



Firms' Beliefs and Learning: Models, Identification, and Empirical Evidence

Victor Aguirregabiria¹ · Jihye Jeon²

Published online: 31 July 2019

© Springer Science+Business Media, LLC, part of Springer Nature 2019

Abstract

This paper reviews recent literature on structural models of oligopoly competition where firms have biased beliefs about the primitives of the model—e.g. demand, costs—or about the strategic behavior of other firms in the market. We describe different structural models that have been proposed to study this phenomenon and examine the approaches that have been used to identify firms' beliefs. We discuss empirical results in recent studies and show that accounting for firms' biased beliefs and learning can have important implications on our measures and interpretation of market efficiency.

Keywords Dynamics · Identification · Learning · Non-equilibrium beliefs · Oligopoly competition · Structural models

JEL Classification C57 · D81 · D83 · D84 · L13

1 Introduction

Firms have uncertainty about current and future realizations of demand, costs, market regulations, and/or the behavior of competitors. They may learn about these elements over time, and this learning process can have substantial implications for their profits and market efficiency. For example, firms may gather better information and

✉ Victor Aguirregabiria
victor.aguirregabiria@utoronto.ca

Jihye Jeon
jjeon@bu.edu

¹ University of Toronto and CEPR, 150 St. George Street., Toronto, ON M5S 3G7, Canada

² Boston University, 270 Bay State Road, Boston, MA 02215, USA

use it in their pricing, entry, and investment strategies to improve their profits and the probability of survival in the market. In today's economies, big data and learning algorithms are becoming important inputs in many industries.¹

The assumption of rational expectations has been the status quo to represent agents' beliefs in many areas in economics, and in particular in industrial organization. Rational equilibrium models of competition incorporate two basic assumptions: Firms form beliefs about uncertain demand, costs, and the behavior of competitors, and given their beliefs the firms maximize expected profits; and firms' beliefs coincide with the actual probability distribution of demand, costs, and competitors' behavior, i.e., rational equilibrium beliefs.

The assumption of rational expectations has the attractive feature that beliefs are endogenously determined in the equilibrium of the model. It also provides econometric identification of belief functions. Nevertheless, this assumption is frequently criticized for ignoring that information acquisition and processing is costly (e.g., Pesaran 1987; Manski 2004). In reality, firms often face significant uncertainty about their rivals' strategies. Firms are different in their ability for collecting and processing information—for similar reasons as they are heterogeneous in their costs of producing goods. As a result, firms are also heterogeneous in their degree of uncertainty and in their speed of learning.

The importance of firms' heterogeneous expectations and the implications for firms' performance and market outcomes have been long recognized in economics, at least since the work of Simon (1958, 1959) and his debate with Muth (1961) on the relative merits of adaptive/behavioral expectations versus rational expectations. However, it has not been until recently that this behavioral approach has received substantial attention in structural models in empirical industrial organization. Part of the reason is that the joint identification of firms' beliefs and structural parameters in payoffs is more challenging when we relax the assumption of rational expectations. Nevertheless, the combination of better data, identification strategies, and econometric techniques has made possible the recent developments in structural models of competition with boundedly rational firms.

An active research area in economics has consisted in modeling firms' beliefs—allowing for their endogeneity but relaxing the assumption of rational expectations. A goal in this literature is to develop useful models that allow for heterogeneity in firms' information and that incorporate the process through which firms acquire knowledge about the environment by searching and processing information and ultimately learning.

¹ Bajari et al. (2018) study the impact of big data on firm performance by examining how the amount of data affects the accuracy of retail product forecasts. Their results show that increases in the time dimension of the data (T) improve demand forecasts, though with diminishing returns to scale. In contrast, increases in the number of products in the data (N) do not generate significant improvements in demand forecasts, which is a surprising and controversial result. They also report that the firm's overall forecast performance—controlling for N and T —has improved over time as the result of better models and methods.

In industrial organization—based on the seminal work of Jovanovic (1982)—there has been work in the development of theoretical models of industry dynamics with Bayesian learning. Macroeconomists have incorporated Bayesian learning (e.g., Evans and Ramey 1992) and adaptive learning (e.g., Sargent 1993, and Evans and Honkapohja 2001, 2012) in dynamic general equilibrium models, and have established conditions for convergence to a rational expectations equilibrium (e.g., Marcet and Sargent 1989a, b). In experimental economics, there is evidence on players' non-equilibrium beliefs from games played in laboratory experiments (e.g., Van Huyck et al. 1990; Heinemann et al. 2009). Behavioral and experimental economists have also developed non-equilibrium models of endogenous beliefs (e.g., the Cognitive Hierarchy model by Camerer et al. 2004, and the Level-k rationality model by Costa-Gomes and Crawford 2006, and Crawford and Iriberri 2007) and learning in games (e.g., experience-weighted attraction learning by Camerer and Ho 1999), and they have estimated structurally these models using data from laboratory experiments. In microeconometrics, there is a large body of work on measuring expectations and using expectations data to relax or validate assumptions about beliefs (Guiso and Parigi 1999; Manski 2004, 2018; Pesaran and Weale 2006; Bover 2015).

Building on these literatures, there has been an increasing interest in empirical industrial organization in the development and estimation of structural models of market competition where firms have non-equilibrium beliefs and try to learn over time, with potentially different degrees of success. Motivated by this interest, there has been also work in the econometrics of games and on the identification of beliefs and payoffs in non-equilibrium models. This paper reviews this recent literature.

In the context of games, firms' learning is related to the game's information structure and how this structure evolves over time. In games with learning, biased beliefs can evolve endogenously to become equilibrium beliefs, and an incomplete information environment may converge to complete information. Learning is also related to multiplicity of equilibria in games. In a model of competition with multiple equilibria, firms are engaged in an adaptive process where they learn how to play an equilibrium of the game (Fudenberg and Levine 1998; Fershtman and Pakes 2012). The specification of a learning process can allow researchers to identify an equilibrium selection mechanism, and this can be an attractive solution to deal with the problem of multiple equilibria in comparative statics exercises using these models (Lee and Pakes 2009).

In the learning literature that we review in this paper, agents (firms) know the model that describes the environment but they have uncertainty about some elements of this environment. Agents may learn over time about these elements when new information arrives. This type of model contrasts with the approach in *evolutionary game theory* to represent learning (Samuelson 1998). Models in evolutionary game theory do not assume that players know the nature of the game, or even that they know that they are playing a game. These models do not specify any particular element that is the object of agents' learning. Instead, evolutionary games in economics consider that agents' actions (strategies) are driven by imitation, either of other agents or of their own behavior in situations that they perceive as similar. Though our paper does not review this important literature, in Sect. 2.2, we discuss

the relationship between *reinforcement learning* algorithms and the concept of learning in evolutionary games.

The theoretical and empirical literature on firms' learning is large and spreads over multiple areas in economics. In this paper, we focus on a specific part of this literature that deals with empirical applications of structural games of oligopoly competition that allow for firms' non-equilibrium beliefs.

There are several recent surveys that cover related topics but with different emphases. Heidhues and Köszegi (2018) provide an excellent survey on behavioral industrial organization with special emphasis on models of consumer demand (see their section 6 that covers supply models). Crawford et al. (2013) review the theoretical and empirical literature of structural models of non-equilibrium strategic thinking in experimental economics, with a focus on evidence from laboratory experiments, but with some discussion of the more scarce evidence from field data. Ching et al. (2017) review the marketing literature on dynamic structural models with Bayesian learning that has concentrated on consumer demand but has some applications to firms' learning (see their Sect. 4). Goldfarb et al. (2012) provide a brief discussion of recent empirical applications of structural non-equilibrium games. Borkovsky et al. (2015) examine several methodological topics in empirical games of oligopoly competition, including firms' biased beliefs and learning.

Beyond the topic of biased beliefs and learning, there is a growing interest in empirical research on structural behavioral models. The recent survey by DellaVigna (2018) provides an excellent review of this literature that includes important methodological aspects on modelling, identification, estimation, and sensitivity analysis.

The rest of the paper is organized as follows. Section 2 presents a model of competition with firms' learning that incorporates the main features in the empirical applications that we cover in this paper. In Sect. 3, we discuss different types of data and identification strategies. Sect. 4 reviews recent empirical applications. We conclude in Sect. 5.

2 Model

In this section we present a framework of firm behavior under uncertainty that incorporates the main features of recent empirical applications on firms' non-equilibrium beliefs and learning. We use this framework to define different types of uncertainty—e.g., on exogenous variables and on the endogenous strategies of competitors—and to introduce different assumptions on firms' beliefs and learning.

The model that we present here is dynamic, in discrete time, and with an infinite horizon, although our discussion of empirical applications in Sect. 4 also includes static games. We can see static games as a particular case of our framework when firms have zero discount factors.

We start with a dynamic model for a single firm in a monopolistic market. Then, we extend the model to a dynamic oligopoly framework. This framework follows the dynamic game of oligopoly competition in the seminal work by Ericson and Pakes (1995), and in subsequent extensions of this framework that incorporate asymmetric

information (Aguirregabiria and Mira 2007; Doraszelski and Satterthwaite 2010; Fershtman and Pakes 2012).

2.1 A Model for a Single Firm

Consider a monopolistic firm that makes an investment decision $a_t \in \mathcal{A}$ at each period $t \in \{1, 2, \dots\}$ to maximize its expected and discounted intertemporal profit, or value. The per-period profit function is $\pi(a_t, k_t, \mathbf{z}_t)$ where k_t is the capital stock, and \mathbf{z}_t is a vector of exogenous state variables that affect demand and costs. The firm has uncertainty about future realizations of the capital stock and of the exogenous state variables.² We assume that capital stock follows a controlled Markov process with transition probability function $p_k(k_{t+1}|a_t, k_t)$. The vector of state variables \mathbf{z}_t follows an exogenous Markov process with transition probability function $p_z(\mathbf{z}_{t+1}|\mathbf{z}_t)$. For notational convenience, we use vector \mathbf{x}_t to represent all the state variables— $\mathbf{x}_t \equiv (k_t, \mathbf{z}_t)$ —and we use function $p_x(\mathbf{x}_{t+1}|a_t, \mathbf{x}_t)$ to represent the true transition probability that governs the evolution of all the state variables: $p_x(\mathbf{x}_{t+1}|a_t, \mathbf{x}_t) \equiv p_k(k_{t+1}|a_t, k_t) p_z(\mathbf{z}_{t+1}|\mathbf{z}_t)$.

The firm's belief about the evolution of the state variables can be different from the true transition probability. Let $b_{\mathcal{H}_t}(\mathbf{x}_{t+1}|a_t, \mathbf{x}_t)$ be the function that represents the firm's belief at period t about the transition of the state variables. This is a transition probability function with arguments $(\mathbf{x}_{t+1}|a_t, \mathbf{x}_t)$. The belief function is indexed by \mathcal{H}_t , which represents the history of past decisions and state variables, $\mathcal{H}_t \equiv \{a_{s-1}, \mathbf{x}_s : s \leq t\}$. The shape of the belief function can be updated as the information set \mathcal{H}_t evolves over time according to its transition rule, $\mathcal{H}_{t+1} = \mathcal{H}_t \cup \{a_t, \mathbf{x}_{t+1}\}$. The nature of this belief updating depends on the particular model of learning. We describe below different learning models.

In principle, the firm's uncertainty might include its own investment behavior in the future. The standard assumption is that the firm believes that, at any period $t + s$ in the future, it will choose investment a_{t+s} to maximize its value given information \mathcal{H}_{t+s} , i.e., $a_{t+s} = \alpha_{\mathcal{H}_{t+s}}^*(\mathbf{x}_{t+s})$ for any $s \geq 0$, where $\alpha_{\mathcal{H}}^*(\mathbf{x})$ represents the optimal decision rule given information \mathcal{H} . Similarly, $V_{\mathcal{H}_t}(\mathbf{x}_t)$ is the value of the firm at period t given current information and beliefs and under the condition that the firm maximizes expected profits forever in the future. For the sake of simplicity, we adopt the notation $b_t(\mathbf{x}_{t+1}|a_t, \mathbf{x}_t)$, $\alpha_t^*(\mathbf{x}_t)$, $V_t(\mathbf{x}_t)$ to represent belief functions, optimal decision rules, and value functions, respectively; but the reader should keep in mind that the time subindex represents the information set \mathcal{H}_t . By Bellman's principle, the value of the firm can be represented recursively as:

$$V_t(\mathbf{x}_t) = \max_{a_t \in \mathcal{A}} \left\{ \pi(a_t, \mathbf{x}_t) + \beta \int V_{t+1}(\mathbf{x}_{t+1}) b_t(\mathbf{x}_{t+1}|a_t, \mathbf{x}_t) d\mathbf{x}_{t+1} \right\} \quad (1)$$

² Uncertainty about the next period capital stock can be associated with stochastic depreciation or with randomness in the transformation of current investment into next period productive capital. For instance, capital can follow the transition rule $k_{t+1} = (1 - \delta_{t+1})k_t + \gamma_{t+1} a_t$, where δ_{t+1} and γ_{t+1} are independently and identically distributed random variables that are unknown to the firm at period t .

The optimal decision rule, $\alpha_t^*(\mathbf{x}_t)$, is the arg-max operator of the expression in brackets $\{ \cdot \}$.

To complete the model, we need to make assumptions about the evolution of the belief function over time. We describe here some of the approaches in the literature that have been most commonly used to model beliefs in single-agent models.

(i) *Rational expectations* The belief function is the actual probability distribution of the state variables. At every period t , $b_t(\mathbf{x}_{t+1}|a_t, \mathbf{x}_t) = p_x(\mathbf{x}_{t+1}|a_t, \mathbf{x}_t)$. That is, at every period t the firm's information set includes the true transition probability of the state variables.

(ii) *Bayesian learning* At the initial period, $t = 0$, the firm has a prior belief function $b_0(\mathbf{x}_1|a_0, \mathbf{x}_0)$ that is exogenous. We describe this prior as a mixture over a collection of L transition probabilities $\mathcal{P} \equiv \{\psi_\ell(\mathbf{x}_{t+1}|a_t, \mathbf{x}_t) : \ell = 1, 2, \dots, L\}$. That is, $b_0(\mathbf{x}_1|a_0, \mathbf{x}_0) = \sum_{\ell=1}^L \lambda_\ell^{(0)}(\mathbf{x}_1, a_0, \mathbf{x}_0) \psi_\ell(\mathbf{x}_1|a_0, \mathbf{x}_0)$, where the set \mathcal{P} and the mixing probabilities $\{\lambda_\ell^{(0)}\}$ that define the prior are exogenously given.³

At any period $t \geq 1$, the firm proceeds as follows. First, it observes the new state \mathbf{x}_t and uses this information to update its beliefs by using Bayes rule. More specifically, the Bayesian updating of the mixing probabilities at period t is:

$$\lambda_\ell^{(t)}(\mathbf{x}_t, a_{t-1}, \mathbf{x}_{t-1}) = \frac{\psi_\ell(\mathbf{x}_t|a_{t-1}, \mathbf{x}_{t-1}) \lambda_\ell^{(t-1)}(\mathbf{x}_t, a_{t-1}, \mathbf{x}_{t-1})}{\sum_{\ell'=1}^L \psi_{\ell'}(\mathbf{x}_t|a_{t-1}, \mathbf{x}_{t-1}) \lambda_{\ell'}^{(t-1)}(\mathbf{x}_t, a_{t-1}, \mathbf{x}_{t-1})} \quad (2)$$

and the beliefs function becomes $b_t(\mathbf{x}'|a, \mathbf{x}) = \sum_{\ell=1}^L \lambda_\ell^{(t)}(\mathbf{x}', a, \mathbf{x}) \psi_\ell(\mathbf{x}'|a, \mathbf{x})$ for any value $(\mathbf{x}', a, \mathbf{x})$. Second, given \mathbf{x}_t and the new belief function b_t , the firm chooses investment a_t to maximize its expected value. If the true probability distribution belongs to the set \mathcal{P} , then under mild regularity conditions this Bayesian learning process implies that b_t converges to the true probability (e.g., Cyert and DeGroot 1974).

(iii) *Adaptive learning* (Sargent 1993; Evans and Honkapohja 2001, 2012). There are different models of adaptive learning. Their common feature is that beliefs are updated according to an ad-hoc rule that weights the past and the new information. For instance, for any value $(\mathbf{x}', a, \mathbf{x})$ the updating of the belief function at period t is:

$$b_t(\mathbf{x}'|a, \mathbf{x}) = (1 - \delta_t) b_{t-1}(\mathbf{x}'|a, \mathbf{x}) + \delta_t K([\mathbf{x}_t, a_{t-1}, \mathbf{x}_{t-1}] - [\mathbf{x}', a, \mathbf{x}]) \quad (3)$$

where $\{\delta_t : t \geq 0\}$ with $\delta_t \in [0, 1)$ is a sequence of parameters that determines the speed of learning, and $K(\cdot)$ is a Kernel function that establishes whether the new information at period t is used to update beliefs only at that point or also at nearby values. The values of the sequence of parameters $\{\delta_t\}$ and the form of the kernel function K define a specific adaptive learning algorithm.⁴ In general, and in

³ Note that the mixing probabilities are different for each value of $(\mathbf{x}_1, a_0, \mathbf{x}_0)$.

⁴ Some examples of kernel functions for adaptive learning are: (i) point kernel: $K([\mathbf{x}_t, a_{t-1}, \mathbf{x}_{t-1}] - [\mathbf{x}', a, \mathbf{x}])$ is the indicator function $1\{[\mathbf{x}_t, a_{t-1}, \mathbf{x}_{t-1}] = [\mathbf{x}', a, \mathbf{x}]\}$; (ii) Gaussian: $\phi([\mathbf{x}_t - \mathbf{x}', a_{t-1} - a, \mathbf{x}_{t-1} - \mathbf{x}]\beta)$, where $\phi(\cdot)$ is the density of the standard normal, and β is a vector of weighting parameters.

contrast to Bayesian learning, this updating rule is not using the new information $\{\mathbf{x}_t, a_{t-1}, \mathbf{x}_{t-1}\}$ optimally.

Under adaptive learning, convergence of the sequence of beliefs to the true probability—to *rational expectations*—requires some additional conditions. For instance, the combination of the following two conditions is sufficient for the convergence of adaptive learning to rational expectations: (1) the true stochastic process for $\{\mathbf{x}_t, a_t\}$ is ergodic; and (2) δ_t goes to zero as t goes to infinity.

The most common specification—often referred to as *recursive least squares learning*—is the one where weights have the form $\delta_t = 1/t$. In this case, as agents accumulate more information as t grows, they eventually stop updating their beliefs. Another common specification consists of a constant weight parameter: $\delta_t = \delta$ for any period t . In this case, the weight put on an observation declines geometrically with the age of the observation. In general, this particular form of adaptive learning does not converge to rational expectations because it “overweights” the influence of the last observation when beliefs are close to the true probability.⁵ However, it is also this constant weighting of the last information that can make adaptive learning faster than Bayesian learning after a structural change in the true transition probability of the state variables.

(iv) *Reinforcement learning* (Erev and Roth 1998). The idea behind reinforcement learning (RL) is that agents update probabilities of playing certain strategies based on previous realized payoffs. In contrast to the learning algorithms that we have presented above, RL is not beliefs-based and it is not based on the explicit maximization of the firm's value.

RL assumes that strategies are “reinforced” by their past payoffs and that the propensity to choose a strategy depends on its *stock of reinforcement* (Erev and Roth 1998; Gureckis and Love 2013). Let $q_t(a, \mathbf{x}_t)$ represent the stock of reinforcement at period t of choice alternative a given that the current state is \mathbf{x}_t . At the initial period, $t = 0$, the firm has initial values for the stock of reinforcement of the different choice alternatives: $\{q_0(a, \mathbf{x}_0) : a = 0, 1, \dots, J\}$. These initial stocks are exogenously given. Then, at every period $t \geq 0$, the RL algorithm proceeds as follows. First, given current stocks $\{q_t(a, \mathbf{x}_t) : a = 0, 1, \dots, J\}$, the decision of the firm at period t is a random draw from a multinomial distribution with choice probabilities $\{p_t(a|\mathbf{x}_t) : a = 0, 1, \dots, J\}$ with

$$p_t(a|\mathbf{x}_t) = \frac{q_t(a, \mathbf{x}_t)}{\sum_{j=0}^J q_t(j, \mathbf{x}_t)} \quad (4)$$

Second, given its choice at period t , the firm observes its current profit $\Pi_t = \pi(a_t, \mathbf{x}_t)$. Then, the firm uses this information to update the reinforcement stock variables. This updating is based on the idea that action-state pairs that generated higher profits

⁵ Suppose that $b_{t-1}(\mathbf{x}'|a, \mathbf{x}) = p_x(\mathbf{x}'|a, \mathbf{x})$ at any point $(\mathbf{x}', a, \mathbf{x})$ such that at period $t-1$ the belief function is unbiased. Then, at period t the learning rule implies that $b_t(\mathbf{x}'|a, \mathbf{x}) - p_x(\mathbf{x}'|a, \mathbf{x}) = \delta [K([\mathbf{x}_t, a_{t-1}, \mathbf{x}_{t-1}] - [\mathbf{x}', a, \mathbf{x}]) - p_x(\mathbf{x}'|a, \mathbf{x})]$, which is different from zero at $[\mathbf{x}', a, \mathbf{x}] = [\mathbf{x}_t, a_{t-1}, \mathbf{x}_{t-1}]$ such that convergence is not achieved.

in the past should have larger reinforcement capital. For instance, for any action-state (a, \mathbf{x}) , the updating can be:

$$q_{t+1}(a, \mathbf{x}) = q_t(a, \mathbf{x}) + R(\Pi_t) K([a_t, \mathbf{x}_t] - [a, \mathbf{x}]) \quad (5)$$

where $R(\Pi_t)$ is an increasing function of current profits, and $K(\cdot)$ is a kernel function. Given a new realization of the state variables, \mathbf{x}_{t+1} , the updated stocks of reinforcement $\{q_{t+1}(a, \mathbf{x}_{t+1}) : a = 0, 1, \dots, J\}$ determine the choice probabilities at $t + 1$, and so on.

2.2 A Dynamic Game for an Oligopoly Industry

Consider an oligopoly industry with N potential entrants. We index firms by $i \in \mathcal{I} = \{1, 2, \dots, N\}$. As in the monopolistic case, the profit of a firm—say i —depends on: the firm's investment decision, a_{it} ; its capital stock, k_{it} ; and a vector of exogenous state variables, \mathbf{z}_t . Importantly, in this dynamic game a firm's profit also depends on the capital stocks of competing firms: the vector $\mathbf{k}_{-it} \equiv \{k_{jt} : j \in \mathcal{I}, j \neq i\}$. The per-period profit function is $\pi_i(a_{it}, \mathbf{x}_t)$ where now the vector of state variables is $\mathbf{x}_t = (k_{it}, \mathbf{k}_{-it}, \mathbf{z}_t)$.

Most of the empirical literature on dynamic games of oligopoly competition considers games of incomplete information under the *independent private values* (IPV) assumption. In this version of the model, the per-period profit function is $\pi_i(a_{it}, \mathbf{x}_t, v_{it})$ where v_{it} is a firm-specific variable that is the private information of firm i , and it is independently distributed across firms and over time, and is independent of \mathbf{x}_t .⁶ At period t , each firm i observes \mathbf{x}_t and its own private information, v_{it} , and chooses its current investment to maximize its expected value.

The sources of uncertainty in a firm's decision problem include: future values of the exogenous state variables and of the own capital stock, as in the single-firm problem; and current and future values of other firms' private information, of their investment decisions, and of their future capital stocks. Uncertainty about other firms' decisions appears not only because of the existence of private information. In models with multiple equilibria, firms can have uncertainty about the selection of the equilibrium: strategic uncertainty, as in Morris and Shin (2002). With bounded rationality, firms can have uncertainty about the level of rationality of the competitors: level- k rationality.

Let $p_{x,i}(\mathbf{x}_{t+1} | a_{it}, \mathbf{x}_t)$ be the *true* transition probability of the state variables conditional on the decision of player i . Under the IPV assumption, this transition probability can be decomposed as follows:

$$p_{x,i}(\mathbf{x}_{t+1} | a_{it}, \mathbf{x}_t) = p_z(\mathbf{z}_{t+1} | \mathbf{z}_t) p_{k_i}(k_{it+1} | a_{it}, k_{it}) \prod_{j \neq i} \sum_{a_{jt} \in A} p_{k_j}(k_{jt+1} | a_{jt}, k_{jt}) p_{a_j}(a_{jt} | \mathbf{x}_t) \quad (6)$$

⁶ Two important exceptions are Fershtman and Pakes (2012) and Asker et al. (2016) which allow for serially correlated private information. We discuss this more general model when we describe below the concept of *Experience-Based equilibrium*, and in section 5.6 where we describe Asker et al. (2016).

This decomposition illustrates the different sources of uncertainty that the firm faces, as to: (1) the future evolution of exogenous state variables, $p_z(\mathbf{z}_{t+1}|\mathbf{z}_t)$; (2) its own future capital stock, $p_{k_i}(k_{it+1}|a_{it}, k_{it})$; (3) the current investment decisions of its rivals, $p_{a_j}(a_{jt}|\mathbf{z}_t)$; and (4) the future capital stocks of the rivals, conditional on their current investments, $p_{k_j}(k_{jt+1}|a_{jt}, k_{jt})$. A firm has beliefs about each of these components. Therefore, we can write a firm's belief function at period t , $b_{it}(\mathbf{x}_{t+1}|a_{it}, \mathbf{x}_t)$, as⁷:

$$b_{it}(\mathbf{x}_{t+1}|a_{it}, \mathbf{x}_t) = b_{z, it}(\mathbf{z}_{t+1}|\mathbf{z}_t) b_{k, it}(k_{it+1}|a_{it}, k_{it}) \prod_{j \neq i} \sum_{a_{jt} \in A} b_{k_j, it}(k_{jt+1}|a_{jt}, k_{jt}) b_{a_j, it}(a_{jt}|\mathbf{x}_t) \quad (7)$$

Let $V_{it}(\mathbf{x}_t, v_{it})$ be the value of the firm at period t given current information and beliefs and under the condition that the firm maximizes expected profits forever in the future. This value function comes from the solution of a dynamic programming problem with the following Bellman equation:

$$V_{it}(\mathbf{x}_t, v_{it}) = \max_{a_{it} \in A} \left\{ \pi_i(a_{it}, \mathbf{x}_t, v_{it}) + \beta \int V_{i, t+1}(\mathbf{x}_{t+1}, v_{it+1}) b_{it}(\mathbf{x}_{t+1}|a_{it}, \mathbf{x}_t) g_i(v_{it+1}) dv_{it+1} d\mathbf{x}_{t+1} \right\} \quad (8)$$

The optimal decision rule, $\alpha_{it}^*(\mathbf{x}_t, v_{it})$, is the arg-max operator of the expression in brackets $\{.\}$.

In this type of model the strategic interactions between firms can include an important informational component. In particular, firms may have incentives to experiment in order to learn about the true transition probabilities, but they may also have incentives to learn by “free-riding” from their rivals’ experimentation.

To complete this dynamic game, the researcher needs to specify the assumptions on the evolution of firms’ beliefs. We describe here the different assumptions that have used in empirical applications of games.

(i) *Rational expectations* For every firm and every period, the belief function is equal to the actual probability distribution of the state variables, $b_{it}(\mathbf{x}_{t+1}|a_{it}, \mathbf{x}_t) = p_{x, it}(\mathbf{x}_{t+1}|a_{it}, \mathbf{x}_t)$. This assumption corresponds to the concept of Markov perfect equilibrium (MPE)—Maskin and Tirole (1988) and Ericson and Pakes (1995).

(ii) *Bayesian learning*⁸ For every firm i , let $\mathcal{P}_i \equiv \{\psi_{\ell, i}(\mathbf{x}'|a_i, \mathbf{x}) : \ell = 1, 2, \dots, L\}$ be a collection of L transition probabilities. The prior belief function for firm i

⁷ Similarly as in the single-firm model, the subindex t in the beliefs function represents the firm's information set \mathcal{H}_{it} and it indicates that this function can be updated over time as new information arrives.

⁸ In this representation of Bayesian updating in a dynamic game, we are implicitly assuming that at every period t firms observe the state vector \mathbf{x}_t but not necessarily the competitors' actions at the previous period, $\mathbf{a}_{-i, t-1}$. If they had this additional information, the Bayesian updating formula would be slightly different. Note that in many dynamic oligopoly models the endogenous state variable k_{it} is a deterministic function of a_{it-1} and k_{it-1} such that observing the sequence of realizations of k_{it} is equivalent to observing the sequence of choices a_{it} .

at period $t = 0$ is a mixture of the distributions in \mathcal{P}_i , where $\{\lambda_{\ell,i}^{(0)}\}$ are the mixing probabilities.

At any period $t \geq 1$, every firm proceeds as follows: First, firms observe the new state \mathbf{x}_t and use this information to update their respective beliefs using Bayes rule. The updated belief function for firm i is $b_{it}(\mathbf{x}'|a_i, \mathbf{x}) = \sum_{\ell=1}^L \lambda_{\ell,i}^{(t)}(\mathbf{x}', a_i, \mathbf{x}) \psi_{\ell,i}(\mathbf{x}'|a_i, \mathbf{x})$ for any value $(\mathbf{x}', a_i, \mathbf{x})$, where Bayesian updating of the mixing probabilities implies:

$$\lambda_{\ell,i}^{(t)}(\mathbf{x}', a_i, \mathbf{x}) = \frac{\psi_{\ell,i}(\mathbf{x}_t|a_{it-1}, \mathbf{x}_{t-1}) \lambda_{\ell,i}^{(t-1)}(\mathbf{x}', a_i, \mathbf{x})}{\sum_{\ell'=1}^L \psi_{\ell',i}(\mathbf{x}_t|a_{it-1}, \mathbf{x}_{t-1}) \lambda_{\ell',i}^{(t-1)}(\mathbf{x}', a_i, \mathbf{x})} \quad (9)$$

If the set of priors \mathcal{P}_i includes the true transition probability, then this Bayesian learning process implies that all the belief functions converge to the true transition probabilities (e.g., Feldman 1987, and on the speed of convergence, Vives 1993). By definition, these true probabilities—together with players' best response behavior—represent a MPE.

(iii) *Adaptive learning* The updating rule of firms' beliefs in the dynamic game under adaptive learning is similar to the formula in the single-firm model. For any player i and period t , and any of value $(\mathbf{x}', a_i, \mathbf{x})$,

$$b_{it}(\mathbf{x}'|a_i, \mathbf{x}) = (1 - \delta_{it}) b_{i,t-1}(\mathbf{x}'|a_i, \mathbf{x}) + \delta_{it} K_i([\mathbf{x}_t, a_{it-1}, \mathbf{x}_{t-1}] - [\mathbf{x}', a_i, \mathbf{x}]) \quad (10)$$

where parameters $\{\delta_{it} \in (0, 1) : t \geq 1\}$ and kernel function K_i have the same interpretation as in the single-firm model.

The convergence to a rational expectations equilibrium—a MPE—of beliefs under adaptive learning has been established by Evans and Honkapohja (1995) and Marcet and Sargent (1989a, b), among others, in different macro models, and by Pakes and McGuire (2001) in dynamic games of oligopoly competition. In these papers, *adaptive learning* can be interpreted also as a recursive algorithm that the researcher can use to compute a rational expectations equilibrium of the model.

(iv) *Reinforcement learning* The RL algorithm in the dynamic game is similar to its single-firm version. The model does not require the specification of beliefs about opponents' strategies. Instead, firms update the conditional choice probabilities based on realized current profits.

Let $\{q_{i0}(a_i, \mathbf{x}_0) : a_i = 0, 1, \dots, J\}$ be a vector with the stocks of reinforcement for firm i at the initial period $t = 0$. Then, at every period $t \geq 0$ firms proceed as follows. First, each firm chooses its own action a_{it} as a random draw from its conditional choice probability function $\{p_{it}(a_i|\mathbf{x}_t) : a_i = 0, 1, \dots, J\}$, with $p_{it}(a_i|\mathbf{x}_t) = q_{it}(a_i, \mathbf{x}_t) / \sum_{j=0}^J q_{it}(j, \mathbf{x}_t)$. Second, each firm observes its own profit, $\Pi_{it} = \pi_i(a_{it}, \mathbf{x}_t, v_{it})$, and uses this information to update the reinforcement stock variables according to the formula, $q_{i,t+1}(a_i, \mathbf{x}) = q_{it}(a_i, \mathbf{x}) + R(\Pi_{it}) K([\mathbf{x}_t, a_{it}] - [a_i, \mathbf{x}])$, where functions R and K have the same interpretation as in the single-firm model above.

RL algorithms are related to evolutionary games in economics (Samuelson 1998). In evolutionary games, there is a finite set of strategy functions that agents

can play, and a dynamic system that determines the proportion of agents that play each strategy at every period. In this dynamic system, the share of agents that play strategy j at period $t + 1$ depends on the share of this strategy at the previous period, and on the payoff that is received by strategy j and by the other strategies at period t . This evolution captures the idea that agents tend to imitate those strategies that have performed better in the past.

(v) *Experience-Based equilibrium* (Fershtman and Pakes 2012). So far, we have assumed that the private information v_{it} is independently and identically distributed. Now, suppose that there is persistent asymmetric information among firms: v_{it} is serially correlated and is observed only by firm i . Since private information is serially correlated, rivals' past actions and states contain relevant information about their current private information. Firms use rivals' past actions and states to predict rivals' current actions.

More specifically, in an Experience-Based equilibrium (EBE), each firm estimates the expected return of an action-state combination by averaging the returns that the firm obtained when that action-state was visited in the past. Given these estimates, each firm chooses the current action that maximizes its estimated expected value.

The concepts of EBE and RL share in common that firms use past payoffs in an action-state to construct ("reinforce") the probability of choosing that action under that state. However, the main difference between the two concepts is that under EBE players explicitly construct beliefs/expectations and maximize expected returns. This feature of EBE makes it more attractive than RL in the context of empirical applications using a structural model and with the purpose of counterfactual experiments. Nevertheless, the greater flexibility of RL makes it, potentially, a more robust approach to predict actual behavior in strategic settings.

(vi) *Fictitious play* (Brown 1951; Cheung and Friedman 1997). Firms form beliefs about rivals' conditional choice probabilities: probability distributions of rivals' actions conditional on current states, $p_{aj}(a_{jt}|\mathbf{x}_t)$. Fictitious play is a learning rule where each firm believes that rivals' actions are sampled from the empirical distribution of their past actions. Therefore, the belief function of firm i about the choice probability of firm j is, for any action-state (a_j, \mathbf{x}) :

$$b_{a_j, it}(a_j|\mathbf{x}) = \frac{\sum_{s=1}^t \omega_{(s,t)} 1\{[a_{jt-s}, \mathbf{x}_{t-s}] = [a_j, \mathbf{x}]\}}{\sum_{s=1}^t \omega_{(s,t)} 1\{\mathbf{x}_{t-s} = \mathbf{x}\}} \quad (11)$$

where $1\{\cdot\}$ is the function and $\{\omega_{(s,t)} : s \leq t\}$ are weights which are non-increasing in the lag index s .

In its original version (Brown 1951), the fictitious play model assumes that the weights $\omega_{(s,t)}$ are the same at every period s such that belief $b_{a_j, it}(a_j|\mathbf{x})$ is just the empirical frequency of action a_j conditional on state \mathbf{x} during periods 1 to t , i.e., $\sum_{s=1}^t 1\{[a_{jt-s}, \mathbf{x}_{t-s}] = [a_j, \mathbf{x}]\} / \sum_{s=1}^t 1\{\mathbf{x}_{t-s} = \mathbf{x}\}$. When all the weights $\omega_{(s,t)}$ are zero except for $s = 1$, we have that $b_{a_j, it}(a_j|\mathbf{x}) = 1\{[a_{jt-1}, \mathbf{x}_{t-1}] = [a_j, \mathbf{x}]\} / 1\{\mathbf{x}_{t-1} = \mathbf{x}\}$, and this is known in the literature as *Cournot learning*. Cheung and Friedman (1997) introduced the more general *weighted fictitious play* as described in equation (11).

(vii) *Rationalizability* (Bernheim 1984; Pearce 1984). The concept of rationalizability imposes two simple restrictions on firms' beliefs and behavior. First, every firm is rational in the sense that it maximizes its own expected profit given beliefs. And second, this rationality is common knowledge—every firm knows that all the firms know that it knows ... that all the firms are rational.

The set of outcomes of the game that satisfy these conditions (the set of rationalizable outcomes) includes all of the Nash equilibria of the game. Every Nash equilibrium is a rationalizable outcome; however, the inverse is not true. In models of market competition with multiple Nash equilibria, the set of rationalizable outcomes can be much larger than the set of Nash equilibria.

(viii) *Cognitive Hierarchy and Level-k Rationality* These models assume that players have different levels of strategic sophistication. Every firm (player) maximizes its subjective expected profit given its beliefs. Firms are heterogeneous in their beliefs and there is a finite number of belief types. Beliefs for each type are determined by a hierarchical structure.

The Cognitive Hierarchy model (Camerer et al. 2004) and the Level-k model (Costa-Gomes and Crawford 2006; Crawford and Iriberri 2007) consider a similar hierarchical structure of beliefs, but there is an important difference. In both models, Level-0 firms believe that strategic interactions are negligible and therefore they behave as in a single-agent model: as if they were monopolists. Level-1 firms believe that the rest of the firms are level-0, and they behave by best responding to these beliefs.

The difference between the Cognitive Hierarchy and the Level-k models is in the beliefs of players at levels k greater than one. In the Level-k model, a level- k player believes that all of the other players are level $(k - 1)$. In the Cognitive Hierarchy model, a level- k player is not certain about the types of the other players but believes that their types come from a probability distribution over levels 0 to $(k - 1)$.

This subtle difference between the two models can have important implications for their predictions. Under very general conditions on the payoff function, the outcome of the game in these models is unique (though different between the two models). This uniqueness is a convenient property for empirical analysis and for counterfactuals.⁹

⁹ In the literature, we find two different approaches for the specification of the behavior of Level-0 players. In empirical applications in industrial organization (Goldfarb and Xiao 2011; Hortaçsu et al. 2017), the assumption is that a level-0 firm behaves as a monopolist. However, a good number of experimental papers assume that Level-0 players behave randomly (Camerer et al. 2004; Costa-Gomes and Crawford 2006). Crawford and Iriberri (2007) consider both types of Level-0 behavior. As shown in that paper, the specification of the behavior of Level-0 players can have important implications on the model predictions.

3 Data and Identification

In this section, we describe the types of data and identification strategies that are used in the estimation of the structural models that were described in Sect. 2. We can distinguish three types of data for the estimation of these structural models: choice data, payoffs data, and beliefs data.

By choice data we mean information of firms' actual decisions, as well as some of the state variables of the problem. For instance, in a dynamic game of investment, the researcher observes firms' investments, their capital stocks, and variables that affect demand and costs. This is the type of data that we find in most empirical applications in industrial organization. Using these data, the researcher can identify the probability distribution of firms' decisions conditional on (observable) state variables: a probabilistic representation of firms' strategies or decision rules.

Under relatively weak conditions, there is a one-to-one relationship between a firm's choice probability function and the firm's expected profit (in a static model) or expected value function (in a dynamic model): The expected payoff function can be identified using choice data.

A firm's expected payoff function depends on its (primitive) payoff function and on its beliefs about the probability distribution of uncertain events. The separate identification of these two functions from observed choices (from the expected payoff) is a key identification problem in this literature. In general, the expected payoff is consistent with many alternative specifications of the payoff function and the belief function. Without some restrictions on these functions there is not point-identification with the use of only choice data.

We first present identification results using choice data only. Then, we explain how supplementing choice data with data on firms' costs or with beliefs data allows the researcher to relax some identification assumptions.

3.1 Identification of Beliefs Using Choice Data

Firms' behavior reveals information not only on managerial preferences but also on their beliefs. We can label this approach as *revealed beliefs* or *revealed preference and beliefs*. This approach is the most commonly used in the structural empirical studies that we survey in this paper.

Aradillas-Lopez and Tamer (2008) study the identification of preferences and beliefs when we relax the assumption of Nash equilibrium and replace it with *rationalizability*. They consider relatively simple but commonly used structural games, such as two-player static discrete choice games, and first-price auction games under independent private values. The specification of players' beliefs is nonparametric and the only restriction on beliefs comes from the assumption of common-knowledge rationality: (i) Every player is rational in the sense that she maximizes her own expected payoff given her beliefs; and (ii) Every player knows that the other players know that she knows ... that all the players are rational. They show that, without further restrictions, choice data cannot point identify preferences and beliefs. However,

under relatively mild conditions, there is set-identification of some parameters in the payoff function.

3.1.1 Imposing Equilibrium Beliefs in a Subset of the Data

A simple and intuitive approach to achieve identification of payoff and belief functions consists of assuming that the data can be separated into two subsets (e.g., two time periods, or two sets of geographic markets)—say R and B (for *rational* and *biased* beliefs, respectively)—with the following properties: (i) For subset R of data beliefs are in equilibrium and correspond to the actual distribution of the variables, i.e., rational expectations; and (ii) The payoff functions in the two subsets of data are the same.

In the case of dynamic models, if subsample R corresponds to time periods before subsample B , then there is the additional assumption that firms do not anticipate their change in beliefs from R to B . This additional assumption is not necessary if subsample R occurs after B , or if they correspond to two sets of independent markets.

Under condition (i) of rational expectations in subset R , we can identify the payoff function from this subset of choice data. Then, given condition (ii) that the payoff function is invariant between the two subsets, we can use the B dataset to identify firm's beliefs without having to impose the restriction of rational expectations for this subset of data. This approach has been used recently by Doraszelski et al. (2018) and Huang et al. (2018).

We describe these applications in Sect. 4 below. Of course, the suitability of this approach depends on how plausible conditions (i) and (ii) are in the context of a specific application.

3.1.2 Static Auction Games

Gillen (2010) studies the nonparametric identification of the distributions of bidders' values and their levels of strategic sophistication in the Level- k first-price auction model proposed by Crawford and Iriberri (2007). The model assumes the same distribution of values for the different level- k types of bidders. The structural parameters of the model consist of the distribution of values— F_V —and the distribution of k -types: $p = (p_1, p_2, \dots, p_K)$. The data consist of the empirical distribution of bids from a sample of homogeneous auctions with the same number of bidders: N .

For any level- k , the model implies a one-to-one relationship between a bidder's value and her bid. Therefore, if all the bidders have the same k -type and the researcher knows this type, then the identification of the distribution of values follows from the standard results in the econometrics of IPV auctions (Guerre et al. 2000). This identification argument extends to the case where bidders have different k -types but the researcher knows the frequency of each type, as represented by the vector p .

Suppose now that the researcher observes two groups of auctions from the same population of bidders, i.e., with the same distributions F_V and p . The difference between these two groups of auctions is in the number of bidders that participate in

each auction: N_1 bidders in the auctions of group 1, and $N_2 \neq N_1$ bidders in group 2. These two samples imply over-identifying restrictions for the distribution F_V , for a given value of p . These over-identifying restrictions can be used to test a proposed value for the distribution of k-types: p . Gillen (2010) shows that, in fact, these over-identifying restrictions can identify the distribution of k-types.

An (2017) applies recent developments in the econometric literature on nonlinear measurement error—Hu (2008) and Hu and Schennach (2008)—to identify the distributions of values and level-k in the Crawford-Iriberri or Level-k auction model. He shows that if the dataset has a panel structure such that each bidder is observed in (at least) three independent auctions, then the distributions of values and of the level of strategic sophistication are nonparametrically point-identified. This is a very useful and powerful identification result for applied researchers.

3.1.3 Static Discrete Choice Games

Goldfarb and Xiao (2011) estimate a Cognitive Hierarchy game of market entry. Note that this model is more general than the Level-k model, and therefore its identification is more challenging. The authors estimate the distribution of firms' types and the profit function using firm-level entry.

Their identification of the structural parameters is based on the sample variation across local markets in firms' entry probabilities and—importantly—in the variance of entry probabilities of firms within the same market. Note that level-0 managers (who are acting as monopolists) have a high probability of entering a given market; level-1 managers, who are best-responding to level-0 managers, have a low probability of entering; and level-2 and higher managers, who are optimizing against a mixture of level-0 and level-1 managers, have an intermediate probability of entering.

While an average entry probability in a market may be attributed either to a combination of high competitive effect and low cognitive types, or to a low competitive effect and high cognitive types, observing that this market has a low variance of entry probabilities suggests high cognitive types, and thus it favors the second explanation.

Xie (2018) studies the identification of static discrete choice games of incomplete information when we do not impose the restriction of Bayesian Nash Equilibrium. The beliefs function is a probability distribution over the set of possible actions of the other players. Xie allows the belief function to be an unrestricted probability distribution, and studies the joint identification of utility and beliefs functions that is based on observed players' behavior.

Suppose that the game is such that a player's number of possible actions is greater than the number of actions of her opponents. Under this condition, and without further restrictions, Xie shows that it is possible to control for this player's beliefs and to set-identify the player's utility function. Furthermore, if the population of markets under study contains variation in players' action spaces across markets, then it is possible to point-identify players' preferences and to test the null hypothesis of equilibrium beliefs.

There are multiple empirical applications of games where players' action spaces vary both across players in the same market and over different markets. In empirical

industrial organization, examples of this type of setting typically appear when firms are multi-product or multi-store and they compete in actions at the product or store level—e.g., prices, advertising—such that the cardinality of the action space depends on the firm's number of products/stores. Firms with different products/stores have different choice sets and the number of products or stores that firms have may vary across local markets.

Xie takes into account that this variation in the number of actions over local markets can be endogenous. In particular, it can be correlated with unobservables for the researcher that affect players' utility in the game. To deal with this endogeneity problem, he proposes a control function approach.

3.1.4 Dynamic Games

In most empirical games, the payoff function includes parameters that capture the effect that competitors' choices have on a player's payoff: strategic interaction parameters. The identification of these parameters requires an exclusion restriction—even under the assumption of rational expectations (Bajari et al. 2010). Therefore, it is relevant to consider whether this exclusion restriction has the power to identify players' beliefs.

Aguirregabiria and Magesan (2019) study identification of dynamic discrete games when firms' beliefs about competitors' actions may be out of equilibrium. Suppose that each player has an observable state variable that enters in her own payoff function, but it does not have a direct effect on the payoff of the other players. Dynamic games of oligopoly competition provide multiple examples of this type of exclusion restriction. For instance, in a standard model of market entry-exit, the incumbent status of a firm in the market affects the firm's profit by determining whether it must pay an entry cost to be active in the market, but it does not have any direct effect on the profits of the other firms. The profits of the other firms are affected only indirectly through this firm's current choice to be active in the market.

These authors show that, in dynamic games, this restriction can be used to identify biases in players' beliefs about other players' behavior and to test the null hypothesis of unbiased beliefs. Consider a market entry game with two potential entrants: firms 1 and 2. Let s_1 be the *special* state variable that enters into the profit function of firm 1 but not into that of firm 2. Under this condition, the probability that firm 2 enters the market depends on the state variable s_1 only because firm 2 believes that this variable affects the behavior of firm 1. The dependence of the choice probability of firm 2 with respect to s_1 reveals information about firm 2's beliefs.

Aguirregabiria and Magesan (2019) show that this information identifies non-parametrically a player's belief as a function of the special state variable up to an intercept and a scale constant. Using this identification result, they construct a non-parametric test for the null hypothesis of equilibrium beliefs.

In the context of single-agent dynamic discrete choice structural models, An et al. (2018) provide sufficient conditions for the identification of agents' subjective beliefs on the Markov transition probabilities of the state variables with the use

of choice data. Their specification of the probabilistic beliefs is nonparametric and can contain any type of bias with respect to the true transition probabilities. The key condition for their identification result is that there is a *special state variable* that enters linearly in the agent's payoff function and for which the agent knows its actual transition probability function.

3.2 Combining Choice Data with Firms' Cost Data

Most industrial organization models assume that firms maximize profits, revenue minus costs. Sometimes the researcher has direct measurements of firms' costs, or some components of these costs. This information—combined with choice data—can be very helpful for the identification of firms' beliefs. Some recent papers have exploited this identification approach.

Hortaçsu and Puller (2008) study firms bidding behavior in the Texas electricity spot market. Their dataset includes direct measurements of firms' marginal costs. The authors use this information to test the hypothesis of equilibrium bidding behavior. Hortaçsu et al. (2017) use the same dataset to estimate a structural Cognitive Hierarchy model of bidding.

With the use of the marginal cost data the authors construct the (Bayesian Nash equilibrium) best-response supply of each firm and compare it to the firm's actual supply using bidding data. They find substantial discrepancies between observed and equilibrium supply curves. Observed supply curves are above and steeper than the best-response curves. Furthermore, there is large firm heterogeneity in the discrepancy between observed and best-response bidding. There is a group of firms with supply curves that are almost completely vertical.

As an identifying restriction, Hortaçsu et al. (2017) assume that these bidders with vertical bidding curves correspond to the Level-0 type of the Cognitive Hierarchy model. This assumption provides an anchor that helps in the identification of the probability distribution of the other k-types.

3.3 Beliefs Data

Beliefs data consist of information on agents' subjective perception of the probability distribution of uncertain events. Though there is a substantial literature on the elicitation of individuals' beliefs either in the laboratory (Levitt and List 2007) or in the field (DellaVigna 2009; Manski 2018), this approach is much less common when using field data of firms.

There are a large number of business surveys that ask managers about their expectations as to future events: such as the evolution of demand and prices, or the growth and survival of their own business. Unfortunately, most of these surveys ask about point predictions and not about probabilistic beliefs.¹⁰ Furthermore, many times, the

¹⁰ See the critique in Manski (2018) of the traditional practice of asking to make point predictions of future events, instead of asking for probabilistic beliefs. See also Potter et al. (2017) on the advantages of probabilistic survey questions.

point predictions do not have a quantitative interpretation. For instance, Bachmann et al. (2013) study firms' uncertainty using business survey data from the German IFO Business Climate Survey (IFO-BCS) and from the Philadelphia Fed Business Outlook Survey (BOS). In these surveys, the questions on expectations have the following form: in the IFO-BCS, "*Expectations for the next 3 months: Your domestic production activities with respect to product X will: 1. increase, 2. roughly stay the same, 3. decrease*"; and in the BOS, "*Your shipments 6 months from now versus current month will: 1. decrease, 2. no change, 3. increase*". It is difficult to incorporate these non-probabilistic (and qualitative) beliefs data in our models.¹¹

Nevertheless, there are a few business surveys that include questions about probabilistic beliefs. An example is the Italian Survey of Investment in Manufacturing. Guiso and Parigi (1999) use these data to study how the entrepreneurs' subjective probability distribution of future demand for firms' products affects capital investment. Consistent with theory, the authors find that the effect of subjective uncertainty on investment is stronger for firms with more market power and with more industry-specific capital.

Firms may have incentives to conceal their actual beliefs and do not report truthful expectations in business surveys. This is, perhaps, not a main concern for beliefs on exogenous events, such as future demand or input prices. However, it can be more problematic for beliefs about their own future behavior, or their competitors' behavior. Experimental economists have been concerned about this problem for a long time (Savage 1971), and they have developed incentive mechanisms for honest revelation of expectations with respect to observable events.

A mechanism that is commonly used in experimental economics to deal with this problem is *Proper Scoring Rules* (PSR) (Nyarko and Schotter 2002; Schotter and Trevino 2014). PSR provide ex-post rewards to agents to reveal their true probabilistic beliefs. If agents are risk neutral and maximize expected payoffs, then PSRs can elicit true beliefs. However, the implementation of PSRs requires the researcher to verify which events do and do not occur. While this is feasible in laboratory experiments, it becomes more complicated and costly in surveys—such that PSR has been rarely applied in field/survey data. A nice exception is the work by Armantier et al. (2015) on individuals' inflation expectations. In the case of eliciting firms' beliefs, an additional issue is the magnitude of the reward that is necessary to generate the revelation of honest beliefs. Revealing its true beliefs may have a substantial cost for a firm in terms of lower profits. In some cases, compensating firms for this expected reduction in profits can be extremely costly for the researcher. As far as we know, there is not any empirical application using PSR to elicit firms' beliefs.

¹¹ Other example of this type of data on firms' beliefs is the survey on entrepreneurs from the French statistical office (INSEE). For new small startups, the survey ask entrepreneurs the following two questions on their beliefs: Question 1: "*What is your view for the next 6–12 months?: 1. the firm will develop, 2. the firm will keep its current balance, 3. I will have to struggle, 4. I will have to shut down the firm, 5. I will sell it, 6. I do not know.*"; Question 2: "*Do you plan to hire in the next 12 months?: 1. yes, 2. no, 3. I do not know*". Landier and Thesmar (2009) use these data to study the relationship between entrepreneurs' optimism (over-confidence) and the use of short-term debt finance.

4 Empirical Applications

In this section, we review recent applications in empirical industrial organization that allow for firms' non-equilibrium or biased beliefs. Though our survey focuses on applications that use structural models, we start by reviewing an important empirical literature on firms' overconfidence that use non-structural models (Sect. 4.1). Then, we describe applications of static games with non-equilibrium beliefs: both discrete choice games of incomplete information (Sect. 4.2) and auction games (Sect. 4.3). The last three subsections deal with dynamic models. In Sect. 4.4, we review applications of dynamic games of oligopoly competition with non-equilibrium beliefs but without an explicit model of learning. Sections 4.5 and 4.6 deal with dynamic models with firms' Bayesian and non-Bayesian learning, respectively.

4.1 Firms' Overconfidence: Non-structural Evidence

Empirical studies in corporate finance have documented that firm managers tend to overestimate their ability (Cooper et al. 1998). Some studies look at the implications of this overconfidence. Malmendier and Tate (2005, 2008) construct a measure of overconfidence that uses the personal investments of CEOs. They use field data to show that CEO overconfidence has a significant impact on important strategic decisions such as investments, mergers, and acquisitions.

Galasso and Simcoe (2011) use the Malmendier-Tate measure of overconfidence to study the impact of CEO overconfidence on innovation decisions. They find that the arrival of an overconfident CEO is correlated with an increase in corporate innovation: as measured by the number of patents, citations received, or expenditure in R&D. Their estimates show that the impact of CEO overconfidence on innovation is stronger under intense product market competition.

Landier and Thesmar (2009) study the role that beliefs play in small entrepreneurs' financial decisions. They use survey data that were collected by Statistics France on a nationally representative sample of French entrepreneurs with information on the entrepreneurs' self-reported beliefs about the future growth of their own business. The authors find substantial heterogeneity in entrepreneurial beliefs, and a positive correlation between optimism (overconfidence) and the use of short-term debt finance rather than the less-risky option of long-term debt finance.

Astebro et al. (2014) distinguish two different types of overconfidence in firms' investment decisions: upward bias in the expected return (*overestimation*); and downward bias in the expected risk (*overprecision*). While the first type of overconfidence can generate excess market entry and investment, the second type can have a negative impact on market entry (and on other investment decisions with limited liability) because it reduces the perceived option value of market entry decisions.

4.2 Static Discrete Choice Games with Non-equilibrium Beliefs

4.2.1 Static Discrete Choice Games with Cognitive Hierarchy Structure

Based on the seminal work by Camerer et al. (2004) on the Cognitive Hierarchy (CH) model and by Costa-Gomes and Crawford (2006) on the Level- k , these models have received substantial attention in empirical applications of games in experimental economics that use laboratory experiments. Several papers have adapted these models to study market competition with the use of field data.

The papers by Goldfarb and Yang (2009) and Goldfarb and Xiao (2011) are among the first empirical studies that structurally estimate non-equilibrium models of strategic thinking. These papers apply and adapt the asymmetric information CH model to games of market entry with the use of field data. Because higher level- k types can make better predictions about rivals' play, Goldfarb and Yang (2009) and Goldfarb and Xiao (2011) interpret the estimated cognitive level of a manager as her strategic ability.

Goldfarb and Yang (2009) estimate a game of technology adoption that is played by internet service providers (ISPs) under the assumptions of the CH model. In 1997, the manager of an ISP that operated in a local market (metropolitan area) had to choose among four possible alternatives to offer their customers internet access: (i) offer only the old technology of 33K modems; (ii) offer the new 56K technology that used Rockwell's modems; (iii) offer the new 56K technology that used US Robotics's modems; or (iv) offer the two formats of modems for the new technology.

The authors adapt the CH model to describe the heterogeneity in the ability or sophistication of ISP managers in this strategic setting. A manager with level-0 sophistication ignores rivals' anticipated decisions and behaves as a monopolist. For k greater or equal than one, a level- k manager assumes that all his competitors are a mixture of lower levels, with those levels drawn from a truncated Poisson distribution.

The estimated results show substantial heterogeneity in the level of sophistication among managers. The estimated distribution of level- k implies a mean value of 2.67. This is at the high end of the range of values that are found in Camerer et al. (2004) for their lab setting. This is reasonable as we expect managers to think and behave more strategically than do undergraduate students in a lab.

Though the CH model fits the data no better than does a model under the restriction of Bayesian Nash equilibrium (BNE), it has interesting implications that are different from those under the BNE. Goldfarb and Yang show that heterogeneity in strategic thinking contributed to the slowing of the diffusion of the 56K technology relative to the diffusion under the counterfactual BNE. They also show that more strategic firms (as measured in 1997) were more likely to survive the dot-com crash in year 2000.

Goldfarb and Xiao (2011) propose and estimate a similar CH model to investigate entry decisions into local US telecommunication markets following the deregulatory Telecommunications Act of 1996, which encouraged greater local competition. First, the authors present reduced-form evidence that shows that—holding other market characteristics constant—more-experienced and better-educated managers

have a lower propensity to enter into very competitive markets. This suggests that better-educated managers are better at predicting competitors' behavior.

Motivated by this pattern in the data, the authors estimate a CH model where the level of strategic sophistication depends on manager characteristics. The authors estimate the firms' types and the profit function with the use of firm-level entry data in 1998 and 2002. The estimation results show that manager characteristics are key determinants of the level of strategic sophistication. In particular, better-educated managers have higher strategic ability. The estimated level of cognitive ability is higher in 2002 than right after the first wave of entry in 1998, which suggests that low-ability managers were more likely to fail during the industry shake-out. This suggests that the behavioral model may be most relevant during the early stages of the life cycle of an industry.

Brown et al. (2012, 2013) estimate a game for the decision of movie distributors as to whether to pre-release a movie to critics. Distributors face a choice between "cold opening" a movie or pre-releasing it to critics in the hope that favorable reviews will increase revenue. This decision can be seen as an asymmetric-information signaling game with verifiable signals. A well-established theoretical result in this class of models (Grossman 1981; Milgrom 1981) is that higher-quality firms have incentives to reveal their true quality to distinguish themselves from lower-quality competitors. Accordingly, in the Bayesian equilibrium of this game, cold-opening should not be profitable because moviegoers will infer low quality for cold-opened movies. Therefore, there should be no cold opening, except possibly by the very worst movie type.

With the use of a sample of 1303 movies, Brown et al. (2013) show that—in contradiction to the Bayesian equilibrium prediction—movies that are not shown to critics ("cold-opened" movies) have between 10% and 30% higher revenue than comparable movies that were critically reviewed before released. This evidence suggests that moviegoers had unrealistically high expectations for cold-opened movies, which is hard to explain using a model with equilibrium beliefs. Furthermore, movie distributors do not appear to take advantage of moviegoers' lack of sophistication, since only 7% of movies were opened cold despite the expected-profit advantage.

To explain this behavior, the authors propose and estimate a cognitive hierarchy model. The estimated model shows that moviegoers have a low level of strategic sophistication. However, movie studios appear to be best responding to consumers that are playing a fully Bayesian-Nash strategy. These results suggest that movie studios would benefit by cold opening more movies.

4.2.2 Other Static Discrete Choice Games with Biased Beliefs

Xie (2018) studies competition in opening hours between two Western-style fast-food restaurants in China. This industry has experienced important structural changes in China in recent years. As a result, firms face substantial strategic uncertainty about the behavior of their competitors. Non-equilibrium beliefs is a plausible hypothesis in this context. The key parameters of interest in this application are the parameters that capture the strategic substitutability/complementarity between

firms' decisions to remain open for 24 h. These parameters can be particularly sensitive to wrongly imposing the assumption of equilibrium beliefs.

Xie studies competition between KFC and McDonalds (the leaders of this industry in China) in their decisions to remain open for 24 h. The author estimates a model of market entry with incomplete information where firms' probabilistic beliefs about the behavior of the competitor are completely unrestricted functions of payoff-relevant state variables. He follows the identification approach that we have described in Sect. 3.1.2, that exploits the asymmetry in firms' choice sets. The number of restaurants of a firm/chain in a local market varies both between firms within a market and across markets such that the cardinality of the action space—how many restaurants in a local market a firm can remain open 24 h—has also the type of variation that is key for Xie's identification strategy.

Xie finds that firms' decisions to remain open 24 h are strategic complements. The estimated results also provide significant evidence that KFC has biased beliefs about the behavior of McDonalds. Accounting for these biased beliefs has important implications for the estimates of the strategic interaction parameters. In particular, it increases significantly the estimated competition effects between the firms in some markets.

4.3 Auctions with Non-equilibrium Beliefs

Hortaçsu and Puller (2008) analyze the bidding behavior of firms in the Texas electricity spot market, where suppliers trade with each other to meet prior contractual obligations and to balance the market. In this market, electricity generating firms submit hourly supply schedules to sell power. Winning sellers earn the price at which aggregate supply bids equal demand.

A key feature of Hortaçsu and Puller's empirical approach is that their dataset contains detailed information not only on firms' bids but also on their marginal costs. Using these data on marginal costs, the authors construct the equilibrium bids of the game (under specific assumptions about the information structure of the game) and compare them to the actual observed bids.

The authors find statistically and economically significant deviations between equilibrium and actual bids. More specifically, while large firms best-respond to other firms' behavior, small firms submit bid functions that are excessively steep. Small firms insufficiently adjust their production quantities to market circumstances, and do not supply much power even when it is ex post profitable to do so.

Based on their interviews with traders in this market, Hortaçsu and Puller argue that this finding is best explained by the relatively low strategic ability in the bidding departments of small firms. There are economies of scale in setting up and maintaining a successful bidding department such that small operators cannot afford the fixed costs of establishing a sophisticated bidding unit. The authors also provide evidence of learning by small firms, with a 10% performance improvement per year.

Very interestingly, this suboptimal behavior by small firms leads to significant efficiency losses in the market. In fact, the inefficiency that is generated by the mistakes of smaller firms is larger than the inefficiency that is generated by the market

power of large firms. This finding has important policy implications for the design of this industry. Market performance could be improved by merging the bidding operations of small firms or by reducing the strategic complexity of this market mechanism.

Motivated by the results that were described in the previous paragraph, Hortaçsu et al. (2017) adapt the Cognitive Hierarchy model to the Texas electricity market. They assume that a level-zero firm simply submits a perfectly inelastic bid function at its contract position. Following Goldfarb and Yang (2009) and Goldfarb and Xiao (2011), they also assume that a higher-level firm best-responds to a truncated Poisson distribution over lower-level types. Because higher-level firms believe that their rivals are higher-level as well, the higher-level firms face a more elastic residual demand curve, which leads them to bid more competitively.

Hortaçsu et al allow the level of sophistication in bidding to depend on the manager's education (similarly as Goldfarb and Xiao 2011) and on firm size (motivated by the results in Hortaçsu and Puller 2008). The information on bidders' marginal costs in this dataset plays a key role in the identification of beliefs.

The authors find very substantial heterogeneity in the level of strategic sophistication across the firms in the Texas electricity market. The strongest determinant of a firm's level of sophistication is its size, though managers' education plays also a statistically significant but small role. Counterfactual experiments show that exogenously increasing the sophistication of low-type firms to the level of median-type firms increases market efficiency by between 9% and 16%.

4.4 Dynamic Games with Non-equilibrium Beliefs

4.4.1 Market Entry

Aguirregabiria and Magesan (2019) study competition in store location between McDonalds (MD) and Burger King (BK) during the early years of the fast-food restaurant industry in the UK. McDonalds opened its first restaurant in the UK in 1974, but it was not until 1981 that it opened outlets outside the London area. Burger King started operating in the UK in 1988. The sample period in this study is 1991-1995.

Estimates of reduced-form models for the decision of opening a new store show that the number of own stores has a strong negative effect on the probability that BK opens a new store but the effect of the competitor's number of stores is economically negligible. In contrast, for MD, the decision of opening a new store is sensitive to BK's existing number of stores in the market. These findings are very robust to different specifications and control variables, and is analogous to evidence in Toivanen and Waterson (2005). This behavior cannot be rationalized by standard equilibrium models of market entry where firms compete and sell substitute products and have equilibrium beliefs about the behavior of competitors.

Aguirregabiria and Magesan consider three possible explanations for this finding: (i) unobserved market heterogeneity for the econometrician that is known to the firms; (ii) positive spillover effects of the competitor's presence; and (iii) firms'

biased beliefs about the competitor's behavior. They propose and estimate a dynamic game of competition in the number of stores that incorporates these three features.

The identification of the structural model in Aguirregabiria and Magesan (2019) follows the approach that was described in Sect. 3.1.3 above. More specifically, they assume that there are two state variables that affect the profit of a firm in a local market (UK district) at period t but not the profit of the competitor at the same period and local market: (i) the number of own stores at period $t - 1$, which affects the firm's cost of setting up a new store in this local market but not its rival's setup cost; and (ii) the geographic distance of this local market to the nearest store that the firm has in other local markets, which affects the firm's economies of density but not its competitor.

Based on these exclusion restrictions, the test of equilibrium beliefs clearly rejects the null hypothesis that BK has unbiased beliefs about MD entry decisions, but it cannot reject the hypothesis that MD has unbiased beliefs. The bias in BK beliefs are mostly in the direction of underestimating the true probability that MD will open a second restaurant store in those markets where it already has one store. Relaxing the assumption of equilibrium beliefs provides more plausible estimates of the structural parameters that capture the effect of competition.

Finally, counterfactual experiments show that having unbiased beliefs would increase BK profits by almost 3% in year 1991. This effect declines over time and it becomes practically zero in 1995.

4.4.2 Dynamic Pricing with Managerial Attention Costs

Ellison et al. (2018) study dynamic price competition in an online platform (Price-watch) where small retailers sell computer components. The authors provide structural estimates of managerial frictions. The platform ranks firms according to price which, along with rapidly changing market conditions, gives firms an incentive to change price frequently. Through reduced-form analysis, the authors establish that there is substantial price inertia and that this inertia is at least partially due to costs of managerial attention.

Based on the evidence, the authors propose and estimate a dynamic game of competitive price adjustment that incorporates two types of managerial costs: (1) a cost of monitoring market conditions: acquiring information on state variables, including the current rank and the distribution of competitors' prices; and (2) a menu cost of physically entering a new price. The main focus of this paper is in the separate identification of the magnitudes of these two costs and of their respective contributions to price dynamics and firms' profits.

A key challenge comes from the fact that the researcher observes price change, but not firms' decisions to monitor information. This is a common issue in empirical applications that introduce inattention as a friction.

To deal with this issue, the authors propose an interesting identification strategy. The strategy is based on a key assumption or exclusion restriction that is related to variables that are predictors for the rank position of the price of a firm: the time duration since the firm's last change in price (SinceChange); and the ratio between

the firm's current sales and the sales at the period of the last price change (QuantityBump). The assumption is that, once the firm decides to monitor and learns about its current rank, these state variables become irrelevant for a firm's optimal price decision.

This assumption has the power to identify the existence of inattention, i.e., existence of monitoring costs. If firms were continuously attentive about their rank, then the probability of price change should not depend on the variables SinceChange and QuantityBump, once we have conditioned on Rank. The dependence of the probability of price change on SinceChange and QuantityBump controlling for Rank is consistent with firms' inattention. Under additional parametric assumptions, this dependence also identifies the structural parameter that measures monitoring cost.

For two groups of firms (out of three groups that are partitioned based on machine-learning techniques) the estimate of monitoring cost is sizable—approximately \$65—while the estimate of the menu cost is negligible: The 95% confidence interval is [\$0, \$4]. Counterfactual experiments show that the composition of the managerial cost—the relative weight of the monitoring cost and the menu cost components—matters for pricing decisions and firms' profits. Reversing these cost components in size (while keeping the total price adjustment cost and the behavior of competitors constant) results in a substantial increase in the firm's profits.

4.5 Dynamic Models with Firms' Bayesian Learning

4.5.1 Pricing with Learning About Competitors' Quality

Ching (2010) incorporates consumer and firm learning in a dynamic oligopoly structural model of competition between a brand-name drug and its generic competitors. In the model, firms are forward-looking and set price in each period. Both firms and consumers are uncertain about the true quality of generic drugs. Firms can use price to control the rate of learning. In particular, on the assumption that consumers are risk-averse, generic firms may have an incentive to price low to encourage more consumers to try their products and resolve the uncertainty. Ching assumes that both firms and consumers use Bayesian updating to learn about the true quality of generic drugs. Each period the brand-name firm acts as the Stackelberg leader, and the generic firms are followers.

4.5.2 Pricing with Learning About Consumer Valuation of Product Attributes

Huang et al. (2015) use a structural learning model to study how car dealers set prices for used cars over time. They argue that the pricing of used cars is more complicated than for new cars because their value depends on mileage, year, and maintenance and it is not clear how consumers perceive and trade-off these dimensions. Car dealers face this demand uncertainty but they can learn over time.

Every period a dealer sets the price for a used car, and then consumers decide whether to buy it. If consumers choose not to buy the used car, this gives the dealer a signal about its unobserved demand component, and she then updates her belief

accordingly. If a dealer is forward-looking, she has an incentive to price the used car high early on, because not selling the car gives her an opportunity to obtain better information about the demand for the car.

With the use of a panel dataset of used-car sales from the dealer CarMax, Huang et al. (2015) find that their structural model can explain demand and pricing patterns well.

4.5.3 Demand Uncertainty and Learning from Competitors' Capacity Choice

Gardete (2016) studies competition in capacity and production between manufacturers in the dynamic random access memory (DRAM) industry. Firms in this industry face significant demand uncertainty before production and capacity decisions can be implemented. Since capacity decisions become common knowledge before production takes place, this information can be used by firms to learn about competitors' signals on uncertain demand.

Gardete proposes and estimates a structural model of competition that has two stages at every period t . In a first stage (capacity investment stage), firms receive private signals, create heterogeneous beliefs about demand, and make investment decisions for their capacities. In a second stage, capacity choices become public, firms update their beliefs about demand using Bayes rule, and then compete in quantities à la Cournot. The paper focuses on the role of market information and the consequences of allowing information sharing in the industry.

The author finds that both firms and customers would benefit if firms shared with competitors their information about the state of demand. This result depends on the degree of demand uncertainty. Firms do not benefit from sharing their information on demand in very stable or in extremely volatile markets, but there is a wide range of intermediate cases where firms' profits increase when they share their private signals on demand.

4.5.4 Firms' Biased Beliefs About the Persistence of Demand Shocks

Goldfarb and Xiao (2018) study exit decisions of restaurants in Texas. First, the authors present empirical evidence that is consistent with the hypothesis that inexperienced restaurant owners do not take into account the transitory nature of weather shocks. They show that good weather has a significant positive effect on restaurant revenue. Weather has also a significant effect on the exit behavior of experienced owners after controlling for revenue: Given the same revenue record, experienced owners are more likely to exit in good weather than in bad weather. This is consistent with experienced owners understanding the transitory nature of weather shocks. In contrast, after conditioning on revenue, weather does not have any power to predict the exit behavior of inexperienced owners.

Motivated by this descriptive evidence, the authors propose and estimate a structural model of market exit where firms are Bayesian learners but firms may have attention costs that affect the number of state variables that they use to construct their strategies. More specifically, in this model, attention costs may prevent a firm from conditioning

on transitory weather shocks. Firms are heterogeneous in their attention costs, and these costs may decline with the owner's experience in the market.

The estimated model shows significant attention costs for inexperienced owners. This inattention is costly to firms. According to the estimated model, more than 3% of the restaurants in the dataset made mistaken exit decisions because of their limited attention.

4.6 Dynamic Models with Firms' Non-Bayesian Learning

4.6.1 Adaptive Learning About Demand

Jeon (2017) analyzes the link between uncertainty about the process of aggregate demand and boom–bust cycles of investment in the container shipping industry. She observes that demand for container shipping becomes increasingly volatile over time—especially after the financial crisis—and that the structural changes in the demand process may cause uncertainty for the firms. Jeon uses a dynamic oligopoly model of firm investment that incorporates uncertainty about the aggregate demand. Firms are uncertain about the parameters that govern the evolution of aggregate demand and they use new realizations of demand to update their beliefs based on adaptive learning. Firms are also allowed to assign heavier weights to more recent observations. An important parameter in firms' belief function is the coefficient that governs relative weights on past observations. To identify this parameter, the paper uses information on shipbuilding and scrap prices to estimate investment costs and scrap values separately from the dynamic model.

Jeon finds that the implied relative weight placed on an observation from 10 years ago is approximately 45% compared to the weight on the most recent observation. By incorporating learning, the study shows that uncertainty about the demand process raises the volatility of investment and the correlation between demand and investment, which helps to explain the fluctuations in investment that are observed in the data. Counterfactual experiments reveal that learning amplifies investment boom–bust cycles through: (i) leading agents to revise expectations more frequently and drastically as they experience demand volatility; and (ii) intensifying firms' strategic incentives.

4.6.2 Learning to Price After Market Deregulation

Huang et al. (2018) study firms' price setting behavior in the Washington State liquor market following the privatization of the market. The state of Washington privatized its previously monopolized market in 2012, which led existing grocery chains to newly enter the market.

The main research question concerns how these new entrants learn about demand and learn to price optimally over time. The authors first document that there are large and heterogeneous price movements in the first 2 years after the privatization. They provide some suggestive evidence that these price changes are due to

firms' learning about demand—including the rate at which prices respond to lagged demand shocks declines over time.

In order to quantify the effect of limited information on pricing and retailer profits, the authors estimate a structural model of demand for liquor. The authors first estimate a random coefficient logit demand model and then estimate firms' marginal costs—or, equivalently, the firms' wholesale costs.

A crucial assumption that is made in the estimation of the costs is that there is no learning and that all retailers price optimally by 2016. This assumption allows the authors to recover costs based on the first order conditions of profit maximization using data prior to 2016 alone. Then, under some parametric assumptions, costs for the later years are imputed. The differences between the optimal prices under these estimated costs and the observed prices in the data are then interpreted to be due to retailers' learning about demand in the new market. The counterfactual experiment reveals that as much as a 13% loss in the profits of inexperienced sellers can be attributed to their failure to price optimally (due to the information frictions).

4.6.3 Other Forms of Learning

Doraszelski et al. (2018) investigate firms' learning about competitors' bidding behavior just after the deregulation of the UK electricity market. Prior to the deregulation, electricity generating firms had to provide frequency response (FR) to the system operator at a fixed price. Deregulation created a market in which firms are allowed to bid for providing FR. The authors argue that—just after deregulation—it was difficult for firms to predict demand and rivals' bidding behavior such that firms faced considerable strategic uncertainty.

The authors distinguish three phases in the evolution of the FR market during the first 4 years after deregulation. In the early phase, bidding behavior was very heterogeneous, and firms made frequent and sizable adjustments in their bids. The middle phase of the FR market shows a dramatic reduction in the range of bids. At the last phase firms' bids become very stable. During these three phases, the demand for FR and the set of market participants were relatively stable. Therefore, the authors argue that the changes in firms' bidding strategies can be plausibly attributed to learning rather than to changes in the environment.

In the second part of the paper, the Doraszelski et al. (2018) study the process of firms' learning by estimating different structural models of learning. First, they estimate demand and cost functions under relatively weak rationality assumptions. Given demand and costs, they propose a model of firms' learning that combines fictitious play (Brown 1951) when firms learn about competitors' strategies, and adaptive learning (Sargent 1993) when they learn about demand parameters.

The authors consider different specifications of the learning model according to: (i) the value of the weighting parameter in fictitious play; and (ii) the demand parameters that firms are learning about (including the case of no uncertainty about demand parameters). They compare the performance of these learning models by using the mean square error of their predictions. The preferred learning model is one where firms have uncertainty about the parameter that captures the price sensitivity

of demand, and where fictitious play places little weight on the past decisions of competitors and learning is relatively quick.

4.6.4 Experience-Based Equilibrium

We conclude this section with a paper by Asker et al. (2016) that is not an empirical application, but presents interesting numerical results on the impact of firms' information sharing in a model of dynamic auctions with asymmetric information. The authors build a dynamic game of firms' participation and bidding in timber auctions. In each period firms decide whether to incur a cost to participate in the auction and bid for the right to harvest an unknown amount of timber. After the auction, firms' inventories are updated based on the outcome of the auction and the amount that they harvest. The source of private information is that each firm's inventory of unharvested timber is unknown to its rivals.

In this context, the authors model information sharing as a periodic complete revelation of the private information of every firm's inventory of timber. The authors consider two main scenarios: a baseline case where it is mandatory for firms to reveal their current inventories every $T > 1$ periods; and the Information Exchange case, where firms reveal information every period.

The authors numerically solve an Experience-Based Equilibrium (EBE) (Fershtman and Pakes 2012) of the model with the use of a reinforcement learning algorithm. The equilibrium of the model illustrates two different effects of revealing private information. First, more information intensifies competition by increasing participation and increasing bids in low-inventory states. However, information sharing also generates dynamic incentives to push the state variables (firms' inventories) to states in which competition is less fierce. For example, when a firm with high inventory knows that its competitor has low inventory, it would either not bid or bid the lowest amount that ensures winning. This second (dynamic) effect dominates the first.

On average, information sharing increases participation and inventories as firms are able to avoid intense competition in low-inventory states more effectively. According to the authors' numerical solutions, the effect of information on social welfare is close to zero. The lower bids do not have any effect on social welfare because they are transfers from the bidding firms to the seller. The positive welfare effect of greater inventories and output is offset by the negative effect of the increase in firms' participation costs.

5 Conclusions

Firms face significant uncertainty about demand and costs and about competitors' strategies. In some empirical applications, the assumption of equilibrium beliefs can be very unrealistic and can generate misleading interpretations of firms' strategic behavior and of our estimation of structural parameters in the profit function. A recent but growing literature in empirical industrial organization proposes and estimates non-equilibrium models of oligopoly competition. These empirical

studies offer new and alternative explanations for observed market outcomes and firm behavior in a broad range of settings: firm entry and competition in a new market or after a substantial change in the environment (e.g., deregulation); new technology adoption; strategic bidding in auctions; investment under demand uncertainty; and price-setting with uncertainty about product quality.

Firms' biased beliefs can be an important source of misallocation in some industries. We need empirical studies that measure the importance of this source of inefficiency. In some industries, firms' information sharing could reduce these inefficiencies, but again we need empirical studies that test this hypothesis.

An important challenge in this literature is that relaxing the assumption of rational beliefs may imply that almost any observed outcome could be explained in terms of exogenous variation in firms' beliefs. This calls for appropriate discipline on the assumptions about beliefs and on identification strategies. It also underlines the usefulness—in this context—of data on firms' expectations and on direct measures of firms' costs. Business surveys with questions about managers' probabilistic beliefs on uncertain events are still quite rare. This information can be very helpful for the identification of these models.

Nevertheless, as helpful as data on firms' beliefs can be, it is rarely sufficient to solve the identification problem in this class of models. Information on beliefs is typically incomplete. It does not cover the complete distribution of uncertain events conditional on state variables—not to mention its evolution over time. In business surveys, firms may have incentives to conceal their beliefs and do not report their truthful expectations. As we mention in Sect. 3.3, standard scoring methods that are used in experimental economics to elicit truthful beliefs are very difficult to implement when dealing with firms. In this context, access to data on firms' costs and profits can be a useful complementary approach to obtain identification.

Acknowledgements The authors would like to thank the comments and suggestions from the Editors, Victor J. Tremblay and Mo Xiao, and from the many generous colleagues who read a first version of this paper, and especially from Yonghong An, Avi Goldfarb, Xinlong Li, Matthew Osborne, Eduardo Souza-Rodrigues, and Erhao Xie.

References

- Aguirregabiria, V., & Mira, P. (2007). Sequential estimation of dynamic discrete games. *Econometrica*, 75, 1–53.
- Aguirregabiria, V., & Magesan, A. (2019). Identification and estimation of dynamics games when players' beliefs are not in equilibrium. *The Review of Economic Studies*, Forthcoming. <https://doi.org/10.1093/restud/rdz013>.
- An, Y. (2017). Identification of first-price auctions with non-equilibrium beliefs: A measurement error approach. *Journal of Econometrics*, 200, 326–343.
- An, Y., Hu, Y., & Xiao, R. (2018). *Dynamic decisions under subjective expectations: A structural analysis*. Baltimore: Department of Economics, Johns Hopkins University.
- Aradillas-Lopez, A., & Tamer, E. (2008). The identification power of equilibrium in simple games. *Journal of Business and Economic Statistics*, 26, 261–283.
- Armantier, O., Bruine, W., Topa, G., van der Klaauw, W., & Zafar, B. (2015). Inflation expectations and behavior: Do survey respondents act on their beliefs? *International Economic Review*, 56, 505–536.

- Asker, J., Fershtman, C., Jeon, J., & Pakes, A. (2016). The competitive effects of information sharing. NBER working paper, No. 22836. National Bureau of Economic Research.
- Astebro, T., Herz, H., Nanda, R., & Weber, R. (2014). Seeking the roots of entrepreneurship: Insights from behavioral economics. *Journal of Economic Perspectives*, 28, 49–70.
- Bachmann, R., Elstner, S., & Sims, E. (2013). Uncertainty and economic activity: Evidence from business survey data. *American Economic Journal: Macroeconomics*, 5, 217–49.
- Bajari, P., Chernozhukov, V., Hortaçsu, A., & Suzuki, J. (2018). The impact of big data on firm performance: An empirical investigation. NBER working paper, No. 24334. National Bureau of Economic Research.
- Bajari, P., Hong, H., Krainer, J., & Nekipelov, D. (2010). Estimating static models of strategic interactions. *Journal of Business and Economic Statistics*, 28, 469–482.
- Bernheim, B. (1984). Rationalizable strategic behavior. *Econometrica*, 52, 1007–1028.
- Borkovsky, R., Ellickson, P., Gordon, B., Aguirregabiria, V., Gardete, P., Grieco, P., et al. (2015). Multiplicity of equilibria and information structures in empirical games: Challenges and prospects. *Marketing Letters*, 26, 115–125.
- Bover, O. (2015). Measuring expectations from household surveys: New results on subjective probabilities of future house prices. *SERIEs Journal of the Spanish Economic Association*, 6, 361–405.
- Brown, G. (1951). Iterative solutions of games by fictitious play. In T. C. Koopmans (Ed.), *Activity analysis of production and allocation. Cowles commission monograph no. 13*. New York: John Wiley.
- Brown, A., Camerer, C., & Lovo, D. (2012). To review or not to review? Limited strategic thinking at the movie box office. *American Economic Journal: Microeconomics*, 4, 1–26.
- Brown, A., Camerer, C., & Lovo, D. (2013). Estimating structural models of equilibrium and cognitive hierarchy thinking in the field: The case of withheld movie critic reviews. *Management Science*, 59, 733–747.
- Camerer, C., & Ho, T. (1999). Experience-weighted attraction learning in normal form games. *Econometrica*, 67, 827–874.
- Camerer, C., Ho, T., & Chong, J. (2004). A cognitive hierarchy model of one-shot games. *Quarterly Journal of Economics*, 119, 861–898.
- Cheung, Y., & Friedman, D. (1997). Individual learning in normal form games: Some laboratory results. *Games and Economic Behavior*, 19, 46–76.
- Ching, A. (2010). Consumer learning and heterogeneity: Dynamics of demand for prescription drugs after patent expiration. *International Journal of Industrial Organization*, 28, 619–638.
- Ching, A., Erdem, T., & Keane, M. (2017). Empirical models of learning dynamics: A survey of recent developments. *Handbook of marketing decision models* (pp. 223–257). Cham: Springer.
- Cooper, A., Woo, C., & Dunkelberg, W. (1998). Entrepreneurs' perceived chances for success. *Journal of Business Venturing*, 3, 97–108.
- Costa-Gomes, M., & Crawford, V. (2006). Cognition and behavior in two-person guessing games: An experimental study. *American Economic Review*, 96, 1737–1768.
- Crawford, V., Costa-Gomes, M., & Iriberri, N. (2013). Structural models of nonequilibrium strategic thinking: Theory, evidence, and applications. *Journal of Economic Literature*, 51, 5–62.
- Crawford, V., & Iriberri, N. (2007). Level-k auctions: Can a nonequilibrium model of strategic thinking explain the winner's curse and overbidding in private-value auctions? *Econometrica*, 75, 1721–1770.
- Cyert, R., & DeGroot, M. (1974). Rational expectations and Bayesian analysis. *Journal of Political Economy*, 82, 521–536.
- DellaVigna, S. (2009). Psychology and economics: Evidence from the field. *Journal of Economic Literature*, 47, 315–372.
- DellaVigna, S. (2018). Structural behavioral economics. In B. Bernheim, S. DellaVigna, & D. Laibson (Eds.), *Handbook of behavioral economics*. Amsterdam: Elsevier.
- Doraszelski, U., Lewis, G., & Pakes, A. (2018). Just starting out: Learning and equilibrium in a new market. *American Economic Review*, 108, 565–615.
- Doraszelski, U., & Satterthwaite, M. (2010). Computable Markov-perfect industry dynamics. *Rand Journal of Economics*, 41, 215–243.
- Ellison, S., Snyder, C., & Zhang, H. (2018). *Costs of managerial attention and activity as a source of sticky prices: Structural estimates from an online market*. Cambridge: MIT, Department of Economics.
- Erev, I., & Roth, A. (1998). Predicting how people play games: Reinforcement learning in experimental games with unique mixed-strategy equilibria. *American Economic Review*, 88, 848–881.

- Ericson, R., & Pakes, A. (1995). Markov-perfect industry dynamics: A framework for empirical work. *Review of Economic Studies*, 62, 53–82.
- Evans, G., & Ramey, G. (1992). Expectation calculation and macroeconomic dynamics. *American Economic Review*, 82, 207–224.
- Evans, G., & Honkapohja, S. (1995). Local convergence of recursive learning to steady states and cycles in stochastic nonlinear models. *Econometrica*, 63, 195–206.
- Evans, G., & Honkapohja, S. (2001). Learning as a rational foundation for macroeconomics and finance. In R. Frydman & E. Phelps (Eds.), *Rethinking expectations: The way forward for macroeconomics*. Princeton: Princeton University Press.
- Evans, G., & Honkapohja, S. (2012). *Learning and expectations in macroeconomics*. Princeton: Princeton University Press.
- Feldman, M. (1987). Bayesian learning and convergence to rational expectations. *Journal of Mathematical Economics*, 16, 297–313.
- Fershtman, C., & Pakes, A. (2012). Dynamic games with asymmetric information: A framework for empirical work. *Quarterly Journal of Economics*, 127, 1611–1661.
- Fudenberg, D., & Levine, D. (1998). *The theory of learning in games*. Cambridge: MIT Press.
- Galasso, A., & Simcoe, T. (2011). CEO overconfidence and innovation. *Management Science*, 57, 1469–1484.
- Gardete, P. (2016). Competing under asymmetric information: The case of dynamic random access memory manufacturing. *Management Science*, 62, 3291–3309.
- Goldfarb, A., Ho, T., Amaldoss, W., Brown, A., Chen, Y., Cui, T., et al. (2012). Behavioral models of managerial decision-making. *Marketing Letters*, 23, 405–421.
- Goldfarb, A., & Xiao, M. (2011). Who thinks about the competition? Managerial ability and strategic entry in US local telephone markets. *American Economic Review*, 101(3130), 3161.
- Goldfarb, A., & Xiao, M. (2018). *Transitory shocks, limited attention, and a firm's decision to exit*. Tucson: Department of Economics, University of Arizona.
- Goldfarb, A., & Yang, B. (2009). Are all managers created equal? *Journal of Marketing Research*, 46, 612–622.
- Grossman, S. (1981). The informational role of warranties and private disclosure about product quality. *Journal of Law and Economics*, 24, 461–489.
- Gillen, B. (2010). *Identification and estimation of level-k auctions*. Pasadena: California Institute of Technology.
- Guerre, E., Perrigne, I., & Vuong, Q. (2000). Optimal nonparametric estimation of first-price auctions. *Econometrica*, 68, 525–574.
- Guiso, L., & Parigi, G. (1999). Investment and demand uncertainty. *Quarterly Journal of Economics*, 114, 185–227.
- Gureckis, T., & Love, B. (2013). *Reinforcement learning: A computational perspective*. New York: New York University.
- Heidhues, P., & Köszegi, B. (2018). Behavioral industrial organization. In B. Bernheim, S. DellaVigna, & D. Laibson (Eds.), *Handbook of behavioral economics*. Amsterdam: Elsevier.
- Heinemann, F., Nagel, R., & Ockenfels, P. (2009). Measuring strategic uncertainty in coordination games. *Review of Economic Studies*, 76, 181–221.
- Hortaçsu, A., & Puller, S. (2008). Understanding strategic bidding in multi-unit auctions: A case study of the Texas electricity spot market. *The RAND Journal of Economics*, 39, 86–114.
- Hortaçsu, A., Luco, F., Puller, S., & Zhu, D. (2017). Does strategic ability affect efficiency? Evidence from electricity markets. NBER working paper, No. 23526. National Bureau of Economic Research.
- Hu, Y. (2008). Identification and estimation of nonlinear models with misclassification error using instrumental variables: A general solution. *Journal of Econometrics*, 144, 27–61.
- Hu, Y., & Schennach, S. (2008). Instrumental variable treatment of nonclassical measurement error models. *Econometrica*, 76, 195–216.
- Huang, Y., Ellickson, P., & Lovett, M. (2018). *Learning to set prices in the Washington state liquor market*. Rochester: University of Rochester, Simon Business School.
- Huang, G., Luo, H., & Xia, J. (2015). *Invest in information or wing it? A model of dynamic pricing with seller learning*. Manuscript. Pittsburgh: Carnegie Mellon University.
- Jeon, J. (2017). *Learning and investment under demand uncertainty in container shipping*. Boston: Department of Economics, Boston University.
- Jovanovic, B. (1982). Selection and the evolution of industry. *Econometrica*, 50, 649–670.

- Landier, A., & Thesmar, D. (2009). Financial contracting with optimistic entrepreneurs. *Review of Financial Studies*, 22, 117–150.
- Lee, R., & Pakes, A. (2009). Multiple equilibria and selection by learning in an applied setting. *Economic Letters*, 104, 13–16.
- Levitt, S., & List, J. (2007). What do laboratory experiments measuring social preferences reveal about the real world? *Journal of Economic Perspectives*, 21, 153–174.
- Malmendier, U., & Tate, G. (2005). CEO overconfidence and corporate investment. *Journal of Finance*, 60, 2661–2700.
- Malmendier, U., & Tate, G. (2008). Who makes acquisitions? CEO overconfidence and the market's reaction. *Journal of Financial Economics*, 89, 20–43.
- Manski, C. (2004). Measuring expectations. *Econometrica*, 72, 1329–1376.
- Manski, C. (2018). Survey measurement of probabilistic macroeconomic expectations: progress and promise. *NBER Macroeconomics Annual*, 32(1), 411–471.
- Marcel, A., & Sargent, T. (1989a). Convergence of least-squares learning in environments with hidden state variables and private information. *Journal of Political Economy*, 97, 1306–1322.
- Marcel, A., & Sargent, T. (1989b). Convergence of least squares learning mechanisms in self-referential linear stochastic models. *Journal of Economic Theory*, 48, 337–368.
- Maskin, E., & Tirole, J. (1988). A theory of dynamic oligopoly, II: Price competition, kinked demand curves, and Edgeworth cycles. *Econometrica*, 3, 571–599.
- Milgrom, P. (1981). Good news and bad news: Representation theorems and applications. *The Bell Journal of Economics*, 12, 380–391.
- Morris, S., & Shin, H. (2002). *Measuring strategic uncertainty*. Princeton: Princeton University.
- Muth, J. (1961). Rational expectations and the theory of price movements. *Econometrica*, 29, 315–335.
- Nyarko, Y., & Schotter, A. (2002). An experimental study of belief learning using elicited beliefs. *Econometrica*, 70, 971–1005.
- Pakes, A., & McGuire, P. (2001). Stochastic algorithms, symmetric Markov perfect equilibrium, and the curse of dimensionality. *Econometrica*, 69, 1261–1281.
- Pearce, D. (1984). Rationalizable strategic behavior and the problem of perfection. *Econometrica*, 52, 1029–1050.
- Pesaran, H. (1987). *The limits to rational expectations*. New York: Basil Blackwell.
- Pesaran, H., & Weale, M. (2006). Handbook of economic forecasting. *Survey Expectations*, 1, 715–776.
- Potter, S., Del Negro, M., Topa, G., & Van der Klaauw, W. (2017). The advantages of probabilistic survey questions. *Review of Economic Analysis*, 9, 1–32.
- Samuelson, L. (1998). *Evolutionary games and equilibrium selection*. Cambridge: MIT press.
- Sargent, T. (1993). *Bounded rationality in Macroeconomics*. Oxford: Oxford University Press.
- Savage, L. (1971). Elicitation of personal probabilities and expectations. *Journal of the American Statistical Association*, 66, 783–801.
- Schotter, A., & Trevino, I. (2014). Belief elicitation in the laboratory. *Annual Review of Economics*, 6, 103–128.
- Simon, H. (1958). The role of expectations in an adaptive or behavioristic model. In M. Bowman (Ed.), *Expectations, uncertainty, and business behavior*. New York: Social Science Research Council.
- Simon, H. (1959). Theories of decision-making in economics and behavioral science. *American Economic Review*, 49, 253–283.
- Toivanen, O., & Waterson, M. (2005). Market structure and entry: Where's the beef? *The RAND Journal of Economics*, 36, 680–699.
- Van Huyck, J., Battalio, R., & Beil, R. (1990). Tacit coordination games, strategic uncertainty, and coordination failure. *American Economic Review*, 80, 234–248.
- Vives, X. (1993). How fast do rational agents learn? *The Review of Economic Studies*, 60, 329–347.
- Xie, E. (2018). *Inference in games without Nash equilibrium: An application to restaurants competition in opening hours*. Toronto: Department of Economics, University of Toronto.