

Convergence of Bayesian learning to general equilibrium in mis-specified models

Citation for published version (APA):

Schinkel, M. P., Tuinstra, J., & Vermeulen, A. J. (2002). Convergence of Bayesian learning to general equilibrium in mis-specified models. *Journal of Mathematical Economics*, 38(4), 483-508. [https://doi.org/10.1016/S0304-4068\(02\)00062-9](https://doi.org/10.1016/S0304-4068(02)00062-9)

Document status and date:

Published: 01/01/2002

DOI:

[10.1016/S0304-4068\(02\)00062-9](https://doi.org/10.1016/S0304-4068(02)00062-9)

Document Version:

Publisher's PDF, also known as Version of record

Please check the document version of this publication:

- A submitted manuscript is the version of the article upon submission and before peer-review. There can be important differences between the submitted version and the official published version of record. People interested in the research are advised to contact the author for the final version of the publication, or visit the DOI to the publisher's website.
- The final author version and the galley proof are versions of the publication after peer review.
- The final published version features the final layout of the paper including the volume, issue and page numbers.

[Link to publication](#)

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal.

If the publication is distributed under the terms of Article 25fa of the Dutch Copyright Act, indicated by the "Taverne" license above, please follow below link for the End User Agreement:

www.umlib.nl/taverne-license

Take down policy

If you believe that this document breaches copyright please contact us at:

repository@maastrichtuniversity.nl

providing details and we will investigate your claim.



ELSEVIER

Journal of Mathematical Economics 38 (2002) 483–508

JOURNAL OF
Mathematical
ECONOMICS

www.elsevier.com/locate/jmateco

Convergence of Bayesian learning to general equilibrium in mis-specified models

Maarten Pieter Schinkel^a, Jan Tuinstra^b, Dries Vermeulen^{a,*}

^a Department of Economics, Universiteit Maastricht, P.O. Box 616, Maastricht 6200 MD, The Netherlands

^b Department of Quantitative Economics, and CeNDEF, Universiteit van Amsterdam, Amsterdam, The Netherlands

Received 1 June 2001; received in revised form 4 April 2002; accepted 12 June 2002

Abstract

A central unanswered question in economic theory is that of price formation in disequilibrium. This paper lays the groundwork for a model that has been suggested as an answer to this question in, particularly, Arrow [Toward a theory of price adjustment, in: M. Abramovitz, et al. (Ed.), *The Allocation of Economic Resources*, Stanford University Press, Stanford, 1959], Fisher [Disequilibrium Foundations of Equilibrium Economics, Cambridge University Press, Cambridge, 1983] and Hahn [Information dynamics and equilibrium, in: F. Hahn (Ed.), *The Economics of Missing Markets, Information, and Games*, Clarendon Press, Oxford, 1989]. We consider sellers that monopolistically compete in prices but have incomplete information about the structure of the market they face. They each entertain a simple demand conjecture in which sales are perceived to depend on the own price only, and set prices to maximize expected profits. Prior beliefs on the parameters of conjectured demand are updated into posterior beliefs upon each observation of sales at proposed prices, using Bayes' rule. The rational learning process, thus, constructed drives the price dynamics of the model. Its properties are analysed. Moreover, a sufficient condition is provided, relating objectively possible events and subjective beliefs, under which the price process is globally stable on a conjectural equilibrium for almost all objectively possible developments of history.

© 2002 Elsevier Science B.V. All rights reserved.

JEL classification: C11; C51; C62; D21; D43; D83

Keywords: Disequilibrium theory; Oligopolistic price setting; Bayesian learning; Conjectural equilibrium; Global stability

* Corresponding author.

E-mail address: d.vermeulen@ke.unimaas.nl (D. Vermeulen).

1. Introduction

In economic theory, a key role in the coordination of behavior is played by prices. As a consequence, the so-called price mechanism is much debated, and the need for it operating freely often stressed. Yet, there are many open research questions on the matter of prices, especially on how they come to take on equilibrium values. For one thing, it is generally left unexplained whose business it actually is to call and change prices. Particularly in models in which price-taking behavior is assumed, this is a pressing question. Reliance on a unique price vector indicates it is left to a single person or institution, and a number of models has been presented in which the central person is in fact an altruistic auctioneer, e.g. in the tâtonnement process, the Edgeworth process, and the Hahn process.¹

Apart from the fact that it seems odd, if not plainly inconsistent, to model all behavior but that of the auctioneer as resulting from constrained rational choice, at least two things meet the eye in these explanations. First, these processes need an exogenous central coordinator to explain the rise of equilibria that are meant to be the outcome of decentralized competitive economies. Second, the conditions these processes need for convergence on equilibrium price values for arbitrary initial prices, i.e. for global stability of the disequilibrium process—have been found to be pretty strong.

A number of suggestions has been made to study the disequilibrium behavior of prices more seriously. An early one is in [Arrow \(1959\)](#), in which Arrow proposed to make price a choice variable of individual firms, that consequently need to come equipped with some local monopoly power, at least as a disequilibrium phenomenon. To Arrow, the construct of perfect competition did not allow for an explanation of price behavior. More recently, [Fisher \(1983\)](#) develops an elaborate model of disequilibrium behavior in which there is clarity on who is setting prices. It is done by dealers, who specialize in differentiated goods, which gives them the local monopoly power to act as a coordinator and set prices. How prices are adjusted with changes in perceptions, however, is not discussed in depth in the monograph, yet indicated as an area of promising further research. Finally, in [Hahn \(1989\)](#) several partial examples are given of perception changes and associated behavior that may indeed be plausible for monopolistically competing price-setters to develop—including a rudimentary version of the behavior we study in this paper. Yet, the consequences of such behavior, particularly when performed in general equilibrium settings, are only hinted upon.

When prices are choice variables of firms, the way firms perceive their market position, and especially changes in these perceptions, can account for the dynamics of prices. This idea is used in the present paper to construct a model of individually rational price adjustment and study its limit behavior, particularly its stability properties. In the present model, each of a number of firms is in monopolistic price competition, but does not have perfect information on the market demand it faces. At each moment in time, based on its information to date on past prices and sales, each firm entertains a demand conjecture instead. Naturally, this conjecture has a structural form different from that of objective demand. Particularly, we consider the most extreme case where firms only consider their own price as an explanatory variable, and do not consider the price effects of competing products.

¹ For an extensive survey of the disequilibrium literature and its problems, as well as pointers to an alternative modeling route on which the present model takes a small step (see [Schinkel, 2001](#)).

Within their conjectured structures, firms learn in a Bayesian way about the value of the demand parameters it has modeled. A fleshed out conjecture serves as a basis for an optimal price through expected profit maximization.

It is shown that for initial beliefs that do not assign zero probability to developments of prices and sales that can actually happen, the incomplete beliefs converge to a finite limit, and, therefore, prices converge as well. This is called ‘No Statistical Surprise’. Convergence takes place on a set of ‘conjectural equilibria’. Under ‘No Statistical Surprise’, therefore, the price process is globally stable in that it reaches an equilibrium for every initial belief-structure. Which particular equilibrium is reached depends on the initial beliefs. This path-dependency result runs solely over beliefs, since the model assumes absence of trade at disequilibrium prices. The stability result does not rely on specific conditions on the structure of objective demand. Instead, the condition of ‘No Statistical Surprise’ is sufficient for the perceived structure to absorb all price effects on objective demand.

The literature on Bayesian or rational learning is quite recent and large. Our paper builds on several of its results. One focus has been the concern to justify the use of rational expectations equilibria. Particularly [Bray and Savin \(1986\)](#), and [Bray and Kreps \(1987\)](#) work in this direction, and establish convergence results for myopic Bayesian learners on rational expectations equilibrium in versions of the cobweb-model. Early work by [Blume and Easley \(1982, 1986\)](#) is also concerned with the influence learning has on the eventual equilibrium situation reached, but in a general equilibrium setting. Particularly, they focus on conditions under which Bayesian learners will identify the true model among several models.

In partial equilibrium models of single firms learning their demand, [Easley and Kiefer \(1988\)](#) among others, study the influence of active learning on firms’ optimization problems. Actively learning firms are aware of the fact that their behavior influences their options for learning. In a discrete game theoretical setting, [Kalai and Lehrer \(1993, 1995\)](#) have obtained results for rational learning behavior. The former reference, [Kalai and Lehrer \(1993\)](#), considers learning in a correctly specified structure, and states conditions under which it converges to a Nash equilibrium of the perfect information game that are similar to ours.

Another, much less extensively traveled route has been to study the influence of structural mis-specification on the convergence process and its equilibria. [Kirman \(1975, 1983, 1995\)](#) sets up an early example of two firms learning, in a least squares way, in a mis-specified structure of their game. However, he does not establish general convergence results. [Nyarko \(1991\)](#) constructs an example of a single, actively learning monopolist whose beliefs do not settle, due to a very particular structural specification error. [Kalai and Lehrer \(1995\)](#) extends the 1993 convergence conditions to structurally mis-specified models to identify the usable notion of equilibrium. However, their article does not present explicit convergence results.

This paper is organized as follows. [Section 2](#) presents the model structure. [Section 3](#) introduces the way in which information is processed, as well as an associated equilibrium concept. [Sections 4 and 5](#) introduce the convergence result, the nature of which is subsequently discussed in [Section 6](#). [Section 7](#) presents the global stability of the price process on the equilibria of the model, introducing the concept of ‘No Statistical Surprise’. [Section 8](#) closes with some concluding remarks on possible extensions of the model.

2. The model

Consider an economy with n different firms. Each firm i has the ability to produce its own commodity, that is, it is supposed to take decisions on, particularly, price and quality of the commodity it produces. In this paper, we will assume that the aggregate of all these strategic choices to be made by firm i are incorporated into one single action space P_i . For technical reasons each P_i is assumed to be a convex and compact metric space.

2.1. Objective demand

The objective demand for each commodity i is assumed to be non-deterministic. In order to model this, let the commodity space of firm i be denoted by X_i . One can think of this space as being the collection of all possible realized demands for its commodity firm i might face. From a technical point of view, we only need few restrictions on this commodity space though, and we will only assume it is a compact metric space.

Suppose that firm i has decided to take action p_i in P_i . We write $p := (p_i)_{i \in N} \in P := \prod_i P_i$ for the entire vector of decisions taken. Now the demand for commodity i is supposed to be given by the density function

$$f_i(x_i|p)$$

with respect to the probability measure ν_i defined on the Borel σ -algebra $\mathcal{B}(X_i)$ generated by the metric on the commodity space X_i . For technical reasons, we assume that for any open set $U \subset X_i$ we have $\nu_i(U) > 0$. Further, by $f_i(x_i|p)$ being defined with respect to ν_i we mean that

$$\int_{X_i} f_i(x_i|p) d\nu_i = 1.$$

We will also assume that the function $f_i : X_i \times P \rightarrow \mathbb{R}$ is continuous.

2.2. Perceived demand

None of the firms is fully aware of the mechanism that generates the demand it faces. Instead, each firm i has a collection Θ_i of ‘worlds’ it deems possible. In world $\theta_i \in \Theta_i$, it conjectures that it serves a demand function that is distributed according to the density function

$$g_i(x_i|p_i, \theta_i)$$

with respect to ν_i . Again, we assume for technical reasons that Θ_i is a compact metric space and that $g_i : X_i \times P_i \times \Theta_i \rightarrow \mathbb{R}$ is continuous.

Note that subjective demand conjectures deviate importantly from objective demand: each firm only considers the effect of its own decision on the demand for its commodity, and neglects the influence of the decisions of the other commodities. In effect, each firm believes that it is a monopolist on its own market. This structural mis-specification reflects incomplete information on the side of the firms. We focus on this extreme situation where

only the effect of a firm's own decision is considered for reasons of exposition. The analysis could be extended to include less severe forms of incomplete information, e.g. structures in which the effects of the actions taken by several of the nearest competitors are included.

2.3. Expected profits

Within its structural mis-specification of how the world works, each firm i believes that there exists a 'true' world. However, it does not know which of possible worlds in Θ_i is the true one. Instead, the firm's perception of the world is stochastic. This means that each firm i has a belief represented by an element of the set $\mathbb{P}(\Theta_i)$ of probability measures on Θ_i . Such a belief $\mu_i \in \mathbb{P}(\Theta_i)$ assigns to each Borel subset A of Θ_i a real number $\mu_i(A)$ that reflects the probability firm i assigns to the event that the real world is an element of A .

Further, let

$$\pi_i(p_i, x_i) \in \mathbb{R}$$

be the net profit of demand x_i when firm i decides to take action p_i (we will assume throughout the paper that π_i is continuous). Then, given a belief μ_i of firm i , the amount $\Pi_i(p_i, \mu_i)$ of money firm i expects to earn is given by

$$\Pi_i(p_i, \mu_i) = \int_{\Theta_i} \int_{X_i} \pi_i(p_i, x_i) g_i(x_i, \theta_i | p_i) d\nu_i d\mu_i.$$

Since each firm i is assumed to be rational it will aim to maximize $\Pi_i(p_i, \mu_i)$ and take an optimal decision. In the remainder of this paper, we need each firm to have a unique optimal decision given the belief μ_i . In other words, we need to know that there is exactly one decision in P_i , which we will denote by $p_i(\mu_i)$, for which $\Pi_i(p_i(\mu_i), \mu_i)$ is larger than or equal to $\Pi_i(p_i, \mu_i)$ for any other possible action p_i of firm i in P_i . In order to guarantee the existence of such a unique optimal decision, we make the following assumption.

Assumption 1. Given the belief μ_i of firm i the function $\Pi_i(p_i, \mu_i)$ is strictly concave.

This is, for example, the case if for each x_i and θ_i the integrand $\pi_i(p_i, x_i) g_i(x_i, \theta_i | p_i)$ is strictly concave in p_i .

Note that $p_i(\mu_i)$ need not maximize expected profits in an objective sense. This is so since, although the world is in fact stochastic, it is stochastic in a way different from perception. More specifically, given the vector $p(\mu) := (p_i(\mu_i))_{i \in N}$ of individual decisions, objective demand is distributed on X_i according to

$$f_i(x_i | p(\mu)),$$

which shows how the true sales opportunities depend on the beliefs of all firms. And in turn these opportunities determine the objective expected net profit. In other words, the objective expected net profit of firm i is in fact given by

$$\int_{X_i} \pi_i(p_i, x_i) f_i(x_i | p(\mu)) d\nu_i.$$

No firm is, of course, capable of tuning its behavior to this true expected net profit.

3. Information processing and conjectural equilibrium

Beliefs are updated according to the Bayesian updating rule, as follows. Suppose that μ_i is the current belief of firm i in $\mathbb{P}(\Theta_i)$. Now the observation of demand x_i in X_i induces the updated belief $B_i(\mu_i)(x_i)$ in $\mathbb{P}(\Theta_i)$ that assigns to a Borel set $A \subset \Theta_i$ the probability

$$B_i(\mu_i)(x_i)(A) := \frac{\int_A g_i(\theta_i | p_i(\mu_i), x_i) \, d\mu_i}{\int_{\Theta_i} g_i(\theta_i | p_i(\mu_i), x_i) \, d\mu_i}$$

provided of course that the denominator is not equal to zero. In order to guarantee that this is the case, independent of the belief μ_i we make the following assumption.

Assumption 2. For all p_i, θ_i and x_i ,

$$g_i(x_i | p_i, \theta_i) > 0.$$

Given this assumption it can be shown that the above formula indeed yields a mapping

$$B_i : \mathbb{P}(\Theta_i) \times X_i \rightarrow \mathbb{P}(\Theta_i),$$

from the space of probability measures times the space of quantities X_i back to the space of probability measures.² This particular updating method, known as Bayesian updating, is firmly founded in probability theory. It is, therefore, sensible from the firms’ perspective to extract information from past observations in this way.

Although it does make perfect sense from the perspective of the firms, the learning process described is ill-founded in objective terms since it is based on an unrecognized structural mis-perception of demand. Hence, in general it cannot be hoped that subjective perceptions will come to explain the true demand for a commodity. Yet, there is a natural candidate for beliefs that are in ‘equilibrium’ with the objective world. Consider a single firm. The firm’s beliefs are in equilibrium if perceived optimal decisions made on the basis of this belief return quantities that are no ground for a revision of beliefs. This is the concept of individual conjectural equilibrium.

Definition 1. An individual conjectural equilibrium for firm i is a belief μ_i for which for all $x_i \in X_i$

$$B_i(\mu_i)(x_i) = \mu_i.$$

Since the observed sales depend on the decisions of all firms, it is quite special for a single firm to be in individual conjectural equilibrium. Yet, if all firms simultaneously are in individual conjectural equilibrium, none has a reason to deviate unilaterally from its decision, since no firm believes it can improve its position by doing so. This leads us to consider the following notion of an equilibrium for our economy.

Definition 2. A conjectural equilibrium is a vector $\mu = (\mu_i)_{i \in N}$ of beliefs such that each μ_i is an individual conjectural equilibrium.

² The technicalities supporting this statement can be found in [Appendices A and B](#).

4. Learning dynamics, infinite histories and beliefs

In Section 3, we saw that firms have mis-specified models of the true state of the world and they are not aware of this false interpretation of their environment. Nevertheless, given their mis-specification of the way the world works, they are aware of the fact that they are not fully informed about the true state of the world. This lack of information is modeled as a probability distribution μ_{i0} (the initial belief) over the collection Θ_i of all worlds that firm i deems possible. This belief reflects the amount of prior information firm i has concerning the true state of the world.

Now since each firm is a profit maximizer and since it is aware of the fact that it is not fully informed, it is eager to learn more about the true state of the world from market experience. It does so in the following way. Given its prior belief μ_{i0} firm i sets its (subjective) optimal decision $p_i(\mu_{i0})$. Once each firm has made this move, the objective demand density function establishes the quantities that can actually be sold given the actions $p_0 := p_i(\mu_{i0})_{i \in N}$. This means that for each firm i a quantity x_{i1} is drawn from the probability measure that assigns to each Borel set $A \subset X_i$ the probability

$$\int_A f_i(x_i | p_0) \, dv_i.$$

This new information is ground for a revision of beliefs via Bayesian updating. Repeating this procedure yields a learning process with the following properties.

At a given time $\tau = 0, 1, \dots$, each individual firm i has recorded a history of consumer demands

$$h_{i\tau} = (x_{it})_{t=1}^{\tau}$$

of finite length τ . This market information is the basis of the belief $\mu_{i\tau}(h_{i\tau})$ of firm i at time τ concerning the state of the world. It then takes a new action $p_i(\mu_{i\tau}(h_{i\tau}))$ based on its current belief. Given the vector $p_{\tau} := (p_i(\mu_{i\tau}(h_{i\tau})))_{i \in N}$ of new decisions, firm i observes a new quantity $x_{i\tau+1}$ drawn from the probability distribution that assigns to each Borel set $A \subset X_i$ the probability

$$\int_A f_i(x_i | p_{\tau}) \, dv_i.$$

Subsequently, beliefs are updated according to the Bayesian updating rule. Formally,

$$\mu_{i\tau+1}(h_{i\tau}, x_{i\tau+1}) := B_i(\mu_{i\tau}(h_{i\tau}))(x_{i\tau+1}).$$

Note that the decision on $p_i(\mu_{i\tau}(h_{i\tau}))$ the firm takes at time τ is a function only of the beliefs at time τ , which in turn derive from the initial beliefs μ_{i0} and the recorded history up until τ . Hence, it is sufficient to record sequences of observed quantities, as the firms do.

We have, thus, constructed a well-specified process in which beliefs lead to perceived optimal decisions p_{τ} , which serve as endogenous signals to obtain new information about the parameters of the distribution of objective demand. This new information, in turn, leads to an update of beliefs and, therefore, to new optimal decisions $p_{\tau+1}$ in a disequilibrium

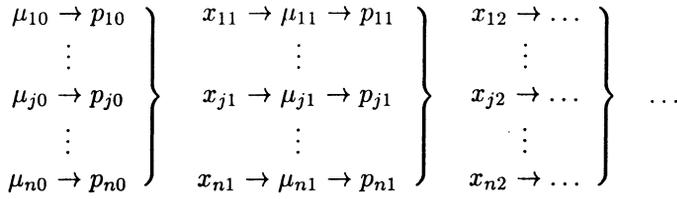


Fig. 1. The development of beliefs and prices over time.

price dynamics that embodies both subjectively rational learning and subjectively rational actions (Fig. 1).

In order to study the dynamic properties of this decision process, we make use of martingale convergence theory. For that purpose, we need to construct an underlying probability space on which we can identify martingales. This is the space of all possible future developments of history a firm *i* foresees at the beginning of time.³ Formally, let

$$H_{i\tau} := \prod_{t=1}^{\tau} X_i$$

be the space of all histories $h_{i\tau}$ of length τ . $\mathcal{B}(H_{i\tau})$ denotes the Borel σ -algebra on $H_{i\tau}$. Further, let $H_i := \prod_{t=1}^{\infty} X_i$ be the space of infinite histories. A specific element of H_i is denoted by h_i . By $\mathcal{B}(H_i)$ we denote the Borel σ -algebra generated by the product topology on H_i .

To complete the probability space of all future histories, we need a measure λ_i on $\mathcal{B}(H_i)$. Formally this λ_i is defined inductively on histories of finite length, combined with infinite extensions. We will now go through this construction step by step. First note that it is in fact sufficient to specify the numbers

$$\lambda_i \left(D_{\tau} \times \prod_{t=\tau+1}^{\infty} X_i \right)$$

for each Borel set D_{τ} in $H_{i\tau}$. Because, once these numbers are known, there is a unique way to extend λ_i to $\mathcal{B}(H_i)$. So we only need to specify the numbers

$$\lambda_{i\tau} \left(\prod_{t=1}^{\tau} D_t \right),$$

where $\lambda_{i\tau}$ is the probability measure induced by the beliefs of firm *i* up till time τ . Once these numbers are known, λ_i follows straightforwardly. In fact,

$$\lambda_i \left(\prod_{t=1}^{\tau} D_t \times \prod_{t=\tau+1}^{\infty} X_i \right) := \lambda_{i\tau} \left(\prod_{t=1}^{\tau} D_t \right),$$

³ We deviate somewhat from the structure generally chosen for this purpose, e.g. in Easley and Kiefer (1988), though in essence the spaces are the same.

the probability that an infinite history starts with a history $h_{i\tau}$ in the set $\prod_{t=1}^{\tau} D_t$. In order to specify these numbers we naturally start with $\lambda_{i0}(\emptyset) := 1$. Further, for $\tau = 1$,

$$\lambda_{i1}(D_1) := \int_{D_1} \int_{\Theta_i} g_i(x_i, \theta_i | p_i(\mu_{i0})) \, d\mu_{i0} \, dv_i.$$

In order to now define $\lambda_{i\tau+1}$ inductively, assume that $\lambda_{i\tau}$ is known. Let $h_{i\tau}$ be a history of length τ . Then the transition probability $\gamma_{i\tau+1}(h_{i\tau})(D_{\tau+1})$ of ending up in $D_{\tau+1} \subset X_i$ provided we have observed history $h_{i\tau}$ is equal to

$$\gamma_{i\tau+1}(h_{i\tau})(D_{\tau+1}) := \int_{D_{\tau+1}} \int_{\Theta_i} g_i(x_i, \theta_i | p_i(\mu_{i\tau}(h_{i\tau}))) \, d\mu_{i\tau}(h_{i\tau}) \, dv_i.$$

The transition probability gives the subjective probability of an observation $x_{i\tau+1}$ being in $D_{\tau+1}$ given that the firm has already observed history $h_{i\tau}$ and subsequently believes that $\mu_{i\tau}(h_{i\tau})$ is the appropriate probability distribution over Θ_i . We then have

$$\begin{aligned} \lambda_{i\tau+1} \left(\prod_{t=1}^{\tau+1} D_t \right) &:= \int_{H_{i\tau}} \int_{X_i} \mathbf{1}_{\prod_{t=1}^{\tau+1} D_t} \, d\gamma_{i\tau+1}(h_{i\tau}) \, d\lambda_{i\tau} \\ &= \int_{H_{i\tau}} \mathbf{1}_{\prod_{t=1}^{\tau} D_t} \int_{X_i} \mathbf{1}_{D_{\tau+1}} \, d\gamma_{i\tau+1}(h_{i\tau}) \, d\lambda_{i\tau} \\ &= \int_{H_{i\tau}} \mathbf{1}_{\prod_{t=1}^{\tau} D_t} \gamma_{i\tau+1}(h_{i\tau})(D_{\tau+1}) \, d\lambda_{i\tau} \\ &= \int_{\prod_{t=1}^{\tau} D_t} \gamma_{i\tau+1}(h_{i\tau})(D_{\tau+1}) \, d\lambda_{i\tau}. \end{aligned}$$

The definition reflects how $\lambda_{i\tau+1}$ derives as the weighted ‘sum’ (i.e. the integral) of all transition probabilities, where the weights are the probabilities $\lambda_{i\tau}$ the firm assigns to the observation that conditions the particular transition probability. The first step easily follows from rewriting the indicator function on the product set as a product of indicator functions. It is then observed that the inner integral equals $\gamma_{i\tau+1}(h_{i\tau})(D_{\tau+1})$. Finally, the indicator function is replaced by the restricted integral.

Now notice that, since the above computation implies that for all sets D_{τ} in $\mathcal{B}(H_{i\tau})$, we have

$$\lambda_{i\tau+1}(D_{\tau} \times X_i) = \int_{D_{\tau}} \gamma_{i\tau+1}(h_{i\tau})(X_i) \, d\lambda_{i\tau} = \int_{D_{\tau}} \mathbf{1}_{H_{i\tau}} \, d\lambda_{i\tau} = \lambda_{i\tau}(D_{\tau}),$$

the measures $\lambda_{i\tau}$ are consistent. Therefore, by the theorem of Kolmogorov, there is a unique probability measure λ_i on $\mathcal{B}(H_i)$ such that

$$\lambda_i \left(D_{\tau} \times \prod_{t=\tau+1}^{\infty} X_i \right) = \lambda_{i\tau}(D_{\tau}).$$

for all Borel sets D_{τ} in $\mathcal{B}(H_{i\tau})$.

An appealing way to think about λ_i is as the probability firm i initially assigns to observing the infinite history $h_i \in H_i$, based on its prior beliefs and its awareness of the learning process it is about the engage in. An example may help to clarify this.

Example 1. A stochastic variable X takes on one of two values, x_1 or x_2 . The probability of x_1 (and, hence, x_2) depends on a parameter θ , that is either θ_1 or θ_2 . Let $\Pr(x_1, x_2|\theta_1) = (1/3, 2/3)$ and $\Pr(x_1, x_2|\theta_2) = (1/2, 1/2)$ be the conditional probabilities of x_1 and x_2 , and suppose $\mu_0 = (1/4, 3/4)$ are the prior beliefs on (θ_1, θ_2) . Over time, a sequence of observations $(x_t)_{t \in \mathbb{N}}$ molds beliefs. We have

$$\gamma_1(X_1 = x_1) = \frac{1}{4} \times \frac{1}{3} + \frac{3}{4} \times \frac{1}{2} = \frac{11}{24} = \lambda_1(X_1 = x_1),$$

$$\gamma_1(X_1 = x_2) = \frac{1}{4} \times \frac{2}{3} + \frac{3}{4} \times \frac{1}{2} = \frac{13}{24} = \lambda_1(X_1 = x_2).$$

Suppose $X_1 = x_1$. Application of Bayes' rule now gives posterior beliefs

$$\mu_1|(X_1 = x_1) = \left(\frac{\frac{1}{4} \times \frac{1}{3}}{\frac{1}{4} \times \frac{1}{3} + \frac{3}{4} \times \frac{1}{2}}, \frac{\frac{3}{4} \times \frac{1}{2}}{\frac{1}{4} \times \frac{1}{3} + \frac{3}{4} \times \frac{1}{2}} \right) = \left(\frac{2}{11}, \frac{9}{11} \right).$$

Similarly, $X_1 = x_2$ would return

$$\mu_1|(X_1 = x_2) = \left(\frac{\frac{1}{4} \times \frac{2}{3}}{\frac{1}{4} \times \frac{2}{3} + \frac{3}{4} \times \frac{1}{2}}, \frac{\frac{3}{4} \times \frac{1}{2}}{\frac{1}{4} \times \frac{2}{3} + \frac{3}{4} \times \frac{1}{2}} \right) = \left(\frac{4}{13}, \frac{9}{13} \right).$$

We then have the conditional transition probabilities

$$\gamma_2(X_2 = x_1|X_1 = x_1) = \frac{2}{11} \times \frac{1}{3} + \frac{9}{11} \times \frac{1}{2} = \frac{31}{66},$$

$$\gamma_2(X_2 = x_2|X_1 = x_1) = \frac{2}{11} \times \frac{2}{3} + \frac{9}{11} \times \frac{1}{2} = \frac{35}{66},$$

$$\gamma_2(X_2 = x_1|X_1 = x_2) = \frac{4}{13} \times \frac{1}{3} + \frac{9}{13} \times \frac{1}{2} = \frac{35}{78},$$

$$\gamma_2(X_2 = x_2|X_1 = x_2) = \frac{4}{13} \times \frac{2}{3} + \frac{9}{13} \times \frac{1}{2} = \frac{43}{78}.$$

The λ -measure for the $t = 2$ paths is now constructed by combining the conditional transition probabilities, as follows.

$$\lambda_2(X_1 = x_1, X_2 = x_1) = \lambda_1(X_1 = x_1)\gamma_2(X_2 = x_1|X_1 = x_1) = \frac{11}{24} \times \frac{31}{66} = \frac{31}{144}.$$

Similarly, we find

$$\lambda_2(X_1 = x_1, X_2 = x_2) = \frac{11}{24} \times \frac{35}{66} = \frac{35}{144},$$

$$\lambda_2(X_1 = x_2, X_2 = x_1) = \frac{13}{24} \times \frac{35}{78} = \frac{35}{144},$$

$$\lambda_2(X_1 = x_2, X_2 = x_2) = \frac{13}{24} \times \frac{43}{78} = \frac{43}{144}.$$

Finally, the posteriors follow from Bayes' rule as

$$\mu_2|(X_1 = x_1, X_2 = x_1) = \left(\frac{4}{31}, \frac{27}{31} \right),$$

$$\mu_2|(X_1 = x_1, X_2 = x_2) = \left(\frac{8}{35}, \frac{27}{35}\right) = \mu_2|(X_1 = x_2, X_2 = x_1),$$

$$\mu_2|(X_1 = x_2, X_2 = x_2) = \left(\frac{16}{43}, \frac{27}{43}\right).$$

5. Convergence of beliefs and actions

The prime interest in this paper is to know whether, given initial beliefs, the process of Bayesian updating will eventually converge to a conjectural equilibrium. That is, we ask whether learning will teach some invariable posterior ideas, or whether perceptions, and thus decisions, will keep on changing forever. In order to address this question we will employ a convergence theorem concerning martingales. For that, we will show that on the probability space $(H_i, \mathcal{B}(H_i), \lambda_i)$ constructed above beliefs indeed form a martingale. To that end, we first need to introduce some notation.

Consider an infinite history $h_i = (x_{it})_{t=1}^\infty$ in H_i . The finite history $h_{i\tau} := (x_{it})_{t=1}^\tau$ in $H_{i\tau}$ is called the *truncation* of h_i till time τ . Further, let A be a Borel set in $\mathcal{B}(\Theta_i)$. Consider the function $\mu_{i\tau}(A)$ from H_i to \mathbb{R} that assigns to an infinite history h_i the real number

$$\mu_{i\tau}(A)(h_i) := \mu_{i\tau}(h_{i\tau})(A).$$

Secondly, notice that the above truncation of infinite histories to histories of length τ induces a natural identification of each element D_τ of the σ -algebra $\mathcal{B}(H_{i\tau})$ with the set

$$D_\tau \times \prod_{t=\tau+1}^\infty X_i$$

in $\mathcal{B}(H_i)$. The subalgebra of $\mathcal{B}(H_i)$ of sets of this form is denoted by $\mathcal{B}_\tau(H_i)$.

First notice that $\mathcal{B}_\tau(H_i)$ is a subset of $\mathcal{B}_{\tau+1}(H_i)$ and that each function $\mu_{i\tau}(A)$ is $\mathcal{B}_\tau(H_i)$ -measurable. We will now show that the sequence $(\mu_{i\tau}(A))_{\tau=1}^\infty$ is even a martingale. In order to do that we need the following.

Lemma 1. *Let ϕ a bounded and $\mathcal{B}_{\tau+1}(H_i)$ -measurable function. Then we have*

$$\int_{H_i} \phi(h_i) d\lambda_i = \int_{H_{i\tau}} \int_{X_i} \phi(h_{i\tau}, x_{i\tau+1}) \int_{\Theta_i} g_i(x_{i\tau+1} \theta_i | p_{i\tau}) d\mu_{i\tau}(h_{i\tau}) dv_i d\lambda_{i\tau}.$$

Proof. Let $D \times D_{\tau+1} \times \prod_{t=\tau+1}^\infty X_i$ be a Borel set in $\mathcal{B}_{\tau+1}(H_i)$. Then

$$\begin{aligned} & \int_{H_i} \mathbf{1}_{D \times D_{\tau+1} \times \prod_{t=\tau+1}^\infty X_i} d\lambda_i \\ &= \int_{H_{i,\tau+1}} \mathbf{1}_{D \times D_{\tau+1}} d\lambda_{i,\tau+1} = \lambda_{i,\tau+1}(D \times D_\tau) = \int_{H_{i\tau}} \mathbf{1}_D \gamma_{i,\tau+1}(h_{i\tau})(D_{\tau+1}) d\lambda_{i\tau} \\ &= \int_{H_{i\tau}} \int_{X_i} \mathbf{1}_{D \times D_{\tau+1}} \int_{\Theta_i} g_i(x_{i\tau+1}, \theta_i | p_{i\tau}) d\mu_{i\tau}(h_{i\tau}) dv_i d\lambda_{i\tau}. \end{aligned}$$

The same equality now easily follows for an arbitrary integrable function. □

Now, we can prove the following.

Theorem 1. *Let A be a Borel set in $\mathcal{B}(\Theta_i)$. Then the sequence $(\mu_{i\tau}(A))_{\tau=1}^{\infty}$ of random variables is a martingale with respect to λ_i .*

Proof. Let A be a Borel set in $\mathcal{B}(\Theta_i)$ and let C be a Borel set in $\mathcal{B}_{\tau}(H_i)$. We have to check that

$$\int_C \mu_{i\tau+1}(A)(h_i) \, d\lambda_i = \int_C \mu_{i\tau}(A)(h_i) \, d\lambda_i.$$

Since C is an element of $\mathcal{B}_{\tau}(H_i)$, we know it can be written as

$$D_{\tau} \times \prod_{i=1} X_i$$

for some Borel set D_{τ} in $H_{i\tau}$. So, since λ_i agrees with $\lambda_{i\tau+1}$ on $\mathcal{B}_{\tau+1}(H_i)$, Lemma 1 yields

$$\begin{aligned} & \int_{D_{\tau} \times \prod_{i=1} X_i} \mu_{i\tau+1}(A)(h_i) \, d\lambda_i \\ &= \int_{D_{\tau} \times X_i} \mu_{i\tau+1}(h_{i\tau}, x_{i\tau+1})(A) \, d\lambda_{i\tau+1} \\ &= \int_{D_{\tau}} \int_{X_i} \mu_{i\tau+1}(h_{i\tau}, x_{i\tau+1})(A) \int_{\Theta_i} g_i(x_{i\tau+1}, \theta_i | p_{i\tau}) \, d\mu_{i\tau}(h_{i\tau}) \, dv_i \, d\lambda_{i\tau}. \end{aligned}$$

Plugging Bayes' rule into this expression yields

$$\int_{D_{\tau}} \int_{X_i} \frac{\int_A g_i(x_{i\tau+1}, \theta_i | p_{i\tau}) \, d\mu_{i\tau}(h_{i\tau})}{\int_{\Theta_i} g_i(x_{i\tau+1}, \theta_i | p_{i\tau}) \, d\mu_{i\tau}(h_{i\tau})} \int_{\Theta_i} g_i(x_{i\tau+1}, \theta_i | p_{i\tau}) \, d\mu_{i\tau}(h_{i\tau}) \, dv_i \, d\lambda_{i\tau}$$

and the two integrals over Θ_i cancel out. Which reduces the above expression to

$$\int_{D_{\tau}} \int_{X_i} \int_A g_i(x_{i\tau+1}, \theta_i | p_{i\tau}) \, d\mu_{i\tau}(h_{i\tau}) \, dv_i \, d\lambda_{i\tau}.$$

To this expression we can apply the theorem of Fubini and switch the order of integration over X_i and A . This yields

$$\begin{aligned} & \int_{D_{\tau}} \int_A \int_{X_i} g_i(x_{i\tau+1}, \theta_i | p_{i\tau}) \, dv_i \, d\mu_{i\tau}(h_{i\tau}) \, d\lambda_{i\tau} \\ &= \int_{D_{\tau}} \int_A \mathbf{1}_{\Theta_i} \, d\mu_{i\tau}(h_{i\tau}) \, d\lambda_{i\tau} = \int_{D_{\tau}} \mu_{i\tau}(h_{i\tau})(A) \, d\lambda_{i\tau}, \end{aligned}$$

where the first equality results from the fact that g_i is a density function with respect to v_i . This concludes the proof. □

The result that beliefs form a martingale may not be very surprising. It states that the nature of Bayesian learning is such that a firm does not expect to change its beliefs in the future. Of course, an actual observation will in general change beliefs, but based upon

current beliefs on future realizations of sales, a firm ex ante predicts it will not. One way to interpret this is as Bayesian learning being sufficient, in that the information present at a given time is used to the full.

Example 2. In our earlier example, it is easy to see that beliefs have the martingale property. The expectation $E_\lambda(\cdot)$ taken with respect to λ is

$$E_\lambda(\mu_1(\theta_1)) = \frac{11}{24} \times \frac{2}{11} + \frac{13}{24} \times \frac{4}{13} = \frac{1}{4} = \mu_0(\theta_1),$$

and similarly

$$E_\lambda(\mu_2(\theta_1)) = \frac{11}{144} \times \frac{4}{31} + \frac{35}{144} \times \frac{8}{35} + \frac{43}{144} \times \frac{16}{43} = \frac{1}{4} = \mu_0(\theta_1).$$

With the above result in hand, we can apply martingale convergence theory to study the limit beliefs of agents, and hence of decisions, as follows.

Take an infinite history h_i in H_i . Let $\mu_{i\tau}(h_i)$ be the probability measure in $\mathbb{P}(\Theta_i)$ that assigns to each Borel set A of Θ_i the real number $\mu_{i\tau}(h_i)(A)$.

Theorem 2. *There exists a Borel set S of infinite histories in H_i with λ_i -probability one, on which the sequence $(\mu_{i\tau}(h_i))_{\tau=1}^\infty$ of probability measures converges weakly to a probability measure $\mu_{i\infty}(h_i)$ for every history h_i in S .*

Proof. We will first construct S . Since Θ_i is compact and metric, we know that there exists a countable basis of the topology. Let \mathcal{U} be the collection of finite intersections of elements of this basis. Take a fixed element U of \mathcal{U} . By [Theorem 1](#), the sequence $(\mu_{i\tau}(U))_{\tau=1}^\infty$ is a martingale. So, by the martingale convergence theorem (see, e.g. [Doob \(1990\)](#), Theorem 4.1 (iv)) there is a set $S(U)$ of infinite histories in H_i with $\lambda_i(S(U)) = 1$ such that $(\mu_{i\tau}(h_i)(U))_{\tau=1}^\infty$ converges for every history h_i in $S(U)$.

Now since \mathcal{U} is the collection of finite intersections of a countable collection, it is a countable set itself. This implies that

$$S := \bigcap_{U \in \mathcal{U}} S(U)$$

has λ_i -probability one, since it is a countable intersection of sets $S(U)$, all having λ_i -probability one.

The construction of the limit probability measure can be done as follows. Take a history h_i in S . First observe that, since Θ_i is compact, Theorem 6.1 of [Billingsley \(1968\)](#) states that $\mathbb{P}(\Theta_i)$ is sequentially compact. So, we know that a subsequence of $(\mu_{i\tau}(h_i))_{\tau=1}^\infty$ converges weakly to some probability measure, say $\mu_{i\infty}(h_i)$. We will show that the original sequence converges weakly to this probability measure. To this end, notice that

$$\mu_{i\tau}(h_i)(U) \rightarrow \mu_{i\infty}(h_i)(U), \quad \text{for all } U \in \mathcal{U}$$

for the original sequence, since this sequence is convergent for every element U of \mathcal{U} by construction of S , and so the above also holds for the weakly convergent subsequence. Moreover, \mathcal{U} is closed under finite intersections and each open set is obviously a countable union of elements of \mathcal{U} since \mathcal{U} contains a countable basis of the topology on Θ_i by construction.

Hence, by Theorem 2.2 of Billingsley (1968), $(\mu_{i\tau}(h_i))_{\tau=1}^\infty$ converges weakly to $\mu_{i\infty}(h_i)$ and the proof is complete since h_i was chosen arbitrarily in S . \square

From now on we will automatically assume that we only consider histories h_i in S whenever we talk about $\mu_{i\infty}(h_i)$. Effectively, we only consider the domain of $\mu_{i\infty}$. We can now prove the following result.

Theorem 3. *The sequence $p_i(\mu_{i\tau}(h_i))_{\tau=1}^\infty$ of actions λ_i -almost-surely converges to the limit decision $p_i(\mu_{i\infty}(h_i))$.*

Proof. By the continuity of p_i established in Lemma 4 of Appendix A, we know that the sequence $p_i(\mu_{i\tau}(h_i))_{\tau=1}^\infty$ of optimal decisions given beliefs at time τ converges to $p_i(\mu_{i\infty}(h_i))$ whenever the sequence $\mu_{i\tau}(h_i)_{\tau=1}^\infty$ of beliefs converges to $\mu_{i\infty}(h_i)$. This happens with λ_i -probability one by Theorem 2. \square

6. The nature of limit beliefs and limit actions

We now know that in our model beliefs, and consequently decisions, converge to limit beliefs and unique limit decisions, respectively, for λ_i -almost-all developments of history. In this section, we will derive some properties of the limit beliefs and decisions. We will show that a limit belief is unique in the sense that, roughly speaking, it only puts weight on worlds that generate the same probability distribution over demands. In other words, any two worlds in the support of a limit belief will have identical probability distributions over demands. Thus, a limit belief selects a unique possible world out of the collection of worlds that are possible initially, up to the identification of worlds of course that generate identical probability distributions. Furthermore, we will show that the limit beliefs obtained support a conjectural equilibrium.

6.1. Unique limit beliefs

For an analysis of the limit properties of beliefs and decisions, consider the following construction. Let μ_i be a probability measure on Θ_i . Evidently Θ_i is a compact set with $\mu_i(\Theta_i) = 1$. So, the collection

$$\mathcal{K} := \{K \subset \Theta_i \mid K \text{ is compact and } \mu_i(K) = 1\}$$

is not empty. Thus, we can define the support of μ_i by

$$\text{supp}(\mu_i) := \bigcap_{K \in \mathcal{K}} K.$$

The only question is whether this set has probability one according to μ_i . To this end, notice that the topology on Θ_i has a countable basis, say \mathcal{B} , since Θ_i is separable and metric. So,

$$\text{supp}(\mu_i) = \bigcap_{B \in \mathcal{B}; \mu_i(B)=0} \Theta_i \setminus B.$$

Hence, $\mu_i(\text{supp}(\mu_i)) = 1$ by the subadditivity of μ_i .

A more colloquial definition of the support of a probability measure μ_i on Θ_i is to say that it is the smallest compact subset K of Θ_i with $\mu_i(K) = 1$. In any event, it enables us to give the following definition.

Definition 3. A belief μ_i does not distinguish between signals if there exists a function $h_i : X_i \rightarrow \mathbb{R}$, such that for any θ_i in $\text{supp}(\mu_i)$ and for all x_i in X_i

$$g_i(x_i | p_i(\mu_i), \theta_i) = h_i(x_i).$$

This condition on μ_i states that every world θ_i in the support of μ_i generates the same density function on X_i . In other words, each signal has the same probability in each world in the support of μ_i . Consequently, no signal x_i will give firm i a reason to change its belief. A more interesting fact is that the converse of this observation is also true. This is reflected in the following result.

Theorem 4. A belief μ_i does not distinguish between signals if and only if

$$B_i(\mu_i)(x_i) = \mu_i$$

holds for all x_i in X_i .

Proof. Suppose that μ_i does not distinguish between signals. Then we can take $h_i : X_i \rightarrow \mathbb{R}$, such that

$$h_i(x_i) = g_i(x_i | p_i(\mu_i), \theta_i), \quad \text{for all } \theta_i \in \text{supp}(\mu_i).$$

Consequently, for any $x_i \in X_i$ and any Borel set A in Θ_i we have

$$\begin{aligned} B_i(\mu_i)(x_i)(A) &= \frac{\int_A g_i(x_i | p_i(\mu_i), \theta_i) \, d\mu_i}{\int_{\Theta_i} g_i(x_i | p_i(\mu_i), \theta_i) \, d\mu_i} = \frac{\int_A h_i(x_i) \mathbf{1}_{\text{supp}(\mu_i)} \, d\mu_i}{\int_{\Theta_i} h_i(x_i) \mathbf{1}_{\text{supp}(\mu_i)} \, d\mu_i} \\ &= \frac{h_i(x_i) \mu_i(A)}{h_i(x_i) \mu_i(\Theta_i)} = \mu_i(A). \end{aligned}$$

Suppose, on the other hand, that μ_i distinguishes between signals. Then we know that there is a pair $\zeta_i, \gamma_i \in \text{supp}(\mu_i)$, and an $x_i^* \in X_i$ for which

$$g_i(x_i^* | p_i(\mu_i), \zeta_i) > g_i(x_i^* | p_i(\mu_i), \gamma_i).$$

So, we can find two positive numbers $U > L \in \mathbb{R}$ and open neighborhoods $N(\zeta_i) \ni \zeta_i$ and $N(\gamma_i) \ni \gamma_i$ such that for all θ_i in $N(\zeta_i)$

$$g_i(x_i^* | p_i(\mu_i), \theta_i) \geq U$$

and for all θ_i in $N(\gamma_i)$

$$g_i(x_i^* | p_i(\mu_i), \theta_i) \leq L.$$

Now notice that $\mu_i(N(\zeta_i)) > 0$ since otherwise $\text{supp}(\mu_i) \setminus N(\zeta_i)$ would be a compact set with μ_i -probability one that is strictly included in $\text{supp}(\mu_i)$. For the same reason $\mu_i(N(\gamma_i)) > 0$.

This implies that

$$\frac{B_i(\mu_i)(x_i^*)(N(\zeta_i))}{B_i(\mu_i)(x_i^*)(N(\gamma_i))} \geq \frac{\int_{N(\zeta_i)} U \mathbf{1}_{\Theta_i} d\mu_i}{\int_{N(\gamma_i)} L \mathbf{1}_{\Theta_i} d\mu_i} = \frac{U\mu_i(N(\zeta_i))}{L\mu_i(N(\gamma_i))} > \frac{\mu_i(N(\zeta_i))}{\mu_i(N(\gamma_i))}.$$

So, at least

$$B_i(\mu_i)(x_i^*)(N(\zeta_i)) \neq \mu_i(N(\zeta_i))$$

or

$$B_i(\mu_i)(x_i^*)(N(\gamma_i)) \neq \mu_i(N(\gamma_i)).$$

In any case, $B_i(\mu_i)(x_i^*)$ does not equal μ_i and the proof is complete. □

The interpretation of this result is straightforward. A belief μ_i does not distinguish between signals if and only if Bayesian updating has no effect on the belief for any possible signal x_i . This fact has important implications. Particularly, since we can show that the limit beliefs $\mu_{i\infty}(h_i)$ in fact are fixed points of the Bayesian updating method, as we will do next, upon the following preliminary result.

Let \mathcal{B} be a countable basis of the topology on X_i . Let W be the collection of sample paths $(x_{it})_{t=1}^\infty$ in H_i for which there is a basis element B in \mathcal{B} such that $\{x_{it} | x_{it} \in B\}$ is finite. We will show first that the following is true.

Lemma 2. $\lambda_i(W) = 0$.

Proof. Let B be an element of \mathcal{B} and let T be a natural number. Define

$$W(B, T) := \{(x_{it})_{t=1}^\infty | x_{it} \notin B\}, \quad \text{for all } t \geq T.$$

Note that this construction is such that $W = \bigcup_{B,T} W(B, T)$. So, W is the countable union of sets $W(B, T)$. Hence, by the subadditivity of λ_i it suffices to prove that $\lambda_i(W(B, T)) = 0$ for any choice of B and T .

To this end, notice that

$$W(B, T) = \prod_{t=1}^{T-1} X_i \times \prod_{t=T}^\infty B^c.$$

Now take some $\tau \geq T$. Denote the subset $\prod_{t=1}^{T-1} X_i \times \prod_{t=T}^\tau B^c$ of the set $H_{i\tau}$ of finite histories up till time τ by $W_{i\tau}$. Then, for a history $h_{i\tau}$ in $W_{i\tau}$, the one-step transition probability $\gamma_{i\tau+1}(h_{i\tau})(B)$ to B is

$$\gamma_{i\tau+1}(h_{i\tau})(B) := \int_B \int_{\Theta_i} g_i(x_i, \theta_i | p_{i\tau}) d\mu_{i\tau}(h_{i\tau}) dv_i \geq \int_B \int_{\Theta_i} \varepsilon d\mu_{i\tau}(h_{i\tau}) dv_i = \varepsilon v_i(B).$$

Here, $\varepsilon > 0$ is chosen such that $g_i(x_i, \theta_i | p_{i\tau}) \geq \varepsilon$ for all x_i and θ_i , which can be done by the compactness of X_i , the continuity of g_i and the assumption that g_i is larger than zero on

X_i . Consequently, $\gamma_{i\tau+1}(h_{i\tau})(B^c) \leq 1 - \varepsilon v_i(B)$. Using this result, we get that

$$\lambda_{i\tau+1}(W_{i\tau+1}) := \int_{W_\tau} \gamma_{i\tau+1}(h_{i\tau})(B^c) d\lambda_{i\tau} \leq (1 - \varepsilon v_i(B))\lambda_{i\tau}(W_{i\tau}).$$

Now, backwards substitution yields

$$\lambda_{i\tau+1}(W_{i\tau+1}) \leq (1 - \varepsilon v_i(B))^{\tau-T+1}\lambda_{iT}(W_{iT}) \leq (1 - \varepsilon v_i(B))^{\tau-T+1}.$$

Further, since B is an open set, we know that $v_i(B) > 0$ by assumption. So, $0 \leq 1 - \varepsilon v_i(B) < 1$ and, hence,

$$\lim_{\tau \rightarrow \infty} \lambda_{i\tau+1}(W_{i\tau+1}) = 0.$$

Finally, notice that

$$0 \leq \lambda_i(W(B, T)) \leq \lambda_i \left(\prod_{t=1}^{T-1} X_i \times \prod_{t=T}^{\tau} B^c \times \prod_{t=\tau+1}^{\infty} X_i \right) = \lambda_{i\tau}(W_{i\tau})$$

where the second inequality follows from the monotonicity of λ_i and the equality from consistency of λ_i with $\lambda_{i\tau}$. Hence, it follows that $\lambda_i(W(B, T)) = 0$. \square

The interpretation of this result is that firms expect a priori that the signals they will receive are persistently exciting. That is, they expect to observe all possible quantities infinitely many times over the course of their learning process, so that they will be able to indeed extract sufficient information from them. The sufficiency of the information is reflected in this theorem.

Theorem 5. *There is a subset Z of S with λ_i -probability one such that the belief $\mu_{i\infty}(h_i)$ does not distinguish between signals for any h_i in Z .*

Proof. Let S be as in Theorem 2 and let W be as in Lemma 2. Write $Z := S \setminus W$. Clearly, $\lambda_i(Z) = 1$, since $\lambda_i(S) = 1$ and $\lambda_i(W) = 0$. Now take a history $h_i = (x_{i\tau})_{\tau=1}^{\infty}$ in Z . Then, since h_i is an element of S we know that $\mu_{i\infty}(h_i)$ exists. We will show that it does not distinguish between signals.

By Theorem 4 it suffices to show that $B(\mu_{i\infty}(h_i))(x_i) = \mu_{i\infty}(h_i)$ for all x_i in X_i . To this end, take an $x_i^* \in X_i$. Then, since $h_i = (x_{i\tau})_{\tau=1}^{\infty}$ is not an element of W , we know that it intersects each element of the basis \mathcal{B} infinitely many times. So, since X_i is metric, this implies that we can find a subsequence $(x_{i\alpha(\tau)})_{\tau=1}^{\infty}$ of $(x_{i\tau})_{\tau=1}^{\infty}$ such that $x_{i\alpha(\tau)} \rightarrow x_i^*$ as $\tau \rightarrow \infty$. Then, on the one hand,

$$B(\mu_{i\alpha(\tau)}(h_{i\alpha(\tau)}))(x_{i\alpha(\tau)+1}) = \mu_{i\alpha(\tau)+1}(h_{i\alpha(\tau)+1}) = \mu_{i\alpha(\tau)+1}(h_i) \rightarrow \mu_{i\infty}(h_i)$$

in the weak topology since the above sequence is a subsequence of $(\mu_{i\tau}(h_i))_{\tau=1}^{\infty}$ which converges to $\mu_{i\infty}(h_i)$ in the weak topology by the choice of S . On the other hand,

$$B(\mu_{i\alpha(\tau)}(h_{i\alpha(\tau)}))(x_{i\alpha(\tau)+1}) \rightarrow B(\mu_{i\infty}(h_i))(x_i^*)$$

since B is continuous by [Theorem 9](#) of [Appendix A](#). Hence, since the space $\mathbb{P}(\Theta_i)$ is Hausdorff, $\mu_{i\infty}(h_i) = B(\mu_{i\infty}(h_i))(x^*)$. \square

Note that if we make the natural assumption that conjectured density functions of demand are uniquely characterized by the value of θ_i , the result implies that the posterior distribution converges to a measure that puts all mass on one particular world θ_i in Θ_i .

Assumption 3. For any $p_i \in P_i$ we have $g_i(x_i|p_i, \zeta_i) = g_i(x_i|p_i, \gamma_i)$ for all $x_i \in X_i$ if and only if $\zeta_i = \gamma_i$.

A measure that puts all probability on one specific world θ_i is called a Dirac measure. We now have the following result.

Corollary 1. Suppose we have [Assumption 3](#). Then $\mu_{i\infty}(h_i)$ is a Dirac measure for every h_i in Z .

Proof. Let h_i be a history in Z . Then $\mu_{i\infty}(h_i)$ does not distinguish between signals by [Theorem 5](#). So, for any pair of worlds ζ_i and γ_i in the support of $\mu_{i\infty}(h_i)$ we have that

$$g_i(x_i|p_i(h_i), \zeta_i) = g_i(x_i|p_i(h_i), \gamma_i)$$

for the unique limit decision $p_i(h_i) := p_i(\mu_{i\infty}(h_i))$ in P_i and all x_i in X_i . Further, by [Assumption 3](#), this can only be the case if $\zeta_i = \gamma_i$. Hence, the support of $\mu_{i\infty}(h_i)$ is inevitably a singleton and $\mu_{i\infty}(h_i)$ is a Dirac measure. \square

6.2. Conjectural equilibrium

Provided that the structure of perceptions satisfies [Assumptions 1–3](#), we have shown that with λ_i -probability one, firm i 's belief is a Dirac measure $\mu_{i\infty}(h_i)$. Consequently, firm i 's limit decision is $p_i(h_i) := p_i(\mu_{i\infty}(h_i))$. Let $\theta_i(h_i)$ be the unique world in the support of $\mu_{i\infty}(h_i)$. The pair $(\theta_i(h_i), p_i(h_i))$ then specifies the limit stochastic view of the world of each firm. That is, each firm i perceives demand to be distributed in the limit as

$$g_i(x_i|p_i(h_i), \theta_i(h_i)).$$

We can now relate our results straightforwardly with our concept of equilibrium. We say that convergence is *almost sure* if it is λ_i -almost sure for every i .

Theorem 6. The learning process almost surely converges to a conjectural equilibrium.

Proof. By [Theorem 5](#) we know that the belief $\mu_{i\infty}(h_i)$ of firm i does not distinguish between signals on Z . So, by [Theorem 4](#) it is a fixed point of the Bayes operator and, hence, an individual conjectural equilibrium. Since this holds for every firm these beliefs form a conjectural equilibrium. \square

7. Objective convergence to conjectural equilibrium

We now know that for almost all developments of history to which a firm initially assigns non-zero probability, its beliefs on the parameters of conjectured market demand, and thereby the decisions it takes, converge to a unique limit belief that puts all mass on a single parameter of conjectured demand. For each firm, the limit decision is an individual conjectural equilibrium.

Since these results hold for every individual firm i , we are indeed close to conclusions on the behavior of the complete economy. However, since the conjectures that firms entertain are structurally mis-specified, their beliefs of possible developments of history need not necessarily match with the objective sequence of market demand they face. Consequently, actual histories may unfold that have λ_i -probability zero for some firms. Firms facing such probability zero histories will be unable to cope with it: Bayesian learning breaks down under such shocking surprises, and convergence fails. In order to exclude the rise of such paths, therefore, we need a condition that relates beliefs to objective probabilities.

The objective probability measure on the space of sample paths of the form $h_i \in H_i$ is potentially influenced by the behavior of all firms through the objective demand functions $f_i(x_i|p)$. In fact, for given initial beliefs μ_0 of the population, the unfolding sequence of individual actions that derives from the firms' sequential individual application of Bayes' rule within their conjectured demand structures, lays out a complete history of the world, when performed in the interrelated objective demand structures. For given priors, the only stochastic influence on the individually observed history h_i is from $f_i(x_i|p_\tau)$ for each τ .

The construction of objective probabilities on space H_i requires an objective probability measure ρ_i on $\mathcal{B}(H_i)$. Like λ_i , ρ_i is formally defined inductively on histories of finite length, combined with infinite extensions. For $\tau = 0$, we naturally have $\rho_{i0}(\emptyset) = 1$. In order to now define $\rho_{i,\tau+1}$ inductively, assume that $\rho_{i\tau}$ is known. Then, given that we have a history $h_{i\tau}$ of length τ , we can define the transition probability $\delta_{i\tau+1}(h_{i\tau})(D_{\tau+1})$ for each Borel subset $D_{\tau+1}$ of X_i as

$$\delta_{i\tau+1}(h_{i\tau})(D_{\tau+1}) = \int_{D_{\tau+1}} f_i(x_i|p_\tau) \, d\nu_i.$$

Again we can define

$$\rho_{i\tau+1} \left(\prod_{t=1}^{\tau+1} D_t \right) := \int_{X_{i=1}^\tau D_t} \delta_{i\tau+1}(h_{i\tau})(D_{\tau+1}) \, d\rho_{i\tau}$$

and apply again the Theorem of Kolmogorov. We now come to a crucial relationship between the objective and subjective probability measures ρ_i and λ_i .

Assumption 4. The probability measure ρ_i is absolutely continuous with respect to probability measure λ_i for every firm i .

In the interpretation that we have offered for the measures λ_i and ρ_i , absolute continuity of ρ_i with respect to λ_i implies that no actual development is possible that was not a priori foreseen as a possibility by the firm concerned. There is, therefore, 'No Statistical Surprise'

on the side of firms. This may seem strong, but is an assumption that is in fact often (implicitly) made in econometric specifications. Moreover, it seems a natural condition necessary for beliefs to settle down, as one can hardly expect beliefs to converge if all the time new and unforeseen events stir up the learning process. ‘No Statistical Surprise’ gives us the following prime result.

Theorem 7. *Beliefs almost surely converge to a conjectural equilibrium in the objective sense.*

Proof. By [Theorem 6](#) we have that the beliefs of each firm i converge to an individual conjectural equilibrium λ_i -almost surely. Since ρ_i is absolutely continuous with respect to λ_i , this convergence is also ρ_i -almost-surely. \square

Finally, again since p_i is continuous, we can conclude the following on the behavior of the individually rational disequilibrium price adjustment process.

Theorem 8. *The decision vector p_τ converges ρ_i -almost surely to a conjectural equilibrium decision vector p_∞ .*

8. Concluding remarks

The groundwork laid in this paper for modeling individually rational disequilibrium price adjustment by introducing elements of imperfect competition, imperfect and incomplete information and learning from self-generated signals, allows for a fairly strong global stability condition for general equilibrium models, that of ‘No Statistical Surprise’. Apart from being intuitively appealing, and doing away with the deus ex machina approach to disequilibrium behavior, this condition extends quite naturally on ‘No Favorable Surprise’, the global stability condition established in [Fisher \(1983\)](#). Our approach also calls for a number of extensions.

In the present model, firms gradually estimate the parameters of their conjectured demand. Naturally, provided they have some monopsony power as well, firms could likewise be taken to learn about supply, proposing purchase prices in the process. ‘No Statistical Surprise’ is likely to be strong enough to obtain convergence results in such a dealer-model as well.

A more demanding extension of the model would be to further specify the relationship between objective demand and supply structures and their subjective counterparts on which behavior is based. The present conjectures consider only the own price effect. Typically, firms would take the prices of several of their nearest competitors into account, applying econometric techniques in which the costs of including additional explanatory variables, or sharper functional forms, are weighted against the expected benefits of more precise predictions, thus, determining the best structural specification to work with. Such an approach would lead to an optimal level of mis-specification and introduce interesting problems concerning the strategic behavior towards rivals.

Related to this is the concept of active learning, where firms reckon with the fact that their prices will provide future information that can be used to increase profits. The

type of non-myopic price setting that results from this has been studied in a partial setting [Easley and Kiefer \(1988\)](#) and [Kiefer and Nyarko \(1989\)](#), where convergence results similar to the ones obtained here are established. The present price adjustment model would benefit greatly from an extension of dealer behavior in this direction, even though we expect ‘No Statistical Surprise’ to be powerful enough to again assure almost sure convergence.

The model presented in this paper relies on a specific and exogenously given structure of the market. Certain firms make it their business to act as intermediaries in the trade of a particular good. Casually, this setup has been defended by an appeal to product differentiation and transaction costs. It is to represent a socially accepted shopping-area structure. Although the identification of commodities with firms, which naturally leads to this market structure, seems quite appropriate in many markets, further specification of these underlying properties of markets is called for. Particularly, the consequences of entry and exit, and the possibility to compete for locally dominant dealerships raises interesting questions. For one thing, efforts to endogenize the market structure may well result in entry conditions that have the model sound more than presently like a disequilibrium story with a competitive ending.

Finally, the present model is altogether silent on the issue of social consequences of the disequilibrium dynamics modeled. In accordance with the observations by [Arrow \(1959\)](#) referred to earlier, it trades efficiency of competitive equilibrium for global stability by introducing monopolistic competition as an essential disequilibrium phenomenon. The precise welfare consequences of this seem a promising area of further research.

In the appendices we have collected those parts of the theoretical framework needed in the paper that would disrupt the flow of the argument too much when presented in the main body of the paper. [Appendix A](#) provides a detailed and complete proof of the continuity of the Bayes operator. [Appendix B](#) provides proofs concerning the support of some of the probability measures used in the text. These two appendices are largely based on [Easley and Kiefer \(1988\)](#).

Acknowledgements

With the usual disclaimer, we acknowledge helpful suggestions and comments by Franklin M. Fisher and an anonymous referee.

Appendix A. Continuity of the Bayes operator

In this appendix, we will show that the Bayes operator defined in [Section 3](#) is continuous. First of all, notice that the denominator in its definition is larger than zero by [Lemma 6](#). So, it is easy to see that $B(\mu_i)(x_i)$ is a non-negative function on the Borel σ -algebra on Θ_i . The σ -additivity of $B(\mu_i)(x_i)$ follows from the σ -additivity of the integral and finally it is obvious that $B(\mu_i)(x_i)(\Theta_i) = 1$. So, $B(\mu_i)(x_i)$ is indeed a probability measure, and the Bayes operator, thus, only takes on values in $\mathbb{P}(\Theta)$. Back to our aim, its continuity, we first need to establish some technicalities.

Lemma 3. *Suppose that we have a sequence $(p_{ik})_{k=1}^\infty$ that converges to some p_i . Then there is a number K such that for all $k > K$, all x_i and all θ_i*

$$|\pi_i(p_{ik}, x_i)g_i(x_i|p_{ik}, \theta_i) - \pi_i(p_i, x_i)g_i(x_i|p_i, \theta_i)| < \varepsilon.$$

Proof. Suppose not. Then for every number n there is a number $k(n) \geq n$ and points $x_i(n)$ and $\theta_i(n)$ such that

$$|\pi_i(p_{ik(n)}, x_i(n))g_i(x_i(n)|p_{ik(n)}, \theta_i(n)) - \pi_i(p_i, x_i(n))g_i(x_i(n)|p_i, \theta_i(n))| \geq \varepsilon.$$

Since both X_i and Θ_i are compact we may assume w.l.o.g that the sequence $x_i(n)_{n=1}^\infty$ converges to a point x_i and the sequence $\theta_i(n)_{n=1}^\infty$ converges to a point θ_i . However, since $k(n) \geq n$ by construction, we know that $p_{ik(n)} \rightarrow p_i$. Hence, taking limits yields

$$0 = |\pi_i(p_i, x_i)g_i(x_i|p_i, \theta_i) - \pi_i(p_i, x_i)g_i(x_i|p_i, \theta_i)| \geq \varepsilon$$

which is a contradiction. □

Lemma 4. *The function $p_i : \mathbb{P}(\Theta_i) \rightarrow P_i$ is continuous.*

Proof. Part (i): First, we will show that the expected payoff function

$$\Pi_i : P_i \times \mathbb{P}(\Theta_i) \rightarrow \mathbb{R}$$

is continuous. Of course, we suppose that $\mathbb{P}(\Theta_i)$ is endowed with the weak topology. Notice that his topology is metrizable by [Theorem 5](#) of Appendix III in [Billingsley \(1968\)](#). Therefore, it is sufficient to establish convergence of Π_i over sequences. So, take a sequence $(p_{ik}, \mu_{ik}) \rightarrow (p_i, \mu_i)$. We want to show that, given $\varepsilon > 0$, there exists a natural number K , such that for all $k \geq K$,

$$|\Pi_i(p_{ik}, \mu_{ik}) - \Pi_i(p_i, \mu_i)| \leq 2\varepsilon.$$

By the triangle inequality, we only need to show that

$$|\Pi_i(p_{ik}, \mu_{ik}) - \Pi_i(p_i, \mu_{ik})| + |\Pi_i(p_i, \mu_{ik}) - \Pi_i(p_i, \mu_i)| \leq 2\varepsilon$$

for sufficiently large k . We will show that both terms on the left hand side of the inequality sign are smaller than or equal to ε for sufficiently large k . The first term reads

$$\begin{aligned} &|\Pi_i(p_{ik}, \mu_{ik}) - \Pi_i(p_i, \mu_{ik})| \\ &= \left\| \int_{\Theta_i} \int_{X_i} [\pi_i(p_{ik}, x_i)g_i(x_i, \theta_i|p_{ik}) - \pi_i(p_i, x_i)g_i(x_i, \theta_i|p_i)] dv_i d\mu_{ik} \right\| \\ &\leq \int_{\Theta_i} \int_{X_i} |\pi_i(p_{ik}, x_i)g_i(x_i, \theta_i|p_{ik}) - \pi_i(p_i, x_i)g_i(x_i, \theta_i|p_i)| dv_i d\mu_{ik}. \end{aligned}$$

Now take K as in [Lemma 3](#). Then, since v_i and μ_{ik} are all probability measures, for each $k \geq K$ the latter expression is smaller than or equal to

$$\int_{\Theta_i} \int_{X_i} \varepsilon \mathbf{1}_{\Theta_i \times X_i} dv_i d\mu_{ik} = \varepsilon.$$

Furthermore, the second term reads

$$\left\| \int_{\Theta_i} \int_{X_i} \pi_i(p_i, x_i) g_i(x_i, \theta_i | p_i) \, dv_i \, d\mu_{ik} - \int_{\Theta_i} \int_{X_i} \pi_i(p_i, x_i) g_i(x_i, \theta_i | p_i) \, dv_i \, d\mu_i \right\|.$$

Because, we assume that $\mathbb{P}(\Theta_i)$ is endowed with the weak topology, it suffices to show that

$$F_p(\theta_i) := \int_{X_i} \pi_i(p_i, x_i) g_i(x_i | p_i, \theta_i) \, dv_i$$

is continuous in θ_i . To that end, take a sequence $\theta_{im} \rightarrow \theta_i$. Let $\varepsilon > 0$ be an arbitrary real number. Let G_p be a positive real number such that

$$\|\pi_i(p_i, x_i)\| \leq G_p, \quad \text{for all } x_i \in X_i.$$

This number exists because $\pi_i(p_i, x_i)$ is continuous in x_i and X_i is compact. Now take a natural number M_p such that for all $m \geq M_p$

$$\|g_i(x_i | p_i, \theta_{im}) - g_i(x_i | p_i, \theta_i)\| \leq \frac{\varepsilon}{G_p}.$$

Then for all $m \geq M_p$

$$\begin{aligned} & \|F_p(\theta_{im}) - F_p(\theta_i)\| \\ &= \left\| \int_{X_i} \pi_i(p_i, x_i) (g_i(x_i | p_i, \theta_{im}) - g_i(x_i | p_i, \theta_i)) \, dv_i \right\| \\ &\leq \int_{X_i} \|\pi_i(p_i, x_i)\| \|g_i(x_i | p_i, \theta_{im}) - g_i(x_i | p_i, \theta_i)\| \, dv_i. \end{aligned}$$

Consequently, since $v_i(X_i) = 1$,

$$\|F_p(\theta_{im}) - F_p(\theta_i)\| \leq \int_{X_i} G_p \frac{\varepsilon}{G_p} \, dv_i = \varepsilon.$$

Part (ii): Now let $(\mu_{ik})_{k=1}^\infty$ be a sequence converging to μ_i in the weak topology. Then, since P_i is a compact metric space, every sequence has a converging subsequence. So, we may assume without loss of generality that $p_i(\mu_{ik})$ converges to some decision p_i^* . We will now show that $p_i^* = p_i(\mu_i)$.

Since $p_i(\mu_{ik})$ is the optimal decision given the belief μ_{ik} , we know that for an arbitrary p_i in P_i it holds that

$$\Pi_i(p_i(\mu_{ik}), \mu_{ik}) \geq \Pi_i(p_i, \mu_{ik}), \quad \text{for all } k.$$

So, by the continuity of Π_i we get that

$$\Pi_i(p_i^*, \mu_i) \geq \Pi_i(p_i, \mu_i),$$

and p_i^* is an optimal action given belief μ_i since p_i was arbitrarily chosen. Hence, $p_i^* = p_i(\mu_i)$ by [Assumption 1](#). □

Furthermore, notice that $g_i : X_i \times P_i \times \Theta_i \rightarrow \mathbb{R}$ is also continuous. So, the function $h : X_i \times \mathbb{P}(\Theta_i) \times \Theta_i \rightarrow \mathbb{R}$ by

$$h(x_i, \mu_i, \theta_i) := g_i(x_i, p_i(\mu_i), \theta_i)$$

is continuous as well. Now suppose we have a sequence $(x_{in}, \mu_{in})_{n=1}^\infty$ converging to some limit (x_i, μ_i) . Define the functions f_n and f from Θ_i to \mathbb{R} by

$$f_n(\theta_i) := h(x_{in}, \mu_{in}, \theta_i) \quad \text{and} \quad f(\theta_i) := h(x_i, \mu_i, \theta_i).$$

Now take an arbitrary $\varepsilon > 0$. We then have the following lemma.

Lemma 5. *There exists a natural number N in \mathbb{N} such that $\|f_n - f\|_\infty < \varepsilon$ for all $n \geq N$.*

Proof. Suppose not. Then there is a subsequence $(f_k)_{k=1}^\infty$ of $(f_n)_{n=1}^\infty$ and a sequence $(\theta_{ik})_{k=1}^\infty$ such that

$$|f_k(\theta_{ik}) - f(\theta_{ik})| \geq \varepsilon, \quad \text{for all } k \in \mathbb{N}.$$

Since Θ_i is compact we may assume that the sequence $(\theta_{ik})_{k=1}^\infty$ converges to some limit θ_i . Then for all $k \in \mathbb{N}$

$$\varepsilon \leq |f_k(\theta_{ik}) - f(\theta_{ik})| = |h(x_{ik}, \mu_{ik}, \theta_{ik}) - h(x_i, \mu_i, \theta_{ik})|.$$

However, since $x_{ik} \rightarrow x_i$, $\mu_{ik} \rightarrow \mu_i$ and $\theta_{ik} \rightarrow \theta_i$, the continuity of h yields

$$\varepsilon \leq |h(x_i, \mu_i, \theta_i) - h(x_i, \mu_i, \theta_i)|$$

so that we arrive at a contradiction. □

The lemma is instrumental in the proof of the following.

Theorem 9. *The Bayes operator is continuous.*

Proof. Suppose that $(\mu_{in}, x_{in})_{n=1}^\infty$ converges to (μ_i, x_i) . It has to be shown that

$$B_i(\mu_{in})(x_{in}) \rightarrow B_i(\mu_i)(x_i)$$

as n goes to infinity. To this end, let F be a closed μ_i -continuous subset of Θ_i . By the Portmanteau theorem it is sufficient to show that

$$B_i(\mu_{in})(x_{in})(F) \rightarrow B_i(\mu_i)(x_i)(F).$$

By the definition of the Bayes operator,

$$B_i(\mu_i)(x_i)(F) = \frac{\int_F g_i(x_i | p(\mu_i), \theta_i) \, d\mu_i}{\int_{\Theta_i} g_i(x_i | p(\mu_i), \theta_i) \, d\mu_i}.$$

Now Lemma 6 in Appendix B guarantees that the denominator is strictly positive. So, since Θ_i itself is an instance of a closed set F whose boundary has measure zero (the

boundary of Θ_i is the empty set after all), it suffices in turn to show that, given $\varepsilon > 0$,

$$\left| \int_F g_i(x_{in}|p(\mu_{in}), \theta_i) d\mu_{in} - \int_F g_i(x_i|p(\mu_i), \theta_i) d\mu_i \right| < 2\varepsilon$$

for sufficiently large n . This is what we set out to do.

First, take N as in Lemma 5. Then for all $n \geq N$,

$$\begin{aligned} & \left| \int_F g_i(x_{in}|p(\mu_{in}), \theta_i) d\mu_{in} - \int_F g_i(x_i|p(\mu_i), \theta_i) d\mu_{in} \right| \\ & \leq \int_F |f_n(\theta_i) - f(\theta_i)| d\mu_{in} \leq \int_{\Theta_i} |f_n(\theta_i) - f(\theta_i)| d\mu_{in} \\ & \leq \int_{\Theta_i} \|f_n - f\|_\infty d\mu_{in} \leq \varepsilon \mu_{in}(\Theta_i) = \varepsilon, \end{aligned}$$

where the last inequality follows from the choice of n and N . So, now we only have to show that for all $\varepsilon > 0$, there exists an $N \in \mathbb{N}$ such that for all $n \geq N$

$$\left| \int_F g_i(x_i|p(\mu_i), \theta_i) d\mu_{in} - \int_F g_i(x_i|p(\mu_i), \theta_i) d\mu_i \right| < \varepsilon.$$

In other words, we have to show the existence of a natural number N such that for every $n \geq N$, μ_{in} is an element of the set of probability measures ν_i with

$$\left| \int_F f(\theta_i) d\nu_i - \int_F f(\theta_i) d\mu_i \right| < \varepsilon.$$

However, since this set is open in the weak topology and $(\mu_{in})_{n=1}^\infty$ weakly converges to μ_i , such an N exists. □

Appendix B. Sufficiently wide world views

For the Bayesian learning process to be well specified, we need that there are no objectively possible events that are assigned probability zero at any time by the firm. A Bayesian learner, namely, would simply not be able to deal with such events. Formally, it means that the denominator of the updating rule might become zero. In this appendix, we will show that Assumption 2 avoids this problem. Although also several somewhat weaker conditions would guarantee that the Bayesian learning process is well defined, we prefer to work with the above condition because of its simplicity. And that it is indeed sufficient is expressed in the following lemma.

Lemma 6. *Let μ_i be a belief in $\mathbb{P}(\Theta_i)$ and suppose that Assumption 2 holds. Let further a decision p_i , a demand x_i , and a Bol set $A \subset \Theta_i$ with $\mu_i(A) > 0$ be given. Then*

$$\int_A g_i(\theta_i|p_i, x_i) d\mu_i > 0.$$

Proof. Take a decision p_i and a demand x_i . Then we know that $g_i(\theta_i|p_i, x_i)$ is a continuous function in the variable θ_i since we even assumed that g_i is continuous in all three variables together. Moreover, Θ_i is compact. So, there exists a real number $\varepsilon > 0$ such that $g_i(\theta_i|p_i, x_i) \geq \varepsilon$ for all $\theta_i \in \Theta_i$. Consequently,

$$\int_A g_i(\theta_i|p_i, x_i) d\mu_i \geq \int_A \varepsilon \mathbf{1}_{\Theta_i} d\mu_i = \varepsilon \int_A \mathbf{1}_{\Theta_i} d\mu_i = \varepsilon \mu_i(A)$$

which is positive since both ε and $\mu_i(A)$ are positive by assumption. \square

References

- Arrow, K.J., 1959. Toward a theory of price adjustment. In: Abramovitz, M., et al. (Eds.), *The Allocation of Economic Resources*. Stanford University Press, Stanford.
- Billingsley, P., 1968. *Convergence of Probability Measures*. Wiley, New York.
- Blume, L., Easley, D., 1982. Learning to be rational. *Journal of Economic Theory* 26.
- Blume, L., Easley, D., 1986. Rational expectations equilibrium: an alternative approach. *Journal of Economic Theory* 34.
- Bray, M., Kreps, D., 1987. Rational learning and rational expectations. In: Feiwel, G. (Ed.), *Arrow and the Ascent of Modern Economic Theory*. New York University Press, New York.
- Bray, M., Savin, N.E., 1986. Rational expectations equilibria, learning, and model specification. *Econometrica* 54.
- Doob, J.L., 1990. *Stochastic Processes*. Wiley, New York.
- Easley, D., Kiefer, N.M., 1988. Controlling a Stochastic Process with Unknown Parameters. *Econometrica* 56.
- Fisher, F.M., 1983. *Disequilibrium Foundations of Equilibrium Economics*. Cambridge University Press, Cambridge.
- Hahn, F., 1989. Information dynamics and equilibrium. In: Hahn, F. (Ed.), *The Economics of Missing Markets, Information, and Games*. Clarendon Press, Oxford.
- Kalai, E., Lehrer, E., 1993. Rational learning leads to Nash equilibrium. *Econometrica* 61 (5).
- Kalai, E., Lehrer, E., 1995. Subjective games and equilibria. *Games and Economic Behaviour* 8.
- Kiefer, N.M., Nyarko, Y., 1989. Optimal control of an unknown linear process with learning. *International Economic Review* 30 (3).
- Kirman, A.P., 1975. Learning by firms about demand conditions. In: Day, R. (Ed.), *Adaptive Economics*. Academic Press, New York.
- Kirman, A.P., 1983. Mistaken beliefs and resultant equilibria. In: Frydman, R., Phelps, E. (Eds.), *Individual Forecasting and Collective Outcomes*. Cambridge University Press, Cambridge.
- Kirman, A.P., 1995. Learning in oligopoly: theory, simulation, and experimental evidence. In: Kirman, P.A., Salmon, M. (Eds.), *Learning and Rationality in Economics*. Blackwell, Oxford.
- Nyarko, Y., 1991. Learning in mis-specified models and the possibility of cycles. *Journal of Economic Theory* 55.
- Schinkel, M.P., 2001. *Disequilibrium Theory: Reflections towards a Revival of Learning*. Unpublished Ph.D. Thesis, Maastricht.