller $i \in N$ ties at the minimum price with $S \in S_i$ other sellers is

$$\pi_{i,S}(x, \omega) = x g_{i,S}(x) D(x, \omega) - C_i(g_{i,S}(x) D(x, \omega), \omega).$$

A **random sharing rule** is again a set of functions $\{h_{i,S}\}_{i \in N, S \in S_i}$ satisfying $h_{i,S} : \mathfrak{R}_+ \rightarrow (0, 1)$ and $\sum_{i \in S} h_{i,S}(x) = 1$ for every $x \in \mathfrak{R}_+$. However, the number assigned to each seller does not represent the share of the market demand received, but the probability of seller i being chosen to meet all the market demand. The complete set of random sharing rules is given by $\{h_{i,S}\}_{i \in N, S \in S_i}$. Given a random sharing rule, the shared profit when seller $i \in N$ ties at the minimum price with $S \in S_i$ other sellers is

$$\pi_{i,S}(x, \omega) = h_{i,S}(x)(x D(x, \omega) - C_i(D(x, \omega), \omega)).$$

Remark 2.1 This way of specifying the sharing rules at price ties contains the deterministic equal sharing rule, and the random winner-takes-all sharing rule, as special cases. If one sets $g_{i,S}(x) = 1/|S|$ for every $x \in \mathfrak{R}_+$, $S \in S_i$ and every $i \in N$, the standard equal sharing rule is restored. If one sets $h_{i,S}(x) = 1/|S|$ for every $x \in \mathfrak{R}_+$, $S \in S_i$ and every $i \in N$, the winner-takes all sharing rule is captured. There are clearly uncountably many other possible sharing rules which are contained in this framework. It may be noted we are imposing that, unlike the cost and demand functions, the sharing rule is independent of the state of the world. All the results which follow could be readily extended to permit the sharing rule to also be state contingent, without many changes, although this would make the notation a little more complicated.

⁸ See for example Vives (1999, pp.117–123), Hoernig (2007) and Bagh (2010).

2.2 The ex post payoffs

Given the specification of the sharing rules at minimum price ties it is now possible to describe the profits in the ex post games. The **monopoly profit** of seller i in state $\omega \in \Omega$ when posting price x is

$$\pi_i(x, \omega) = xD(x, \omega) - C_i(D(x, \omega), \omega).$$

The **deterministic shared profit** when seller $i \in N$ ties at the minimum price with $S \in S_i$ other sellers in state $\omega \in \Omega$ given deterministic sharing rules $\{g_{i,S}\}_{i \in N, S \in S_i}$ is

$$\pi_{i,S}(x, \omega) = xg_{i,S}(x)D(x, \omega) - C_i(g_{i,S}(x)D(x, \omega), \omega).$$

The **random shared profit** when seller $i \in N$ ties at the minimum price with $S \in S_i$ other sellers in state $\omega \in \Omega$ given random sharing rules $\{h_{i,S}\}_{i \in N, S \in S_i}$ is

$$\pi_{i,S}(x, \omega) = h_{i,S}(x)(xD(x, \omega) - C_i(D(x, \omega), \omega)) = h_{i,S}(x)\pi_i(x, \omega).$$

The following mild conditions are imposed upon the monopoly and tied profit functions.

Assumption 2.4 For every $i \in N$, $\omega \in \Omega$ and $S \in S_i$, the functions $\pi_i(\cdot, \omega)$ and $\pi_{i,S}(\cdot, \omega)$ satisfy the following:

- (i) They are left lower semicontinuous.
- (ii) There exist prices at which they achieve a strictly positive value.

Remark 2.2 The conditions in Assumption 2.4 are widely known in the literature and are discussed in detail in, amongst others, Baye and Morgan (2002) and Bagh (2010). Left lower semicontinuity permits that a function can jump downwards as it is approached from the left, but not from the right. From Assumption 2.1, the cost functions of all the sellers are strictly convex, so any discontinuities generated in the profit functions must come about from discontinuities in the demand function. Taken together, the two conditions in Assumption 2.4 guarantee that there exist solutions to the equations $\pi_i(x, \omega) = 0$ and $\pi_{i,S}(x, \omega) = 0$ in the region $0 \leq x \leq \bar{x}(\omega)$.

Given the specifications of all the different profit functions, the full payoffs in the discontinuous game can be summarized. Fix a vector of strategies $f = (f_1, \dots, f_n) \in L$, the ex post payoff which seller i obtains in state $\omega \in \Omega$ is

$$u_i(f, \omega) = \begin{cases} \pi_i(f_i(\omega), \omega), & \text{if } f_i(\omega) < f_j(\omega) \text{ for all } j \neq i; \\ \pi_{i,S}(f_i(\omega), \omega), & \text{if } i \text{ ties with } S \in S_i \text{ sellers at min. price;} \\ 0 & \text{if there is a } j \text{ s.t. } f_j(\omega) < f_i(\omega). \end{cases}$$

2.3 The ex ante payoffs

Let us now turn to consider how, at the ex ante stage, before each seller has received the information regarding which element in P_i the state of the world is in, the sellers evaluate their expected payoff. As noted in the introduction, standard Bayesian expected payoffs are difficult to work with in this context. Given this difficulty, we consider a well-known alternative expected payoff: maximin expected utilities. If a seller knows that the state of the world is contained in $E \in P_i$, we consider the case where the seller is pessimistic and assigns all the probability associated with event E , which is $\mu(E)$, to the minimum ex post payoff in E . To this end, let X be the set of probability distributions over Ω :

$$X = \left\{ x \in \mathfrak{R}^\Omega : x(\omega) \geq 0 \text{ for every } \omega \in \Omega \text{ and } \sum_{\omega \in \Omega} x(\omega) = 1 \right\}$$

and let M_i be the set of probability distributions which **agree with seller i 's private information**:

$$M_i = \{x \in X : x(E) = \mu(E) \text{ for every } E \in P_i\}.$$

Clearly $M_i \subseteq X$. Fixing a vector of strategies $f = (f_1, \dots, f_n) \in L$ the **ex ante payoff** of seller i is:

$$U_i(f) = \min_{x \in M_i} \left[\sum_{\omega \in \Omega} x(\omega) u_i(f, \omega) \right].$$

An alternative, but equivalent expression of the ex ante payoffs, which is often easier to analyze, is

$$U_i(f) = \sum_{E \in P_i} \mu(E) \left[\min_{\omega \in E} u_i(f, \omega) \right].$$

It is worth noting that if seller i 's information partition P_i contains all singletons, then maximin expected utility coincides with standard Bayesian expected utilities. Therefore, the type of ex ante evaluation of the payoffs we are considering can be seen as a generalization of Bayesian expected utilities.

Remark 2.3 Gilboa and Schmeidler (1989) provided an axiomatic characterization of this type of maximin decision rule in the context of ambiguity aversion. It is well known that although this type of decision rule is quite simple it is capable of explaining the Ellsberg (1961) violations of standard subjective expected utility theory. In recent years, maximin expected utilities, and similar variants, have been used in a wide range of papers, including Correia-da-Silva and Hervés-Beloso (2009), de Castro and Yannelis (2011) and He and Yannelis (2015a). More generally, maximin expected utilities are a generalization of families of choice problems under ignorance (where

some relevant parameter is unknown to the decision maker). An excellent summary of philosophical choice in the context of ignorance, including maximin expected utilities, is contained in Peterson (2017).

Now that we have defined the ex ante payoffs, we can define the equilibrium concept. A set of strategies $f = (f_1, \dots, f_n) \in L$ is a **pure strategy price equilibrium** if, for every $i \in N$

$$U_i(f) \geq U_i(f'_i, f_{-i}) \text{ for every } f'_i \in L_i.$$

The primitives of a **price game with asymmetric information and sharing rules** can be summarized by $G = (N, \Omega, (C_i, P_i)_{i \in N}, \{r_{i,S}\}_{i \in N, S \in S_i}, D, \mu), r \in \{g, h\}$. In a price game with asymmetric information and sharing rules, G , two sellers $i, j \in N$ will be described as being of the same **information type** if $P_i = P_j$. In words, sellers i and j are of the same information type if their information partitions are the same. For any game G , let T^G be the unique partition of the player set N with the two properties: (i) if $T \in T^G$ and $i, j \in T$ then i and j are of the same information type (ii) if $T, T' \in T^G$ and $i \in T, j \in T'$ then i and j are *not* of the same information type.

3 Existence of pure strategy price equilibrium

Before turning to the existence results, a little bit of notation will be useful. Fix a price game $G = (N, \Omega, (C_i, P_i)_{i \in N}, \{r_{i,S}\}_{i \in N, S \in S_i}, D, \mu), r \in \{g, h\}$ and for each $i \in N, S \in S_i$ and $\omega \in \Omega$ let the real numbers $l_{i,S}(\omega)$ and $m_i(\omega)$ be defined as:

$$l_{i,S}(\omega) = \min\{x | 0 \leq x < \bar{x}(\omega) \text{ and } \pi_{i,S}(x, \omega) \geq 0\}$$

$$m_i(\omega) = \min\{x | 0 \leq x < \bar{x}(\omega) \text{ and } \pi_i(x, \omega) \geq 0\}.$$

$l_{i,S}(\omega)$ defines the lowest breakeven price for player i when tied at the lowest price in the market with coalition S . $m_i(\omega)$ defines the lowest breakeven price for player i when the firm posts the lowest price in the market and becomes the monopoly supplier.

Let $C = \{S \subseteq N : |S| \geq 2\}$. For each $S \in C$, and $\omega \in \Omega$, define the interval $I_S(\omega)$ to be

$$I_S(\omega) = \cap_{i \in S} (l_{i,S}(\omega), m_i(\omega)).$$

One final assumption is imposed upon the primitives of a price game with a deterministic sharing rule. This assumption simply requires that, with a deterministic sharing rule, regardless of which coalition a seller ties with, their demand forthcoming is decreasing in the price tied at.

Assumption 3.1 Fix a price game $G = (N, \Omega, (C_i, P_i)_{i \in N}, \{r_{i,S}\}_{i \in N, S \in S_i}, D, \mu)$ with a deterministic sharing rule $r = g$, for every $i \in N$ and $S \in S_i$, the functions $g_{i,S}(\cdot)D(\cdot, \omega)$ are decreasing on $(0, \bar{x}(\omega))$.

The following existence results use Lemmas 7.1–7.3 which, for brevity, are contained in the Appendix.

Proposition 3.1 *Fix a price game $G = (N, \Omega, (C_i, P_i)_{i \in N}, \{r_{i,S}\}_{i \in N, S \in S_i}, D, \mu)$ with a deterministic sharing rule $r = g$, if for every $E \in P(\Sigma)$,*

$$\bigcap_{\omega \in E} I_N(\omega) \neq \emptyset$$

then the game G possesses infinitely many pure strategy price equilibria.

Proof Suppose for every $E \in P(\Sigma)$, $\bigcap_{\omega \in E} I_N(\omega) \neq \emptyset$. We shall directly construct one such equilibrium. For each $E \in P(\Sigma)$, select an $x^E \in \bigcap_{\omega \in E} I_N(\omega)$. Define the strategies of the sellers to be

$$f_i^*(E) = x^E \quad \text{for every } i \in N.$$

By construction, these strategies are P_i -measurable for each seller. Hence, $f_i^* \in L_i$ for every $i \in N$.

It follows from Lemma 7.2 that $\pi_{i,N}(x^E, \omega) \geq 0$ for every $\omega \in E$. Therefore, for every $i \in N$, and $E \in P_i$,

$$\mu(E)[\min_{\omega \in E} u_i(f^*, \omega)] \geq 0$$

and

$$U_i(f^*) = \sum_{E \in P_i} \mu(E)[\min_{\omega \in E} u_i(f^*, \omega)] \geq 0.$$

Suppose, for some $E \in P_i$, seller i decided to deviate by choosing

$$f_i(E) = x < x^E.$$

Across all the states in E all the other sellers post price x^E . As a consequence, seller i would post the minimum price in the market whenever the state of the world is in E .

However, as $x < x^E < m_i(\omega)$ for every $\omega \in E$, $\pi_i(x, \omega) < 0$ for every $\omega \in E$ and

$$\mu(E)[\min_{\omega \in E} u_i(f_i, f_{-i}^*, \omega)] < 0.$$

This is not a profitable deviation. Alternatively, suppose for some $E \in P_i$ seller i decided to deviate by choosing

$$f_i(E) = x > x^E.$$

Then, for all the states in E , i never posts the minimum price in the market and $u_i(f_i, f_{-i}^*, \omega) = 0$ for every $\omega \in E$. Hence

$$\mu(E)[\min_{\omega \in E} u_i(f_i, f_{-i}^*, \omega)] = 0.$$

This is not a profitable deviation. As $\cap_{\omega \in E} I_N(\omega)$ is the intersection of finitely many open sets, $\cap_{\omega \in E} I_N(\omega)$ is also an open set, and it follows that there are infinitely many pure strategy price equilibria. \square

One might reasonably ask when the conditions of Proposition 3.1 are likely to be satisfied. If the cost functions of the sellers are not too different, the sharing rule is not too different from the standard equal sharing rule, and the market demands across the different states are quite similar, then the game is likely to satisfy the conditions. An example of this is presented in the next section. The intuition behind the result is that sometimes, the primitives of the game are such that one can select a set of prices which are $P(\Sigma)$ -measurable and which are also Nash equilibria in the games defined by the ex post payoffs. The next result demonstrates that even if one cannot make a selection of prices which are $P(\Sigma)$ -measurable and are Nash equilibria in the games formed by the ex post payoffs, there may still exist a pure strategy price equilibrium.

Proposition 3.2 *Fix a price game $G = (N, \Omega, (C_i, P_i)_{i \in N}, \{r_{i,S}\}_{i \in N, S \in S_i}, D, \mu)$ with a deterministic sharing rule $r = g$, if the following two conditions are satisfied:*

- (i) *For every $T \in T^G$, $|T| \geq 2$.*
- (ii) *For every $\omega \in \Omega$, $\cap_{S \in C} I_S(\omega) \neq \emptyset$.*

Then, the game G possesses infinitely many pure strategy price equilibria.

Proof Let G be a game satisfying the conditions in the proposition. We shall directly construct one such equilibrium of the game. Let $\{x_\omega\}_{\omega \in \Omega}$ be a set of real numbers such that, for every $\omega \in \Omega$,

$$x_\omega \in \cap_{S \in C} I_S(\omega).$$

For each $i \in N$ and $E \in P_i$ define the strategies of the sellers to be

$$f_i^*(E) = \max_{\omega \in E} x_\omega.$$

By construction, these strategies are P_i -measurable for each seller. Hence $f_i^* \in L_i$ for every $i \in N$. In each state of the world, a seller either does not post the minimum price in the market, or ties at the minimum price with some coalition.

From the definition of the strategies, for each $\omega \in E \in P_i$

$$f_i^*(\omega) = \max_{\omega \in E} x_\omega > l_{i,S}(\omega') \quad \text{for every } \omega' \in E \text{ and } S \in S_i.$$

It follows from Lemma 7.2, that $\pi_{i,S}(f_i^*(\omega), \omega) \geq 0$ for every $\omega \in E$ and $S \in S_i$. Therefore, for each $\omega \in E$

$$u_i(f^*, \omega) \geq 0$$

and

$$\min_{\omega \in E} u_i(f^*, \omega) \geq 0$$

and as $\mu(E) > 0$

$$\mu(E)[\min_{\omega \in E} u_i(f^*, \omega)] \geq 0$$

and consequently

$$U_i(f^*) = \sum_{E \in P_i} [\min_{\omega \in E} u_i(f^*, \omega)] \geq 0.$$

Suppose, for some $E \in P_i$, i decided to deviate to $f_i(E) < \max_{\omega \in E} x_\omega$. Because $f_i(E) < \max_{\omega \in E} x_\omega$, in one state in E i posts the unique minimum price in the market and obtains negative profit. Consequently

$$\min_{\omega \in E} u_i(f_i, f_{-i}^*, \omega) < 0.$$

This is not a profitable deviation. Alternatively, suppose for some $E \in P_i$, i decided to deviate to $f_i(E) > \max_{\omega \in E} x_\omega$. As $|T| \geq 2$, for every $T \in T^G$, i never posts the minimum price in the market in any state in E , and

$$\min_{\omega \in E} u_i(f_i, f_{-i}^*, \omega) = 0.$$

This is not a profitable deviation. As $\cap_{S \in C} I_S(\omega)$ is the intersection of finitely many open intervals, it is also an open set and it follows that one can construct infinitely many pure strategy price equilibria. □

The conditions in Proposition 3.2 will be satisfied when the sellers’ cost functions are not too different, and the sharing rule is not too far from the standard equal sharing rule. However, unlike the conditions in Proposition 3.1, the market demand can vary greatly across the different states, and the conditions in Proposition 3.2 can still be satisfied.

The intuition behind the two conditions in the proposition is as follows. Condition (i) says that no seller has a monopoly over their information partition. This solves the well-known “open-set” technical problem in Bertrand games where one seller may wish to price as close as possible, but slightly below, the price of another seller (in which case the best response is an empty set). Condition (ii) makes it possible for us to find a set of prices at which, in each state of the world, a seller knows that whichever other sellers he may tie with he is always guaranteed non-negative profit. As we are searching for sufficient conditions for the existence of a pure strategy equilibrium, it might be asked whether the conditions in (i) and (ii) can be significantly weakened in any way. This does not seem possible because it is easy to find games which satisfy

one of the conditions, but violate the other, and fail to possess a pure strategy price equilibrium.

An interesting question is how the price equilibrium constructed in the proof of Proposition 3.2 compares with the price equilibrium in a market with no ambiguity. Under condition (ii) of the proposition, the set $I_N(\omega)$ is non-empty for every $\omega \in \Omega$. If there was no ambiguity in the market, and the sellers could identify each state of the world, then the price equilibrium would have all the sellers setting the same price in $I_N(\omega)$. The reason this does not always happen in the equilibrium constructed in the proof of Proposition 3.2 is because the ambiguity aversion makes them reluctant to lower their prices when uncertain of the state of the world. Consequently, with ambiguity, one could observe the sellers posting different prices in the market.

The next result gives similar conditions to Proposition 3.2 that guarantee equilibrium existence with a random sharing rule.

Proposition 3.3 *Fix a price game $G = (N, \Omega, (C_i, P_i)_{i \in N}, \{r_{i,S}\}_{i \in N, S \in S_i}, D, \mu)$ with a random sharing rule $r = h$, if the following two conditions are satisfied:*

- (i) *For every $T \in T^G, |T| \geq 2$.*
- (ii) *$C_i(\cdot, \omega) = C_j(\cdot, \omega)$ for every $i \neq j$ and every $\omega \in \Omega$.*

Then, the game G possesses a pure strategy price equilibrium.

Proof As $C_i(\cdot, \omega) = C_j(\cdot, \omega)$ for every $i \neq j$ and every $\omega \in \Omega$, then $m_i(\omega) = m_j(\omega) = m(\omega)$ for every $i \neq j$. For each $i \in N$ and $E \in P_i$, define the strategies of the sellers to be

$$f_i^*(E) = \max_{\omega \in E} m(\omega).$$

By construction these strategies are P_i -measurable for each seller. Hence, $f_i^* \in L_i$ for every $i \in N$.

Given these strategies, for each $\omega \in \Omega$ seller i either does not post the minimum price in the market or i ties with some coalition at the minimum price. If i ties at the minimum price with coalition $S \in S_i$ in state ω he obtains $\pi_{i,S}(f_i^*(\omega), \omega) = h_{i,S}(f_i^*(\omega))\pi_i(f_i^*(\omega), \omega)$. From the construction of the strategies, and Lemma 7.3, it follows that for every $E \in P_i$, and every $\omega \in E, \pi_i(f_i^*(\omega), \omega) \geq 0$. Hence, for every $\omega \in \Omega$,

$$u_i(f^*, \omega) \geq 0$$

and for every $E \in P_i$

$$\min_{\omega \in E} u_i(f^*, \omega) \geq 0.$$

Furthermore, as $\mu(E) > 0$,

$$\mu(E)[\min_{\omega \in E} u_i(f^*, \omega)] \geq 0$$

and

$$U_i(f^*) = \sum_{E \in P_i} [\min_{\omega \in E} u_i(f^*, \omega)] \geq 0.$$

Suppose, for some $E \in P_i$, i decided to deviate to $f_i(E) < \max_{\omega \in E} m(\omega)$. Then

$$\min_{\omega \in E} u_i(f_i, f_{-i}^*, \omega) < 0$$

because the minimum is at least as low as i obtaining the monopoly profit in one state at a price less than the corresponding $m(\omega)$. Consequently

$$\mu(E)[\min_{\omega \in E} u_i(f_i, f_{-i}^*, \omega)] < 0.$$

This is not a profitable deviation. Suppose, for some $E \in P_i$, i decided to deviate to $f_i(E) > \max_{\omega \in E} m(\omega)$. Then, because for each $T \in T^G$, $|T| \geq 2$, i never posts the minimum price in the market in any state in E and consequently

$$\min_{\omega \in E} u_i(f_i, f_{-i}^*, \omega) = 0$$

and

$$\mu(E)[\min_{\omega \in E} u_i(f_i, f_{-i}^*, \omega)] = 0.$$

This is not a profitable deviation. □

It is worth noting that the existence results in Propositions 3.2 and 3.3 are both direct and constructive. Often with discontinuous games, we can establish that an equilibrium exists, but actually finding an equilibrium can be difficult. This illustrates the advantage of using maximin expected utilities at the ex ante stage. With maximin expected utilities, we simply have to find prices at which undercutting in one possible state of the world may make a seller worse off to prevent them wanting to undercut (because they consider the most pessimistic outcome when contemplating a deviation). With Bayesian expected utilities, this is no longer the case, and the equilibrium existence problem is significantly more complicated. This is probably one of the reasons why little work has been undertaken to examine equilibrium existence in incomplete information Bertrand games.

Proposition 3.4 *Fix a price game $G = (N, \Omega, (C_i, P_i)_{i \in N}, \{r_{i,S}\}_{i \in N, S \in S_i}, D, \mu)$ with a deterministic sharing rule $r = g$, if the following two conditions are satisfied:*

- (i) *For every $T \in T^G$, $|T| \geq 2$.*
- (ii) *$C_i(\cdot, \omega) = C_j(\cdot, \omega)$ for every $i \neq j$ and every $\omega \in \Omega$.*

Then, the game G possesses infinitely many pure strategy price equilibria.

Proof As $C_i(\cdot, \omega) = C_j(\cdot, \omega)$ for every $i \neq j$ and every $\omega \in \Omega$, it follows that $m_i(\omega) = m_j(\omega) = m(\omega)$ for every $i \neq j$. From Lemma 7.1, $l_{i,S}(\omega) < m(\omega)$ for every $S \in \mathcal{S}_i$ and $i \in N$. Therefore

$$\bigcap_{S \in \mathcal{C}} I_S(\omega) \neq \emptyset \quad \text{for every } \omega \in \Omega.$$

The game G satisfies all the conditions of Proposition 3.2 and it follows that there are infinitely many pure strategy price equilibria. \square

3.1 Discussion of the results

An initial point worth noting about Propositions 3.1–3.4 is that the proofs of the existence results are direct and constructive. This is in contrast to many general existence results for discontinuous games, such as Reny (1999) and He and Yannelis (2015b), which give conditions on the primitives that guarantee the existence of equilibrium, but do not show how to find such equilibria. Furthermore, the equilibrium existence results in the aforementioned papers do not apply to the game studied here because in our model sellers can have tied payoffs which are greater than non-tied payoffs (when both are positive) which violates the concept of payoff security.

It is well known that in complete information price games with strictly convex costs, there may exist a continuum of pure strategy Nash equilibria (Dastidar 1995; Bagh 2010). In Proposition 3.1, we show that if each ex post game has a continuum of equilibria, and it is possible to make a selection from each of the ex post equilibrium sets which is measurable with respect to the intersection of the sellers' information algebras, then the ex ante game possesses a pure strategy price equilibrium. This is quite intuitive because if such a selection exists, then each seller knows they are playing a Nash equilibrium regardless of which state of the world occurs. In this sense, the equilibrium constructed in the proof of Proposition 3.1 is often termed "ex post stable." This means no seller would wish to change their price after the uncertainty is resolved and the state of the world is revealed to them.

Proposition 3.2 reveals the advantage of maximin expected utilities for demonstrating existence of equilibrium. Condition (ii) of the result means that there is a set of prices in each state of the world such that each seller knows if they tie with any other number of sellers then they are guaranteed non-negative profits. By selecting prices from these sets, and having the sellers play the maximum of such prices when they are uncertain of the state of the world, the sellers are guaranteed non-negative profits. The maximin expected utilities stop them wanting to undercut their rivals' prices for fear of obtaining negative profits when uncertain of the state of the world. With Bayesian expected utilities, this may not be the case, and a seller may be willing to lower their price if the expected profit is an improvement. Condition (i) in Proposition 3.2 is essentially a technical requirement which permits us to avoid "open-set" problems of the nature discussed in Blume (2003): when one seller has a cost advantage and would like to post a price as close as possible, but just below, their rival's price (and consequently the best response set is not defined). If condition (i) is satisfied, then

each seller knows there is another seller of the same type in the market who will post the same price.

In a complete information price game with a random sharing rule, the equilibrium would tend to involve the sellers tying at the monopoly breakeven price (Baye and Morgan 2002), as long as the sellers have the same cost functions. What Proposition 3.3 does is extend this concept to an incomplete information setting. Condition (ii) ensures that the sellers have the same cost functions, and hence, the same monopoly breakeven prices. In the constructed equilibrium, each seller plays the maximum monopoly breakeven prices among the states of the world which they know could occur. At these prices, sellers are guaranteed non-negative profits. Again, maximin expected utilities means sellers are not willing to lower their prices, for fear of getting negative profits. Condition (ii), as in Proposition 3.2, ensures that there are no “open-set” problems in which a seller’s best response is not properly defined.

4 Illustrative examples

This section contains three detailed examples, which demonstrate the wide variety of different price games with asymmetric information and sharing rules that the existence results in the previous section can be applied to. They should also help readers understand the specific workings of the results in the previous section.

4.1 Example 1

Consider a price game with asymmetric information with two sellers, $N = \{1, 2\}$, and two states of the world $\Omega = \{\omega_1, \omega_2\}$. The information partitions of the sellers are $P_1 = \{\{\omega_1\}, \{\omega_2\}\}$ and $P_2 = \{\Omega\}$, so seller 1 is informed of the state of the world, whereas seller 2 is totally uninformed. The probability distribution, μ , is $\mu(\omega_1) = \mu(\omega_2) = 1/2$. The cost function of the sellers is $C_i(x, \omega) = x^2$ for every $\omega \in \Omega$, $i \in N$. The sharing rule at price ties is deterministic and given by

$$g_{1,N} = \frac{3}{5} \quad \text{and} \quad g_{2,N} = \frac{2}{5}.$$

The state-contingent market demands are piecewise affine and given by $D(x, \omega_1) = \max\{0, 10 - x\}$ and $D(x, \omega_2) = \max\{0, 11 - x\}$. Given these market primitives, direct calculation of the $m_i(\omega)$ numbers in state 1 yields

$$m_1(\omega_1) = m_2(\omega_1) = 5.$$

Calculation of the $l_{i,S}(\omega)$ numbers, given the sharing rule is different from equal sharing, yields

$$l_{1,N}(\omega_1) = 3\frac{3}{4} \quad \text{and} \quad l_{2,N}(\omega_1) = 2\frac{6}{7}.$$

Therefore, $I_N(\omega_1) = (3\frac{3}{4}, 5)$. Repeating these calculations for state ω_2 yields

$$m_1(\omega_1) = m_2(\omega_1) = 5\frac{1}{2}$$

and

$$l_{1,N}(\omega_2) = 4\frac{1}{8} \quad \text{and} \quad l_{2,N}(\omega_2) = 3\frac{1}{7}.$$

Therefore, $I_N(\omega_2) = (4\frac{1}{8}, 5\frac{1}{2})$. As $P(\Sigma) = \{\Omega\}$, it follows that

$$\cap_{\omega \in \Omega} I_N(\omega) = (4\frac{1}{8}, 5) \neq \emptyset$$

and the game satisfies all the conditions of Proposition 3.1. One pure strategy price equilibrium of the game is

$$f_i^*(\omega) = 4\frac{1}{2} \quad \text{for every } \omega \in \Omega \quad \text{and} \quad i \in N.$$

4.2 Example 2

Consider a price game with asymmetric information with four sellers $N = \{1, 2, 3, 4\}$ and three states of the world $\Omega = \{\omega_1, \omega_2, \omega_3\}$. The information partitions of the sellers are

$$P_i = \{\{\omega_1, \omega_2\}, \{\omega_3\}\} \quad \text{for } i = 1, 2$$

and

$$P_i = \{\{\omega_1\}, \{\omega_2, \omega_3\}\} \quad \text{for } i = 3, 4.$$

The probability distribution, μ , is $\mu(\omega_1) = \mu(\omega_2) = \mu(\omega_3) = 1/3$. The cost function of the sellers is $C_i(x, \omega) = x^2$ for every $\omega \in \Omega$ and $i \in N$. The state-contingent market demands are piecewise affine and are $D(x, \omega_1) = \max\{0, 10 - x\}$, $D(x, \omega_2) = \max\{0, 20 - x\}$ and $D(x, \omega_3) = \max\{0, 30 - x\}$. The sharing rule is deterministic and is given by

$$g_{i,N} = \frac{1}{4} \quad \text{for every } i \in N.$$

For three seller coalitions:

$$\begin{aligned}
 g_{1,\{S\}} &= \frac{5}{12} \quad \text{if } |S| = 3 \\
 g_{i,S} &= \frac{7}{24} \quad \text{if } |S| = 3, \{1\} \in S \quad \text{and } i \neq 1 \\
 g_{i,S} &= \frac{1}{3} \quad \text{for every } i \in S \quad \text{if } |S| = 3 \quad \text{and } \{1\} \notin S.
 \end{aligned}$$

Finally, for two seller coalitions:

$$g_{i,S} = \frac{1}{2} \quad \text{for every } i \in S \quad \text{if } |S| = 2.$$

Given these primitives, after calculating the $m_i(\omega)$ and $l_{i,S}(\omega)$ numbers, one finds

$$\begin{aligned}
 \cap_{S \in C} I_S(\omega_1) &= (3\frac{1}{3}, 5) \neq \emptyset \\
 \cap_{S \in C} I_S(\omega_2) &= (6\frac{2}{3}, 10) \neq \emptyset
 \end{aligned}$$

and

$$\cap_{S \in C} I_S(\omega_3) = (10, 15) \neq \emptyset.$$

Therefore, the game satisfies all the conditions of Propositions 3.2 and 3.4. One pure strategy price equilibrium of the game is

$$f_i^*({\omega}_1, {\omega}_2) = 8 \quad \text{and} \quad f_i^*({\omega}_3) = 12 \quad \text{for } i = 1, 2$$

and

$$f_i^*({\omega}_1) = 4 \quad \text{and} \quad f_i^*({\omega}_2, {\omega}_3) = 12 \quad \text{for } i = 3, 4.$$

4.3 Example 3

Consider a price game with asymmetric information with four sellers $N = \{1, 2, 3, 4\}$ and three states of the world $\Omega = \{\omega_1, \omega_2, \omega_3\}$. The information partitions of the sellers are

$$P_i = \{\Omega\} \quad \text{for } i = 1, 2$$

and

$$P_i = \{\{\omega_1\}, \{\omega_2, \omega_3\}\} \quad \text{for } i = 3, 4.$$

The probability distribution, μ , is $\mu(\omega_1) = \mu(\omega_2) = \mu(\omega_3) = 1/3$. The cost function of the sellers is $C_i(x, \omega) = x^2$ for every $\omega \in \Omega$ and $i \in N$. The state-contingent market demands are discontinuous and are given by

$$D(x, \omega_1) = \begin{cases} 10 - x, & \text{if } x < 7; \\ \max\{0, 15 - x\} & \text{if } x \geq 7. \end{cases}$$

$$D(x, \omega_2) = \begin{cases} 20 - x, & \text{if } x < 15; \\ \max\{0, 25 - x\} & \text{if } x \geq 15. \end{cases}$$

$$D(x, \omega_3) = \begin{cases} 30 - x, & \text{if } x < 20; \\ \max\{0, 35 - x\} & \text{if } x \geq 20. \end{cases}$$

The sharing rule is random and is given by

$$h_{i,N} = \frac{1}{4} \text{ for every } i \in N.$$

For three seller coalitions:

$$h_{1,\{S\}} = \frac{5}{12} \text{ if } |S| = 3$$

$$h_{i,S} = \frac{7}{24} \text{ if } |S| = 3, \{1\} \in S \text{ and } i \neq 1$$

$$h_{i,S} = \frac{1}{3} \text{ for every } i \in S \text{ if } |S| = 3 \text{ and } \{1\} \notin S.$$

Finally for two seller coalitions:

$$h_{i,S} = \frac{1}{2} \text{ for every } i \in S \text{ if } |S| = 2.$$

Given these primitives, the game satisfies all the conditions of Proposition 3.3. One pure strategy price equilibrium of the game is

$$f_i^*(\Omega) = 15 \text{ for } i = 1, 2$$

and

$$f_i^*({\omega_1}) = 5 \text{ and } f_i^*({\omega_2, \omega_3}) = 15 \text{ for } i = 3, 4.$$

5 Related literature

This work contributes to three active strands of research: firstly, the collection of works that provide conditions for existence of pure strategy Nash equilibria in games where

sellers compete in prices to supply homogeneous products;⁹ secondly, the class of the literature that generalises equilibrium existence results to permit arbitrary sharing rules at price ties; and thirdly, the growing literature that considers the important implications of incomplete information for workhorse models of price competition. The latter remains an under-developed topic, partly due to the well-known technical tensions that arise when working with discontinuous payoffs, ambiguity and asymmetric information in competition.¹⁰ For instance, only recently has the issue of equilibrium existence under incomplete information in a Cournot oligopoly, even with continuous payoffs, been addressed by Einy et al. (2010).

In the first strand, Dastidar (1995) demonstrates that with strictly convex symmetric cost functions and equal sharing at price ties, there exists a continuum of pure strategy price equilibria where sellers earn positive profits and a single zero-profit equilibrium. With cost asymmetries amongst sellers, existence is guaranteed and may become unique. Hoernig (2002) shows that there also exists a continuum of mixed strategy equilibria. Baye and Morgan (2002) consider a model with winner-takes-all at price ties. They provide weak conditions on the market primitives, including the assumption of an initial breakeven price on the monopoly profit function that guarantee the existence of pure strategy zero-profit equilibrium.

Saporiti and Coloma (2010) consider more involved cost structures. They develop Dastidar's (1995) framework to analyze symmetric sellers with identical convex variable costs and (possibly avoidable) fixed costs. Drawing on Panzar's (1989) notion of subadditivity, they provide necessary and sufficient conditions for the existence of pure strategy Bertrand equilibria. Specifically, when the cost function is not subadditive at the output level associated with the oligopoly breakeven price, a (possibly asymmetric) pure strategy equilibrium exists. From the opposite direction, existence of pure strategy equilibrium precludes the cost function from being subadditive at every output level above the oligopoly breakeven point.

Dastidar (2011a) generalizes Saporiti and Coloma (2010) to show that symmetric superadditive cost functions guarantee the existence of pure strategy Bertrand equilibrium, without assuming convexity of variable costs. Conversely, there exists no Bertrand equilibrium if costs are symmetric and subadditive. Introducing cost-asymmetries, however, restores existence in either pure or mixed strategies (Dastidar 2011b).

Yano and Komatsubara (2017) study duopolists' choices between simultaneous-move Bertrand competition and leader-follower Stackelberg competition, when products are homogeneous but production technologies vary. Following Dastidar (1995), sellers must meet all forthcoming demand at their chosen price with equal sharing at price ties. When the market is tight, such that each seller's capacity is low relative to the market demand, Bertrand competition is driven by the (Stackelberg) first-mover advantage. When the market is less tight, the followers' disadvantage diminishes quicker for the least efficient seller, who subsequently assumes the follower

⁹ Baye and Kovenock (2008) present a detailed account of this literature.

¹⁰ See He and Yannelis (2017b, p.1421) for a list of recent work on equilibrium existence with discontinuous payoffs.

position. As tightness decreases further, the efficient seller may also accept a follower position, inducing multiple equilibria where either firm can become the leader.

The second strand considers the link between sharing rules at price ties and equilibrium existence. Hoernig (2007) considers general properties for sharing rules that permit the application of Reny's (1999) benchmark existence results for pure and mixed strategy equilibria. A pure strategy equilibrium exists whenever the sharing rule is weakly tie-decreasing, coalition monotone, the sum of payoffs is upper semi-continuous and non-tied payoffs are continuous. Tie-decreasing requires that the payoff each tied player receives is non-increasing when an additional player joins the tie, which generates quasiconcave payoffs. Coalition monotone imposes that the sum of the tied players' payoffs increases with the number of tied sellers. With symmetric non-tied payoffs, all players receive zero payoff. With asymmetric non-tied payoffs, one player may earn positive profit. Bagh (2010) provides alternative conditions for equilibrium existence for a wide range of arbitrary but deterministic sharing rules. Continuity and monotonicity requirements in Hoernig (2007) are replaced with strictly convex costs and mild assumptions on the demand function. In both Hoernig (2007) and Bagh (2010), sellers' information is complete.

The third strand concerns the growing economic recognition of the importance of ambiguity for individual choice, which can be traced to Ellsberg's (1961) famous urn experiments.¹¹ When players cannot attach probabilities to outcomes, choices need not be consistent with standard subjective expected utility maximisation and the traditional game theoretic machinery breaks down (Fellner 1961; Pulford and Colman 2007). Gilboa and Schmeidler (1989) developed Savage's (1954) subjective expected utility framework to provide an axiomatic characterisation of maximin expected utility, according to which decision making is driven by the minimum utility across the set of probabilities that are consistent with the available information (Al-Najjar and Weinstein 2009).

Incomplete information has drastic implications. For instance, the Bertrand paradox of marginal cost pricing with only two sellers, is quickly overturned. When sellers have symmetric uncertainty regarding their rivals' costs but know their own cost, equilibrium involves pricing above cost (Spulber 1995). When marginal costs are constant but uncertain and sellers are risk averse, losses are overweighted in the expected utility function. Therefore, equilibrium must involve positive expected profits that persist as market concentration decreases (Wambach 1999). The disruptive effects of incomplete information are not confined to price competition. Einy et al. (2010) consider Bayesian output-setting oligopolists with incomplete information regarding state-contingent production costs and market demand. They provide conditions for the existence and uniqueness of pure strategy Nash equilibrium without permitting negative prices.

Maximin preferences address the growing discontent with Bayesian utilities under uncertainty and feature several favourable properties (He and Yannelis 2015a). Firstly, in discontinuous games with incomplete information and Bayesian players, equilibrium existence often follows from strengthening the restrictions on the conditions in

¹¹ See Al-Najjar and Weinstein (2009) for a critical assessment of the ambiguity aversion literature and Sugden (1991) for a survey of the roots of rational choice.

the complete information analogue. For instance, He and Yannelis (2015b) introduce a stronger notion of Reny's (1999) payoff security and impose quasiconcave expected utilities to guarantee equilibrium existence with Bayesian players under asymmetric and incomplete information. He and Yannelis (2017a), however, recently demonstrated that maximin utilities overcome the need to introduce additional assumptions. Specifically, maximin equilibrium existence "follows immediately from the existence of equilibrium in every ex-post game" (He and Yannelis 2017a, p.120). Moreover, if we were to impose the conditions on non-tied payoffs outlined in Hoernig (2007) and remove our requirement of measurability of prices with respect to sellers' private information, equilibrium existence follows directly from He and Yannelis (2017a, Proposition 1).

Secondly, maximin preferences overcome the well-known conflict between incentive compatibility and efficiency that arises with Bayesian expected utility preferences (de Castro and Yannelis 2011; Holmstrom and Myerson 1983). Moreover, maximin is the unique preference relation that ensures any efficient outcome is incentive compatible. The intuition follows that in equilibrium agents equalise utilities across all possible states. Therefore, individuals have no incentive to misreport their private information once it is received.¹²

Correia-da-Silva and Hervés-Beloso (2009) study maximin preferences in an extension of an Arrow–Debreu asymmetric information general equilibrium model. They analyse trade contracts when individuals choose between lists of bundles of products. The bundle that is delivered is state dependent, where agents hold incomplete and asymmetric information over future states. Decisions are prudent in the sense that agents consider the worst possible outcome and this expectation may be confirmed. In equilibrium individuals, choose bundles that yield constant utility across uncertain states. Therefore, consumption decisions need not be measurable with respect to the information of each individual but the contingent lists must be. Correia-da-Silva and Hervés-Beloso (2012) depart from prudent equilibrium to consider agents who always receive (and, therefore, expect to receive) the cheapest consumption bundle on the list. They impose that at least one agent can verify the true state, which prevents false information disclosures. They establish equilibrium existence in such an environment. de Castro et al. (2017) also study implementation under ambiguity using maximin preferences. They show that maximin equilibrium is consistent with individually rational non-cooperative strategies under ambiguity.

6 Concluding remarks

The aim of this work has been to analyze a general incomplete information extension of the classical Bertrand price game which permits a wide variety of sharing rules at price ties and asymmetries of information amongst the sellers of the type usually only studied in the general equilibrium literature. The results indicate that, if no seller has a monopoly over their private information, sellers' costs are not too different, and the sharing rule at price ties is not too different from the equal sharing rule, then there exists

¹² See de Castro and Yannelis (2011), Sect. 1, for an appraisal of the applications of maximin preferences.

a continuum of pure strategy price equilibria. In this sense, the model demonstrates that, even if sellers have asymmetric information about the state of the world, and the sharing rule is different from equal sharing, there may still be positive existence results similar to those presented by Dastidar (1995), and generalized in the subsequent literature on price competition. Furthermore, our results permitted discontinuities in both the monopoly and tied profit functions arising from discontinuities in the market demand. We conclude this work with several possibilities for extending the model which the authors hope to pursue in future research.

- Although maximin decision rules have been recently applied to a wide range of economic models, there is an alternative way of modelling ambiguity which is to use capacities. A capacity has the properties of a probability measure, but may violate the standard additivity properties. This alternative way of modelling probability beliefs can be parametrized, as in Eichberger and Kelsey (1999), with the parameter measuring the degree of confidence an individual has about the different possible probabilities they face. Our model has only considered maximin preferences, which although simple and tractable, represent the polar extreme in which individuals are wholly unable to assign beliefs to unknown probabilities. It would be interesting to see whether any of the results could be extended to permit capacity-type modelling of ambiguity beliefs.
- The results in this work indicate that there may be a continuum of pure strategy price equilibria. However, these results do not indicate if some of these pure strategy price equilibria are more convincing, or have more desirable properties, than others. One particular refinement of the Nash equilibrium set which has been applied to Bertrand price games is coalition-proofness.¹³ The idea of coalition-proof equilibrium was introduced by Bernheim et al. (1987) and is appropriate in settings where players can communicate with each other to enact coalitional strategy deviations, but cannot commit to these deviations.¹⁴ As a result, any proposed deviations have to be self-enforcing. Chowdhury and Sengupta (2004) established the existence of coalition-proof equilibrium in homogeneous-good Bertrand games under quite general conditions. In the incomplete information Bertrand game studied here, we could admit the possibility that sellers could communicate at the ex ante stage to enact coalition deviations. What the appropriate definition of coalition-proofness should be in this context and whether such a refinement of the equilibrium set exists are interesting open questions.¹⁵
- Throughout the paper, we have focussed on maximin expected utilities at the ex ante stage. As was noted, these remedy some of the difficult technical issues which arise with Bayesian utilities. But it remains an interesting open question whether equilibrium existence could be established with Bayesian expected utilities. Unfortunately, the existence results of He and Yannelis (2015b) cannot be

¹³ Although recently Andersson et al. (2014) have studied an alternative refinement of the equilibrium set in price games, based on strategic uncertainty, and have shown that it selects prices different from the coalition-proof equilibrium.

¹⁴ Coalition-proofness has been applied to a wide range of different games, including cooperative games, as in Habis and Herings (2011).

¹⁵ In defining coalition-proof equilibrium at the ex ante stage, one would have to take into account the possibility that sellers could share their private information as in Wilson (1978).

applied to such a game because non-tied payoffs can be higher than tied payoffs and, consequently, the Bertrand game tends not to be better reply secure. Indeed, even the complete information game studied in Dastidar (1995) is not better reply secure. However, it may be possible to understand whether an equilibrium exists with Bayesian expected utilities by exploiting the well-known properties of the ex post price games.

- Although we have assumed that the sellers tying at the minimum price meet all the market demand forthcoming, there is a large literature on price games with capacity constraints and rationing rules: the Bertrand–Edgeworth approach to price competition. This literature has provided important insights into the competitive limit in large markets, and when we should expect price-making behaviour to result in outcomes close to the competitive equilibrium of markets. As far as the authors are aware, there has been little research on extending the standard Bertrand–Edgeworth price game to permit asymmetric information. This is probably due to the technical difficulties as equilibria in Bertrand–Edgeworth games are usually only in mixed strategies.¹⁶ It would be interesting, however, to see whether the Bertrand–Edgeworth game could be generalized to permit asymmetries of information of the type we have studied here.

Open Access This article is distributed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made.

Appendix

This appendix contains three lemmas, and their proofs, which were used to prove the main results.

Lemma 7.1 *Fix a price game $G = (N, \Omega, (C_i, P_i)_{i \in N}, \{r_{i,S}\}_{i \in N, S \in S_i}, D, \mu), r \in \{g, h\}$. Then*

- (i) *the real numbers $l_{i,S}(\omega)$ and $m_i(\omega)$ are well defined.*
- (ii) *if $r = g$, the sharing rule is deterministic, then $l_{i,S}(\omega) < m_i(\omega)$.*

Proof (i) From Assumption 2.2, $D(0, \omega) = \bar{y}(\omega) > 0$. Hence, $\pi_{i,S}(0, \omega) < 0$ and $\pi_i(0, \omega) < 0$. From Assumption 2.4, the functions achieve strictly positive values and are left lower semicontinuous, so the sets

$$L_{i,S}(\omega) = \{x | 0 \leq x < \bar{x}(\omega) \text{ and } \pi_{i,S}(x, \omega) = 0\}$$

and

$$M_i(\omega) = \{x | 0 \leq x < \bar{x}(\omega) \text{ and } \pi_i(x, \omega) = 0\}.$$

¹⁶ If one imposes strong conditions upon the demand function, then pure strategy price equilibria may exist even with capacity constraints (Tasnadi 1999)

are non-empty and their minimums are well defined. Let $l_{i,S}(\omega) = \min L_{i,S}(\omega)$ and $m_i(\omega) = \min M_i(\omega)$.

(ii) Suppose $l_{i,S}(\omega) = a$ and $m_i(\omega) = b$. From the definition of $l_{i,S}(\omega)$, we have

$$a g_{i,S}(a) D(a, \omega) - C_i(g_{i,S}(a) D(a, \omega)) = 0$$

which implies

$$a = \frac{C_i(g_{i,S}(a) D(a, \omega))}{g_{i,S}(a) D(a, \omega)}.$$

Similarly, from the definition of b , we have

$$b D(b, \omega) - C_i(D(b, \omega)) = 0$$

which implies

$$b = \frac{C(D(b, \omega))}{D(b, \omega)}.$$

Suppose a contradiction to the statement in the lemma: $m_i(\omega) \leq l_{i,S}(\omega)$. Then, $b \leq a$ and

$$\frac{C_i(D(b, \omega))}{D(b, \omega)} \leq \frac{C_i(g_{i,S}(a) D(a, \omega))}{g_{i,S}(a) D(a, \omega)}.$$

However, as $b \leq a$, $D(b, \omega) \geq D(a, \omega) > g_{i,S}(a) D(a, \omega)$, this contradicts the strict convexity of the cost function $C_i(\cdot, \omega)$. Hence, $l_{i,S}(\omega) < m_i(\omega)$. □

Lemma 7.2 Fix a price game $G = (N, \Omega, (C_i, P_i)_{i \in N}, \{r_{i,S}\}_{i \in N, S \in S_i}, D, \mu)$ with a deterministic sharing rule $r = g$, if $x' > l_{i,S}(\omega)$ then $\pi_{i,S}(x', \omega) \geq 0$.

Proof If $x' \geq \bar{x}(\omega)$ then $D(x', \omega) = 0$ and $\pi_{i,S}(x', \omega) = 0$. Suppose $x' > l_{i,S}(\omega)$ and $x' < \bar{x}(\omega)$. As $x' > l_{i,S}(\omega)$, there is an x'' such that $x' > x'' > l_{i,S}(\omega)$ and $\pi_{i,S}(x'', \omega) \geq 0$. As $\pi_{i,S}(x'', \omega) \geq 0$, we have

$$x'' \geq \frac{C_i(g_{i,S}(x'') D(x'', \omega))}{g_{i,S}(x'') D(x'', \omega)}.$$

As $x' > x''$, the strict convexity of the cost function and Assumption 3.1 yields

$$\frac{C_i(g_{i,S}(x'') D(x'', \omega))}{g_{i,S}(x'') D(x'', \omega)} \geq \frac{C_i(g_{i,S}(x') D(x', \omega))}{g_{i,S}(x') D(x', \omega)}.$$

Hence

$$x' > x'' \geq \frac{C_i(g_{i,S}(x'') D(x'', \omega))}{g_{i,S}(x'') D(x'', \omega)} \geq \frac{C_i(g_{i,S}(x') D(x', \omega))}{g_{i,S}(x') D(x', \omega)}$$

which implies $\pi_{i,S}(x', \omega) > 0$. □

Lemma 7.3 Fix a price game $G = (N, \Omega, (C_i, P_i)_{i \in N}, \{r_{i,S}\}_{i \in N, S \in S_i}, D, \mu)$. If $x' > m_i(\omega)$ then $\pi_i(x', \omega) \geq 0$.

Proof If $x' \geq \bar{x}(\omega)$ then $D(x', \omega) = 0$ and $\pi_i(x', \omega) = 0$. Suppose $x' > m_i(\omega)$ and $x' < \bar{x}(\omega)$. As $x' > m_i(\omega)$, there is an x'' such that $x' > x'' > m_i(\omega)$ and $\pi_i(x'', \omega) \geq 0$. As $\pi_i(x'', \omega) \geq 0$, we have

$$x'' \geq \frac{C_i(D(x'', \omega))}{D(x'', \omega)}.$$

As $x' > x''$, again the strict convexity of the cost function yields

$$\frac{C_i(D(x'', \omega))}{D(x'', \omega)} \geq \frac{C_i(D(x', \omega))}{D(x', \omega)}.$$

Hence

$$x' > x'' \geq \frac{C_i(D(x'', \omega))}{D(x'', \omega)} \geq \frac{C_i(D(x', \omega))}{D(x', \omega)}$$

which implies $\pi_i(x', \omega) > 0$. □

References

- Allingham, M. (2002). *Choice theory: a very short introduction*. New York: Oxford University Press.
- Al-Najjar, N. I., & Weinstein, J. (2009). The ambiguity aversion literature: a critical assessment. *Economics and Philosophy*, 25(3), 249–284.
- Andersson, O., Argenton, C., & Weibull, J. (2014). Robustness to strategic uncertainty. *Games and Economic Behavior*, 85, 272–288.
- Bagh, A. (2010). Pure strategy equilibria in Bertrand games with discontinuous demand and asymmetric tie-breaking rules. *Economics Letters*, 108, 277–279.
- Baye, M., & Kovenock, D. (2008). Bertrand competition. In S. Durlauf & L. Blume (Eds.), *The new Palgrave dictionary of economics* (2nd ed.). Palgrave Macmillan.
- Baye, M., & Morgan, J. (2002). Winner-take-all price competition. *Economic Theory*, 19, 271–282.
- Bernheim, D., Peleg, B., & Whinston, M. (1987). Coalition-proof Nash equilibria I: concepts. *Journal of Economic Theory*, 42, 1–12.
- Blume, A. (2003). Bertrand without fudge. *Economics Letters*, 78(2), 167–168.
- Cerreia-Vioglio, S., Maccheroni, F., Marinacci, M., & Montrucchio, L. (2013). Ambiguity and Robust statistics. *Journal of Economic Theory*, 148(3), 974–1049.
- Chowdhury, P., & Sengupta, K. (2004). Coalition-proof Bertrand equilibria. *Economic Theory*, 24, 307–324.
- Correia-da-Silva, J., & Hervés-Beloso, C. (2009). Prudent expectations equilibrium in an economy with uncertain delivery. *Economic Theory*, 39, 67–92.
- Correia-da-Silva, J., & Hervés-Beloso, C. (2012). General equilibrium with uncertain delivery. *Economic Theory*, 51, 729–755.
- Dastidar, K. (1995). On the existence of pure strategy Bertrand equilibrium. *Economic Theory*, 5, 19–32.
- Dastidar, K. (2011a). Existence of Bertrand equilibrium revisited. *International Journal of Economic Theory*, 7, 331–350.
- Dastidar, K. (2011b). Bertrand equilibrium with subadditive costs. *Economics Letters*, 112(2), 202–204.
- de Castro, L. I. & N. C. Yannelis, (2011) Ambiguity aversion solves the conflict between Efficiency and incentive compatibility, University of Manchester Discussion Paper.
- de Castro, L., Liu, Z., & Yannelis, N. C. (2017). Implementation under ambiguity. *Games and Economic Behavior*, 101, 20–33.

- Debreu, G., & Scarf, H. (1963). A limit theorem on the core of an economy. *International Economic Review*, 4, 235–246.
- Edwards, R. A. (2019). Pricing and obfuscation with complexity averse consumers. *Oxford Economic Papers*, 71, 777–798.
- Eichberger, J., & Kelsey, D. (1999). E-capacities and the Ellsberg paradox. *Theory and Decision*, 46, 107–140.
- Einy, E., Haimanko, O., Moreno, D., & Shitovitz, B. (2010). On the existence of Bayesian cournot equilibrium. *Games and Economic Behavior*, 68(1), 77–94.
- Ellsberg, D. (1961). Risk, ambiguity and the savage axioms. *Quarterly Journal of Economics*, 75, 643–669.
- Fellner, W. (1961). Distortion of subjective probabilities as a reaction to uncertainty. *Quarterly Journal of Economics*, 75(4), 670–689.
- Gilboa, I., & Schmeidler, D. (1989). Maximin expected utility with non unique prior. *Journal of Mathematical Economics*, 18, 141–53.
- Habis, H., & Herings, P. J. J. (2011). Transferable utility games with uncertainty. *Journal of Economic Theory*, 146(5), 2126–2139.
- He, W., & Yannelis, N. C. (2015a). Equilibrium theory under ambiguity. *Journal of Mathematical Economics*, 61, 86–95.
- He, W., & Yannelis, N. C. (2015b). Discontinuous games with asymmetric information: an extension of Reny's existence theorem. *Games and Economic Behavior*, 91, 26–35.
- He, W., & Yannelis, N. C. (2016). Existence of equilibria in discontinuous Bayesian games. *Journal of Economic Theory*, 162, 181–194.
- He, W., & Yannelis, N. C. (2017a). A remark on discontinuous games with asymmetric information and ambiguity. *Economic Theory Bulletin*, 5(1), 119–126.
- He, W., & Yannelis, N. C. (2017b). Equilibria with discontinuous preferences: new fixed point theorems. *Journal of Mathematical Analysis and Applications*, 450, 1421–1433.
- Hoernig, S. (2002). Mixed Bertrand equilibria under decreasing returns to scale: an embarrassment of riches. *Economics Letters*, 74(3), 359–362.
- Hoernig, S. (2007). Bertrand games and sharing rules. *Economic Theory*, 31, 573–585.
- Holmstrom, B., & Myerson, R. B. (1983). Efficient and durable decision rules with incomplete information. *Econometrica*, 51(6), 1799–1819.
- Panzar, J. (1989). Technological determinants of firm and industry structure. *Handbook of Industrial Organization*, 1, 4–59.
- Peterson, M. (2017). *An Introduction to Decision Theory*. Cambridge: Cambridge University Press.
- Pulford, B. D., & Colman, A. M. (2007). Ambiguous games: evidence for strategic ambiguity aversion. *Quarterly Journal of Experimental Psychology*, 60, 1083–1100.
- Reny, P. (1999). On the existence of pure and mixed strategy nash equilibria in discontinuous games. *Econometrica*, 67, 1029–1056.
- Saporiti, A., & Coloma, G. (2010). Bertrand competition in markets with fixed costs. *B.E. Journal of Theoretical Economics*, 10(1), 27.
- Savage, L. J. (1954). *The foundations of statistics*. New York: Wiley.
- Spulber, D. (1995). Bertrand competition when Rivals' costs are unknown. *Journal of Industrial Economics*, 43, 1–11.
- Sugden, R. (1991). Rational choice: a survey of contributions from economics and philosophy. *The Economic Journal*, 101(407), 751–785.
- Tasnadi, A. (1999). Existence of pure strategy nash equilibrium in Bertrand–Edgeworth oligopolies. *Economics Letters*, 63(2), 201–206.
- Vives, X. (1999). *Oligopoly pricing: old ideas and new tools*. Cambridge: MIT Press.
- Wambach, A. (1999). Bertrand competition under cost uncertainty. *International Journal of Industrial Organization*, 17, 941–951.
- Wilson, R. (1978). Information, efficiency and the core of an economy. *Econometrica*, 4, 807–816.
- Yano, M., & Komatsubara, T. (2018). Price competition or price leadership. *Economic Theory*, 66(4), 1023–1057.