

Duopoly Competition, Escape Dynamics and Non-cooperative Collusion*

Batlome Janjgava, Sergey Slobodyan[†]

Abstract

In this paper, we study an imperfect monitoring model of duopoly under similar settings as in Green and Porter (1984), but here firms do not know the demand parameters and learn about them over time through the price signals. We investigate how a deviation from rational expectations affects the decision making process and what kind of behavior is sustainable in equilibrium. We find that the more common information firms analyze to update their beliefs, the more room is for implicit coordination. This might propagate escapes from the Cournot-Nash Equilibrium and the formation of cartels without explicit cooperative motives. In contrast to Green and Porter (1984), our results show that in a model with learning, breakdown of a cartel happens even without a demand shock. Moreover, in this model an expected price serves as an endogenous price threshold, which triggers a price war. Finally, by investigating the durations of the cooperative and price war phases, we find that in industries with a higher Nash equilibrium output and a lower volatility of firm-specific shocks, it is easier to maintain a cartel and harder to break it down.

Keywords: Beliefs, Escape Dynamics, Implicit Collusion, Self-Confirming Equilibrium, Learning

JEL Classification: D83, D43, L13, L40

*We would like to thank Krešimir Žigic and Levent Çelik for valuable comments and suggestions. All remaining errors are the responsibility of the authors. Batlome Janjgava is grateful for the financial support received from the World Bank Research Fellowship grant. E-mail: Batlome.Janjgava@cerge-ei.cz, Sergey.Slobodyan@cerge-ei.cz.

[†]Center for Economic Research and Graduate Education - Economics Institute, a joint workplace of Charles University in Prague and the Academy of Sciences of the Czech Republic. Address: CERGE-EI, P.O. Box 882, Politických veznu 7, Prague 1, 111 21, Czech Republic.

Abstrakt

V tomto článku studujeme nedokonalý monitorovací model duopolů v podobném prostředí jako Green a Porter (1984), zde ale firmy neznají parametry poptávky a poznávají je v čase pomocí cenových signálů. Zkoumáme, jak odchylka od racionálních očekávání ovlivňuje rozhodovací proces a jaký druh chování je udržitelný v rovnováze. Zjistíme, že čím jsou informace, které firmy analyzují při aktualizování svých názorů, všeobecnější, tím větší je prostor pro implicitní koordinaci. To může zvětšit množství niků z Cournot-Nashova ekvilibria a posílit utváření kartelů bez explicitních motivů pro kooperaci. Na rozdíl od Greena a Portera (1984) naše výsledky ukazují, že v modelu s učením dochází k rozpadu kartelů i bez poptávkového šoku. Navíc v tomto modelu očekávaná cena slouží jako endogenní hraniční cena, která vyvolá cenovou válku. Zkoumáním délek trvání kooperace a cenové války pak zjistíme, že v odvětvích s vyššími výstupy Nashovy rovnováhy a nižší volatilitou šoků specifických pro firmy je jednodušší udržovat kartel a těžší ho rozbít.

The difficulty lies not so much in developing new ideas as in escaping from old ones...

– *John Maynard Keynes, The General Theory of Employment Interest and Money, 1935*

1 Introduction

In this paper, we study an imperfect monitoring model of duopoly where firms do not know the demand parameters and learn about them over time through the price signals. Related literature includes the work of Stigler (1964), Green and Porter (1984), Slade (1989), and Ellison and Scott (2009). Stigler (1964) argues that a cartel might break down under imperfect information with unobservable firm action as a consequence of secret price cutting, because the rivals fail to distinguish between cheating and a negative demand shock. Green and Porter (1984) extend Stigler’s ideas of secret price cutting by introducing imperfect information (only the price is observable, but both the demand shock and rivals’ supply are not observable) into a dynamic model of oligopoly. Green and Porter (1984) show that when structural parameters are known and firms can calculate cooperative and Cournot-Nash equilibrium outcomes, we can have both collusive and Cournot-Nash equilibria as a non-cooperative rational expectations equilibrium outcome.¹

In practice firms never actually know the demand curve (Balvers and Cosimano, 1990). This implies that firms cannot calculate the Cournot-Nash equilibrium payoff and use it as a benchmark punishment payoff to deter deviations from an agreed upon output. Therefore, when structural parameters are unknown, the only sustainable equilibrium concept is the Nash equilibrium.² Nevertheless, we often observe tacit collusion and cooperative behavior in practice.

These observations lead us to consider the same model as Green and Porter (1984) but with unknown structural parameters and study whether learning itself can create room for implicit coordination and propagate cooperation. Moreover, we are interested in several other issues such as What are the main determinants and facilitators of cooperative behavior when structural parameters are unknown? Can we observe price wars? When do they occur? What determines the durations of the cooperative and price war phases? Which industries are more

¹Collusive equilibrium is supported by postulating the existence of a trigger price and a punishment period of a fixed duration, during which firms switch to the Cournot-Nash equilibrium for fixed amount of time whenever the price goes below the trigger price. After the fixed punishment period firms revert to the cartel-like behavior. However, no mechanism that allows coordinating on the trigger price and the punishment duration is provided. Abreu et al. (1986) extends Green and Porter’s analysis allowing for more general strategies and show that trigger strategies are indeed optimal.

²Kaneko (1982) shows that in supergames when there is no monitoring and ”when a player does not even know what single-period payoffs she has achieved thus far ... the only possible SPEs are strategy profiles that call for single-shot Nash actions in all periods” (Friedman, 2000, p. 65).

susceptible to cooperative behavior?

The extensive IO literature on learning can be classified into *active-learning* models, where the researcher is concerned with whether the agents could learn the true state or whether the deviations from complete learning remain significant in the long-run (Rothschild (1974), Balvers and Cosimano (1990), Rustichini and Wolinsky (1995), Keller and Rady (1999)); *experimentation* models, where researchers are interested in when and how myopically optimal actions can differ from the actions of agents who recognize that their current actions may influence their future expectations (Grossman et al. (1977), Balvers and Cosimano (1990), Keller and Rady (1999)); *signal-jamming* models, where rivals' actions are unobservable and agents attempt to manipulate their actions to affect the direction in which rivals update their beliefs (Riordan (1985), Fudenberg and Tirole (1986), Mirman et al. (1993a)); and *strategic information manipulation* models, where rivals' actions are observable and agents cannot influence the direction in which rivals update their beliefs but instead can affect the the informativeness of observed signals and thus the extent to which rivals are likely to update their expectations³ (Sakai (1985), Mirman et al. (1993b), Mirman et al. (1994)).

In this literature, unknown parameters are usually modeled as constants whose true values are not observable by the firms. The firms have some prior beliefs over the distribution of possible values of the parameters which are updated using Bayes' law upon the arrival of new information. In the models of active learning, however, the unknown parameters are often modeled as random variables. This allows the researcher to assess learning in a changing environment where there is positive probability that the unknown parameters may change. This often makes the learning "perpetual", or not ceasing even in the long run (Balvers and Cosimano, 1990, Rustichini and Wolinsky, 1995). In our model unknown parameters are constants. However, our learning is "perpetual" because we assume that the firms believe in the possibility of structural changes and incorporate this belief into their estimation model. We use the same belief updating mechanism — Bayes' law — as the above learning models, but we are especially interested in the effect of learning on optimal behavior and do not focus on the learning process per se. Therefore, we assume that firms are not concerned with *experimentation* or *signal-jamming* when making decisions and take their beliefs as remaining unchanged in the future. This makes firms in our model anticipated rather than expected profit maximizers in the sense of Kreps (1998).

³Researchers study under what conditions it is optimal for agents to influence through their current actions, the informativeness of observed signals to make it easier or more difficult to distinguish the "good" and "bad" states of the demand.

The firms in our paper adopt the Kalman filtering algorithm for belief updating. In our linear model, the shocks are normally distributed, and using Kalman filter is equivalent to using Bayes' law for prior updating (Balvers and Cosimano, 1990, Sargent, 1999). There are several papers from IO literature that use the Kalman filter for updating beliefs, among them Slade (1989), Slade (1990), Balvers and Cosimano (1990). From this list the closest paper to our work is by Slade (1989), who presents a way of generating price wars in a price-setting duopoly with unknown demand parameters. In Slade's model unknown parameters are random variables and may change with some small probability. When demand parameters change, firms start updating their beliefs about them using the Kalman filter, and this propagates price wars. Similar to Green and Porter (1984), the price war in Slade (1989) arises as a consequence of some exogenous process. Frequency of its occurrence is determined by the demand shocks or shifts in demand parameters, since the history of actions does not alter the distribution of exogenous shocks. We, however, show that even when unknown parameters are constant, price wars could happen endogenously, as a result of belief updating. Our findings are consistent with Ellison (1994), who reports that the Green and Porter model prediction that price wars happen only due to exogenous shocks in equilibrium is not supported by the data. In our model, price wars occur when the mean actual price falls below the mean forecasted price; thus, the mean forecasted price serves as an endogenous trigger of a price war and firms do not have to commit a priori to some agreed upon price trigger as in Green and Porter (1984).

Another difference that distinguishes our paper from Slade (1989) is as follows. In Slade (1989) model, if the unknown parameters do not change quite often, the learning ceases and the firms can calculate a new stationary Nash equilibrium. In our case learning never ceases; therefore, the firms never know the true values of unknown parameters and thus are never able to calculate the Nash equilibrium of the model. The recurrent escapes from a neighborhood of the Nash equilibrium of our model occur in the direction of the cooperative equilibrium, while in Slade's model, the direction of belief updating is determined by the exogenous shift in demand parameters.

In our model the equilibrium concept is the self-confirming equilibrium (SCE) developed and studied by Fudenberg and Levine (1993, 2009) who found widespread application in the models of adaptive learning, cf. Sargent (1999), Williams (2001), Cho et al. (2002), Williams (2003), Kasa (2004), Sargent et al. (2009), and Ellison and Scott (2009). In these models, the adaptive learning mechanism combined with the SCE concept is capable of generating recurrent

rapid deviations of beliefs from the SCE followed by a slow return. These recurrent deviations, called ‘escape dynamics’ after Sargent (1999), were used in the cited papers to explain various economic phenomena such as price volatility, business cycles, disinflation, and currency crises.

The paper which is closest to ours from the adaptive learning literature is by Ellison and Scott (2009), which studies a Hotelling’s duopoly model of non-renewable resource markets under adaptive learning. They illustrate the possibility of escapes from the Nash equilibrium towards a cooperative equilibrium and based on this explain the observed price volatility in non-renewable resource markets. Ellison and Scott analyze escape dynamics using the large deviation theory popularized and adopted for economic application by Williams (2001) and argue that the observed escape from the NE is mostly due to resource scarcity. In contrast to Ellison and Scott (2009), in this paper we show that occasional escapes may happen exclusively due to adaptive learning without the presence of resource scarcity; thus, learning itself is capable of propagating coordination on the cooperative equilibrium not only in exhaustible resource markets but also in regular commodity markets when rivals’ supply is unobservable and firms do not know the parameters of the inverse demand function.

Additionally, instead of focusing on price volatility generated by the self-referential learning process, we investigate dependence of expected duration of different game stages — competition, collusion, price war — on the structural parameters of the model. Based on a simple one-dimensional Brownian motion approximation approach, developed in Kolyuzhnov et al. (2006), we find simple analytical solutions for the belief thresholds that trigger escapes from Nash equilibria, and using the approach we find an analytical expression of expected time before escape to cooperation. Our analytical results based on the approximation using one-dimensional Brownian motion are confirmed by the simulations of the original problem.

The rest of the paper is structured as follows. We first describe the model and the sustainable equilibrium under Rational Expectations (RE). Then in Section 3, we develop firms’ control problem, motivate and define the self-confirming equilibrium (SCE), show that the SCE coincides with the Nash equilibrium (NE), give intuitive insights on possible sources of escapes from the NE, and illustrate them by numerical simulations. In Section 4, we formally analyze the mean and escape dynamics of beliefs, derive an analytical expressions of the belief trigger and derive analytical expression for mean exit time from the NE. Next, in Section 5, we characterize the nature of belief trajectories and study numerically the duration of the Cooperative Equilibrium (CE) and price war as well as the mean exit time from the NE. Section 6 concludes.

2 The Model Setup

Consider an oligopolistic industry comprised of two firms producing a single, homogeneous good. It is assumed that firms possess homogeneous production technology with constant marginal cost c . Each period, firm i produces $q_{in} = \hat{q}_{in} + \omega_{in}$ units of production where \hat{q}_{in} is controlled by the firm and ω_{in} a random disturbance. The controllable term \hat{q}_{in} is the expected profit maximizing output while ω_{in} is a firm-specific Gaussian shock, $\omega_{in} \sim N(0, \sigma_2^2)$. The industry faces a linear inverse demand schedule:

$$y_n = a - bQ_n + \omega_n, \quad (1)$$

where a and b are positive constants, $Q_n = q_{in} + q_{jn}$ is industry supply, and ω_n is the aggregate demand shock. In order to demonstrate that escapes could happen without an aggregate demand shock, we set its variance to zero.

A firm does not know the intercept (a) and slope (b) of the inverse demand curve. Additionally, it is uncertain regarding rivals' supply. This industry structure may describe the situation in many non-renewable resource extraction industries which are increasingly dominated by global conglomerates, and the information about actual production levels is unreliable in the short run as with, e.g., the OPEC countries' oil production levels.

Because of uncertainty of the structural parameters (a and b), the firms depart from rational expectations and use an adaptive learning algorithm to form expectations instead. The firms attempt to learn the unknown parameters from past observations: they form perceptions about the economic environment, construct beliefs about unknown parameters based on their perceptions and adopt a belief updating mechanism consistent with their belief system. Time varying beliefs affect firms' actions (quantity produced) which in turn affect prices. The firms then use newly available price and quantity data points to update their beliefs.

This interconnection of price signals, beliefs and actions establishes a self-referential process. Because of this self-referential nature under adaptive learning, our equilibrium concept is the self-confirming equilibrium (SCE) developed and studied by Fudenberg and Levine (1993, 2009), which was introduced into the adaptive learning literature by Sargent (1999). Typically, the SCE is defined as a point where agents' assumption of the orthogonality of the regression error term to regressors is satisfied; therefore, even though the agents' model of their economic environment might be misspecified, at the SCE they do not have an incentive to revise their subjective model which generates "correct" orthogonality prediction. In this paper, the agents will be using a

misspecified model in their learning, which makes the SCE an equilibrium concept appropriate for our purposes.

To define the SCE first of all we have to understand how a firm builds its subjective model based on the rational expectations equilibrium of the game. Therefore, in the next section we define the REE of our model and postpone the definition of the SCE till Section 3 where we deal with a firm's control problem. In our model, we do not consider strategic manipulations: the firms are assumed to believe that they cannot influence the rivals' decisions even indirectly through the price signals, and thus derive their actions based purely on the currently estimated demand function.

2.1 *Equilibrium under Rational Expectations*

In this subsection we abstract from learning, equip both firms with rational expectations, and solve for the rational expectation equilibrium (REE) of the model. We consider both the non-cooperative (Nash equilibrium (NE)) and the cooperative (collusive equilibrium (CE)) rational expectations equilibria of our model, and then investigate the stability of each equilibrium in our model settings under RE.

In the Nash equilibrium, each period the firm solves the following profit maximization problem taking the rivals' output as given:

$$\begin{aligned} \max_{\{\hat{q}_{ik}\}_n^\infty} \mathbb{E} \left[\sum_{k=n}^{\infty} \beta^{k-n} (y_k - c) q_{ik} \right] \\ \text{s.t. } y_k = a - b(q_{ik} + q_{jn}). \end{aligned}$$

Solving the above control problem for the symmetric equilibrium provides industry supply and price in the Nash equilibrium:⁴

$$\bar{Q} = \frac{2(a-c)}{3b}, \quad \bar{y} = \frac{a+2c}{3}. \quad (2)$$

In the collusive equilibrium, the joint expected profit is maximized and the individual firms

⁴In the rest of the paper we use \bar{x} and \tilde{x} for the equilibrium values of variable x in the NE and in the CE respectively.

share it equally since they are assumed to have the same marginal cost:

$$\begin{aligned} \max_{\{\tilde{Q}_k\}_n^\infty} \mathbb{E} \left[\sum_{k=n}^{\infty} \beta^{k-n} (y_k - c) Q_k \right] \\ \text{s.t. } y_k = a - bQ_k, \end{aligned}$$

Here, the symmetric collusive equilibrium industry supply and price are given as follows:

$$\tilde{q} = \frac{a - c}{4b}, \quad \tilde{y} = \frac{a + c}{2}. \quad (3)$$

As is well known, the static joint profit maximization problem is not “stable”. There is always temptation to produce more than the agreed upon output and undercut the price. Friedman (1971) provides the solution in supergames to this instability issue when the actions are observable. He shows that when a certain relation between discount factor and the number of firms holds, the CE can be sustained as an equilibrium by credible punishment threats.

However, we cannot use the oligopoly theory of Friedman (1971) in supergames to solve for the rational expectation equilibrium of the model since here actions are not observable. Moreover, uncertainty about structural parameters makes it impossible for firms to calculate equilibrium payoffs such as monopoly and Cournot-Nash payoffs. This effectively mean that firms do not have an equilibrium implied benchmark payoff such as the Cournot-Nash payoff to revert to as a punishment to deter deviations from an agreed upon output. Therefore, firms also cannot use the price-trigger strategies of Green and Porter (1984). These considerations make it clear that in our setting only the NE can be sustained as an equilibrium concept.⁵

After considering the REE of the model and reviewing the instabilities involved in creating the cartel, we are ready to introduce learning into the model and consider its implications for equilibrium outcomes. However, if the equilibrium concept is NE, it is not clear whether the collusive equilibrium can be observed in industries with the unobservability of rivals’ actions and this is the question addressed in the following sections.

⁵For the same reason, Nash is the equilibrium concept in the models with private information about costs (see for example Chakrabarti (2010) for further discussions).

3 Learning the Self-confirming Equilibrium

3.1 *The Firm's Control Problem*

Under adaptive learning a firm does not know the true parameters of the inverse demand curve but estimates them from the history of realized prices and own quantities produced. Here, we construct a firm control problem that governs the structure, estimation and consistency of its beliefs.

Every firm understands that the only sustainable equilibrium strategy of the game is to play the Nash strategy and that this is common knowledge. These assumptions lead the firms to postulate the time-invariance of their average beliefs. This establishes them as anticipated profit rather than expected profit maximizers in the sense of Kreps (1998), meaning that if we define a vector $\gamma_{in} = (\gamma_{in}^0, \gamma_{in}^1)'$ as a firm i 's belief about the intercept and slope of the inverse demand function at the time period n , then the belief evolution equation can be written as

$$\gamma_{in} = \gamma_{in-1} + \eta_{in}, \quad (4)$$

where $\eta_{in} \sim N(0, V_i)$. A firm's belief evolution equation (4) follows a locally constant random walk process. This type of belief formation is commonly used in economic literature (see Stock and Watson (1996), Sargent (1999), Cho et al. (2002)).

The absence of demand shocks and the time-invariance of beliefs lead firms to form the perception that the rivals' supply is time-invariant on average. This gives firms enough basis to build a subjective model that is consistent both with their beliefs and with the REE of the model. To distinguish between the "truth" and the firm's perceptions about the economy, we denote firm i 's observed price by y_{in} .⁶ Based on the firm's perception of constancy of rivals' supply, the perceived law of motion (PLM) of price is given by the following equation:

$$y_{in} = \gamma_{in}^0 + \gamma_{in}^1 q_{in} + u_{in}, \quad (5)$$

where $\gamma_{in}^0 = a - b\hat{q}_{jn}$, $\gamma_{in}^1 = -b$, and $u_{in} = -b\omega_{jn}$.

At the beginning of each period, firms determine the controllable part (\hat{q}_{in}) of its own supply

⁶Here, $y_n = y_{in}$ but y_n and y_{in} differ in how the total variation in price is decomposed into endogenous and exogenous components off the equilibrium path.

(q_{in}) by solving the following optimization problem:

$$\begin{aligned} \max_{\{\hat{q}_{ik}\}_n^\infty} \hat{\mathbb{E}} \left[\sum_{k=n}^{\infty} \beta^{k-n} (y_{ik} - c) q_{ik} | y^{n-1} \right] \\ \text{s.t. } y_{ik} = \gamma_{ik}^0 + \gamma_{ik}^1 q_{ik} + u_{ik}, \end{aligned} \quad (6)$$

where β is a common discount factor, y^{n-1} is a history of realized prices up to the period $n-1$, $y^n = \{y_0, y_1, y_2, \dots, y_n\}$ and $\hat{\mathbb{E}}[\cdot]$ denotes the firm's subjective expectation.

The expected profit maximization problem above assumes that a firm treats its beginning-of-period beliefs about the structural parameters as the true ones and solves the control problem (6) as if it had rational expectations.⁷ This assumption transforms the firm's dynamic optimization problem into a static one. Each period the firm solves the following static problem under the subjective expectations conditional on the history of price signals:

$$\begin{aligned} \hat{\pi}_{in} = \max_{\hat{q}_{in}} \hat{\mathbb{E}}[(y_{in} - c) q_{in} | y^{n-1}] \\ \text{s.t. } y_{in} = \gamma_{in}^0 + \gamma_{in}^1 q_{in} + u_{in}. \end{aligned} \quad (7)$$

The solution yields the firm's supply function dependent on its current beliefs:

$$\hat{q}_{in} = \max \left\{ \frac{\hat{\mathbb{E}}[\gamma_{in}^0 | y^{n-1}] - c}{-2 \hat{\mathbb{E}}[\gamma_{in}^1 | y^{n-1}]}, 0 \right\}. \quad (8)$$

A firm understands that its own action depends on its beliefs about the inverse demand function parameters and needs to adopt some estimation model to estimate them. Based on its perceptions and control problem, the firm builds the following estimation model:

$$\begin{aligned} y_{in} &= x'_{in} \gamma_{in} + u_{in}, \\ \gamma_{in} &= \gamma_{in-1} + \eta_{in} \end{aligned} \quad (9)$$

where $x_{in} = (1, q_{in})'$, $u_{in} \sim N(0, \sigma^2)$ and $\sigma^2 = b^2 \sigma_2^2$.

Each period, after the production decision is made, firms observe the price y_n and use Bayes' law to update their beliefs $\hat{\gamma}_{in+1} = \hat{\mathbb{E}}[\gamma_{in+1} | y^n]$ about the unknown parameters using their estimation model (9). The optimal Bayesian updating framework, based on the estimation model (9), is provided by the Kalman filtering technique and is described by the following

⁷In other words, under $\hat{\mathbb{E}}[\cdot]$ a firm fixes its future belief vectors at the current estimated value.

recursive algorithm:

$$\begin{aligned}\hat{\gamma}_{in+1} &= \hat{\gamma}_{in} + \frac{\hat{P}_{in}}{1 + x'_{in}\hat{P}_{in}x_{in}} x_{in}(y_{in} - x'_{in}\hat{\gamma}_{in}), \\ \hat{P}_{in+1} &= \hat{P}_{in} - \frac{\hat{P}_{in}x_{in}x'_{in}\hat{P}_{in}}{1 + x'_{in}\hat{P}_{in}x_{in}} + \sigma^{-2}V_i,\end{aligned}\tag{10}$$

where $\hat{P}_{in} = \sigma^2 \text{cov}[\gamma_{in} - \hat{\gamma}_{in}]$.⁸

3.2 *Self-confirming Equilibrium*

The firm's policy rule (8) under the PLM defines the actual law of motion (ALM) of the price given the true relationship (1) between price, y_n , and industry supply as follows:

$$y_n = a - b[q_{in}(\hat{\gamma}_{in}) + q_{jn}(\hat{\gamma}_{jn})].\tag{11}$$

Every firm treats the subjective model (5) it has in mind as if it were the true one. Therefore, they take the errors u_{in} to be orthogonal to the regressors x_{in} . This orthogonality is generally not observed in the data. However, there exists an equilibrium defined as a time-invariant vector of parameters $\bar{\gamma}_i$ for which the assumed orthogonality condition is satisfied, as the true mathematical expectation with respect to ALM of price (11) is indeed zero:

$$E[x_{in}(y_n - x'_{in}\bar{\gamma}_i)] = 0.\tag{12}$$

The Self-Confirming Equilibrium is defined as the unique symmetric vector of beliefs in which both firms' orthogonality assumptions are confirmed by observations.

Taking expectations in (12) produces the following expression:

$$\begin{bmatrix} \gamma_i^0 \\ \gamma_i^1 \end{bmatrix} = \begin{bmatrix} a - bE[q_{jn}] + \rho b E[q_{in}] \\ -b - \rho b \end{bmatrix},\tag{13}$$

where $\rho = \text{cov}[q_{jn}, q_{in}] / \text{var}[q_{in}]$.

In the SCE, the beliefs are fixed and thus there is no correlation between the firms' actions,

⁸Note that such behavior is optimal only if the firm does not take into account the effect of its own learning on its future beliefs and prices. In terms of Wieland (2000), the firms are engaged in "passive" learning.

therefore $\rho = 0$. A unique self-confirming equilibrium is then given by

$$\bar{\gamma}_i = \bar{\gamma}_j = (\bar{\gamma}_i^0, \bar{\gamma}_i^1)' = (\bar{\gamma}_j^0, \bar{\gamma}_j^1)' = \left(\frac{2a+c}{3}, -b \right)'. \quad (14)$$

From (8), at SCE both firms produce

$$\bar{q} \equiv \hat{q}_{in}(\bar{\gamma}) = \hat{q}_{jn}(\bar{\gamma}) = \frac{a-c}{3b}. \quad (15)$$

The SCE thus coincides with the Nash equilibrium (2). The reason for this is the fact that both equilibria are derived assuming that a firm considers rivals' actions to be independent of its own actions and constant over time. Therefore, deviation from the NE is equivalent to deviation from the SCE. These two concepts will be used interchangeably in what follows in the paper.

The expression for ρ in (13) represents the linear estimate of the conjectural variation that measures the degree of rivals' reaction to the firm's action. The belief that conjectural variation is zero eliminates any incentives to deviate from the SCE.⁹ However, if conjectural variations were non-zero, the beliefs would start to deviate from the SCE. A particular form of fast and sudden deviations from SCE is called escapes in economic literature after Sargent (1999).

Since the beliefs are estimated, they are a vector random variable. As both firms use the same price history and have the same PLM (5) to update beliefs, there are common factors in the forecast errors made by firms. These common factors allow the firms' beliefs to become correlated. The correlation of beliefs implies correlation of actions and might become a source for escaping from the SCE. Since firms are homogenous and use the same price signal to update their beliefs we expect ρ to be non-negative during an escape.

As mentioned above, the escapes are fast, which can be explained by the presence of a following positive feedback loop. If the firms' beliefs become correlated sufficiently strongly (ρ is sufficiently positive), the correlation between firms' actions increases so much that the beliefs are pushed even further from the SCE (ρ becoming even larger). Intuitively, we can think of the limit point of the escape as the point where firms' actions get perfectly correlated; $\rho = 1$. As is known from IO literature, ρ equal to one corresponds to the cooperative equilibrium. Beliefs that satisfy the orthogonality condition (12) with $\rho = 1$ have the form

$$\tilde{\gamma}_i = \tilde{\gamma}_j = (\tilde{\gamma}_i^0, \tilde{\gamma}_i^1)' = (\tilde{\gamma}_j^0, \tilde{\gamma}_j^1)' = (a, -2b)'. \quad (16)$$

⁹This point is proved in more rigorous way in Section 4.1, where we derive E-stability conditions of the SCE.

Given this belief vector each firm's produced output equals the CE output:

$$\tilde{q} \equiv \hat{q}_{in}(\tilde{\gamma}) = \hat{q}_{jn}(\tilde{\gamma}) = \frac{a - c}{4b}.$$

Thus, we can think of escapes as movements from the Nash equilibrium towards the cooperative equilibrium.

4 Escape Dynamics

In this section, we describe formally the most probable escape path and show that indeed escapes happen towards the direction of cooperative equilibrium. We derive the price thresholds that trigger the escapes from the Nash equilibrium towards the collusive equilibrium and vice versa. Then, we characterize mean escape time analytically and investigate the factors affecting the frequency of escapes. To study escape dynamics we draw on the approach developed in Kolyuzhnov et al. (2006).

In Section 3.1, the Kalman filter updating scheme was provided as the belief updating mechanism. Another updating mechanism often used in the literature in models with structural changes is constant gain recursive least squares learning (CG RLS), which is equivalent to weighted least square (WLS) with the weights geometrically falling with the age of data points. In order to facilitate the comparison with Ellison and Scott (2009), who use CG RLS, we follow Sargent and Williams (2005) in designing a Kalman filter algorithm that asymptotically has the same expected dynamics as the CG RLS. This equivalence is achieved if the belief updating scheme (10) is approximated as follows:

$$\begin{aligned}\hat{\gamma}_{in+1} &= \hat{\gamma}_{in} + \epsilon P_{in} x_{in} u_{in}, \\ \hat{\gamma}_{jn+1} &= \hat{\gamma}_{jn} + \epsilon P_{jn} x_{jn} u_{jn}, \\ P_{jn+1} &= P_{jn} + \epsilon \left(\sigma_2^{-2} \hat{V}_i - P_{in} M_i(\gamma_{in}) P_{in} \right), \\ P_{jn+1} &= P_{jn} + \epsilon \left(\sigma_2^{-2} \hat{V}_j - P_{jn} M_j(\gamma_{jn}) P_{jn} \right),\end{aligned}\tag{17}$$

and matrix \hat{V}_i is selected so that $\hat{V}_i = \sigma_2^2 M_i(\bar{\gamma}_i)^{-1}$. Here $u_{in} = y_n - x'_{in} \hat{\gamma}_{in}$ is a forecast error, ϵ the gain in a constant gain algorithm, $M_i(\gamma_{in}) = E[x_{in} x'_{in}]$, and $P_{in} = \epsilon^{-1} \hat{P}_{in}$. The expressions for j are analogous. Thus the defined approximate Kalman filter is asymptotically equivalent to the constant gain recursive least squares algorithm with gain ϵ . Evans et al. (2010, Proposition




FIGURE 1: Simulated time path of price for different specifications of the slope parameter, $a = 2$, $\sigma_2^2 = 0.01$, $\epsilon = 0.005$.

1, 2) show that the constant gain adaptive learning asymptotically approximates the optimal Bayesian estimation procedure if the beliefs are given by (4). Even if the firm believes the parameters γ_i to be constant over time, a form of the constant gain learning rule could be shown to be the (maximally) robust optimal predictor.

Simulating the system (17), we observe in Figure 1 that if the beliefs escape the Self-Confirming Equilibrium, they escape towards the cooperative equilibrium and then return back to the SCE. Figure 2 shows that escape from the SCE $\bar{\gamma} = (1.333, -0.100)'$ (see (14)) happens towards the belief vector $\tilde{\gamma} = (2.000, -0.200)'$ (see (16)) that corresponds to full collusion. This confirms our intuition in Section 3.2 about the limit point of escape in the belief space. Moreover, the escape dynamics towards the cooperative equilibrium lies along the line that connects the cooperative and the Nash equilibrium beliefs.

The observed escapes from the SCE are not a result of the instability of the belief updating process described by (17). To elaborate on this let us consider an average dynamics of (17) which is constructed as follows. Fix the beliefs $\gamma_{i,j}$ and $P_{i,j}$. Generate a process of prices y_n that would be obtained for these beliefs from equation (11). Then, take the expectation of the right-hand side of (17) with respect to the probability distribution induced by the derived stochastic process for y_n . Repeat this process for all possible beliefs and form the following




FIGURE 2: Simulated belief trajectory (circle) and line connecting CE and NE beliefs (dashed), $a = 2$, $b = 0.1$, $\sigma_2^2 = 0.01$, $\epsilon = 0.005$.

system of ordinary differential equations (ODE):

$$\begin{aligned}
 \dot{\gamma}_i &= P_i \bar{g}_i(\gamma_i, \gamma_j) \\
 \dot{\gamma}_j &= P_j \bar{g}_j(\gamma_i, \gamma_j) \\
 \dot{P}_i &= \sigma^{-2} \hat{V}_i - P_i M_i(\gamma_i) P_i, \\
 \dot{P}_j &= \sigma^{-2} \hat{V}_j - P_j M_j(\gamma_j) P_j,
 \end{aligned} \tag{18}$$

where $\bar{g}_i(\gamma_i, \gamma_j) = \mathbb{E}[x_{in}(y_n - x'_{in} \gamma_i)]$.

As discussed in Evans and Honkapohja (2001), solution paths of the resulting system of ODE approximate the behavior of the real-time learning dynamics (which takes place in discrete time) asymptotically as the gain parameter goes to zero. For more details on deriving the approximating ODE for the Kalman filter that is asymptotically equivalent to the constant gain RLS, consult Sargent and Williams (2005). The system (18) is referred to as the mean dynamics ODE, and its solution paths are called mean dynamics trajectories or simply mean dynamics. By construction, mean dynamics approximates the average behavior of the firm's belief when the updating process is given by (17).

A unique steady state of (18) is given by $\gamma_i^* = \gamma_j^* = \bar{\gamma}$ and $P_i^* = P_j^* = \bar{P} = M(\bar{\gamma})^{-1}$, which coincides with the SCE. In Appendix A, it is shown that the steady state of the system (18) is stable for any parameter values; we say that the SCE beliefs are E-stable. Evans and

Honkapohja (2001) show that if a stationary point of the mean dynamics ODE is stable, then the corresponding real-time dynamics with a decreasing gain converges to the stationary point with a probability that could be made as close to one as desired. For initial beliefs within a sufficiently small neighborhood of the SCE, the firms on average will have no incentives to deviate from the Nash equilibrium. Instead, if the firms are using constant gain learning as in this paper, they will not be able to learn the beliefs $\gamma_{i,j}^* = \bar{\gamma}$ and $P_{i,j}^* = M(\bar{\gamma})^{-1}$ exactly, but their beliefs will converge to an invariant distribution centered on $(\bar{\gamma}, \bar{P})$. Given E-stability of the SCE, we can say that the observed escapes are not due to its instability.

Simulations in Figure 1 clearly show that for some parameter values the beliefs may escape the SCE towards cooperative equilibrium. It is likely, therefore, that a threshold exists such that an escape from the SCE is triggered when the beliefs reach this threshold. We turn next to finding this belief threshold and studying the likelihood of escape from the SCE as a function of the structural parameters of the model.

4.1 *Analysis of Mean Dynamics and Trigger of Cooperation*

Since firms are homogeneous and their mean dynamics equations are the same, in order to find the belief threshold, it suffices to restrict our analysis to a symmetric case where $\gamma_j = \gamma_i$. We then analyze the following mean dynamics equations:

$$\begin{aligned}\dot{\gamma} &= P\bar{g}(\gamma), \\ \dot{P} &= \sigma^{-2}\hat{V} - PM(\gamma)P,\end{aligned}\tag{19}$$

where the belief vector γ is two-dimensional.

Following Kolyuzhnov et al. (2006), the escapes in the system (19) occur predominantly along the direction of the dominant eigenvector \bar{v} corresponding to the largest eigenvalue of \bar{P} .¹⁰ We restrict our attention to the line that starts at the SCE and extends in the direction \bar{v} . In parametric form, this line is given by the following expression:

$$\Gamma = \{\gamma \mid \gamma = \bar{\gamma} + \delta\bar{v}, \delta \in \mathbb{R}\}.\tag{20}$$

¹⁰If the constant gain parameter ϵ is not too small, the behavior of equation (17) is almost one-dimensional because λ_1 , the dominant eigenvalue of \bar{P} , is significantly larger than the second one. As a result, increments in the equation (17), given by

$$P_n g(\gamma_n, \xi_n) = \sum a_i \lambda_i v_i \approx a_1 \lambda_1 v_1$$

are concentrated in a narrow cone formed around v_1 , the dominant eigenvector of \bar{P} (the eigenvalues and eigenvectors of P_n are essentially the same as those of \bar{P} near the SCE). For similar ideas one can also consult the method of principal component analysis.

Because the dominant eigenvalue of \bar{P} ($\bar{\lambda}_1 = 4545.40$) is significantly larger than the second one ($\bar{\lambda}_2 = 0.02$), solution paths of (19) along the escape route are likely to lie very close to Γ . Therefore, Γ represents the dominant escape path.

The escape vector \bar{v} can be written as $((1 - 1/\bar{\lambda}_1)^{-1}\bar{q}, -1)^T$, where $\bar{\lambda}_1$ is the largest eigenvalue of \bar{P} . The direction \bar{v} is closely aligned with the vector $\tilde{\gamma} - \bar{\gamma}$, whenever $\bar{\lambda}_1$ tends to infinity. The latter can happen whenever σ_2^2 tends to zero. Therefore, the likelihood of observing escapes from the Nash equilibrium (beliefs given by $\bar{\gamma}$) towards exactly the cooperative equilibrium (beliefs $\tilde{\gamma}$) is higher the lower is the volatility of firm-specific shocks.

The belief updating procedure minimizes volatility of a forecast error u_{in} , which suggests that the belief updating is sensitively dependent on the behavior of the mean forecast error. Therefore, to find the belief trigger, we should understand how beliefs γ and the mean forecast error $u = E[u_{in}]$ interact with each other. Since the right-hand side of the mean dynamics equations (18) depends on $\bar{g}_i(\gamma_i, \gamma_j) = E[x_{in}u_{in}]$, the mean forecast error can be used to judge whether the belief updating process is accelerating or stabilizing. When the mean forecast error is growing fast the belief updating accelerates, while we should observe the opposite when the growth of the mean forecast error slows down. Since we observe escapes from the SEC, even though the SCE is a stable steady-state of the system (18), there should exist a neighborhood of the SCE such that the mean forecast error is declining over time ($\dot{u}/u < 0$) within it and growing ($\dot{u}/u > 0$) outside of it.

The above logic is illustrated in Figure 3 where mean belief trajectories are plotted with different initial beliefs $\gamma \in \Gamma$ together with the curves $\dot{u} = 0$ and $u = 0$ near the SCE. We see that on the segment ANE of line Γ , the mean forecast error is declining; $\dot{u}/u < 0$. If the initial beliefs are located in this segment, the average forecast error is expected to decline, meaning that beliefs are expected to move towards the SCE. However, for an initial belief $\gamma \in \Gamma$, outside of the ANE segment, the beliefs diverge from the SCE since $\dot{u}/u > 0$, and they are expected to move along Γ towards the CE. We see that the neighborhood of $\bar{\gamma}$, where $\dot{u}/u < 0$, behaves like the region of attraction of the SCE along the escape direction; outside of this region the beliefs escape the Self-Confirming Equilibrium.

When firms escape the Nash equilibrium and start cooperating, the industry supply shrinks. Therefore, the belief threshold γ^* that triggers cooperation and at the same time escape from




FIGURE 3: Phase diagram of beliefs, $a = 2$, $b = 0.1$, $\sigma_2^2 = 0.01$. Two types of mean belief trajectories are depicted: one with initial beliefs inside the attraction region of SCE along the dominant escape path and another for initial beliefs outside the attraction region of SCE along the dominant escape path. Mean belief trajectories are shown by arrows.

the SCE should satisfy the following equation:

$$\begin{cases} \dot{u} = \nabla u(\gamma) \dot{\gamma} > 0, \\ \dot{Q} = \nabla Q(\gamma) \dot{\gamma} < 0. \end{cases}$$

The $\dot{Q} < 0$ restriction is redundant because the escape of beliefs from the SCE implies a cooperative behavior: for $\gamma \in \Gamma$, $\dot{u} > 0$ implies $\dot{Q} < 0$. This implies that we can find the belief threshold $\gamma^* \in \Gamma$ by solving an algebraic equation $\dot{u} = 0$ for δ^* :

$$\dot{u} = \nabla u(\bar{\gamma} + \delta^* \bar{v}) \bar{P} \bar{g}(\bar{\gamma} + \delta^* \bar{v}) = 0.$$

This yields the following approximate expression for δ^* :

$$\delta^* \approx 6 \frac{\sigma_2^2}{\bar{q}^2} b. \quad (21)$$

The analysis above and in the next section relies on the one-dimensional nature of escaping beliefs to a significant degree. In case of a strictly one-dimensional dynamics, we could state the following proposition regarding the direction and the size of deviation which triggers escape from the neighborhood of the SCE.

PROPOSITION 1 *The escape from the SCE should occur in the direction given by the vector \bar{v} . The escape happens if the beliefs deviate from the SCE by more than $\delta^* \|\bar{v}\|$.*

Because the true dynamics is not one-dimensional, both results in Proposition 1 are correct only asymptotically, as the constant gain ϵ converges to zero. How good is the assumption of one-dimensionality? In Figure 4 we plot histograms of directions of escape from the ball $E = \{\gamma \mid \|\gamma - \bar{\gamma}\| \leq \delta^* \|\bar{v}\|\}$ for different values of the constant gain ϵ .¹¹ For every escape point γ^{esc} , we write $\gamma^{esc} - \bar{\gamma}$ as a vector proportional to $(\kappa, 1)$. If the escapes happen exactly along the direction \bar{v} , κ equals -6.668 . As is clear from the graph, simulated escapes are distributed symmetrically around the theoretical direction; moreover, the distribution of escape directions becomes concentrated around the mean as the value of the constant gain ϵ declines ($\text{Var}(\kappa)$ is increasing with ϵ). Therefore, considering the escape dynamics as essentially one-dimensional is a valid approximation and could be used to derive the other values of interest.

4.2 Time Until Cooperation and Comparative Statics

In this section we use the expression for δ^* to characterize the expected time until an escape to cooperation happens and investigate its dependence on the structural parameters.

DEFINITION 1 *The mean exit time from the NE is defined as a mean value of N^ϵ where $N^\epsilon = \inf\{n \mid \|\gamma_n - \bar{\gamma}\| > \|\gamma^* - \bar{\gamma}\|\}$.*

The mean exit time is the expected time until the beliefs leave the ball of radius $\|\gamma^* - \bar{\gamma}\|$ around the SCE for the first time. This ball includes the region of attraction of the SCE. After this time, a very fast transition to cooperation follows.

In addition to the mean dynamics approximation (18), it is possible to approximate the real-time dynamics under adaptive learning by the following one-dimensional continuous time diffusion process as described in Appendix B:¹²

$$d\hat{\varphi}_t = -\hat{\varphi}_t dt + \sqrt{\epsilon \lambda_{\bar{\Sigma}}} dz_t, \quad (22)$$

where $\hat{\varphi}_t = \hat{v}^T(\gamma_t - \bar{\gamma})/\|\hat{v}\|$, $\gamma_t = (\gamma_{in}^T, \gamma_{jn}^T)^T$, $\lambda_{\bar{\Sigma}} = 2b^2\sigma_2^2\bar{\lambda}_1$, z_t is a one-dimensional Wiener process.

¹¹The boundary of this ball, a sphere $\{\gamma \mid \|\gamma - \bar{\gamma}\| = \delta^* \|\bar{v}\|\}$, intersects the line Γ at a distance $\delta^* \|\bar{v}\|$ from the SCE.

¹²The diffusion process (22) is a special case of the well-known Ornstein-Uhlenbeck process.




FIGURE 4: Frequency distribution of directions of escape from the ball $E = \{\gamma \mid \|\gamma - \bar{\gamma}\| \leq \|\gamma^* - \bar{\gamma}\|\}$ and the direction of \bar{v} (dashed), $a = 2$, $b = 0.1$, $\sigma_2^2 = 0.01$.

This approximation allows us to derive an analytical expression for the expected number of real-time periods till escape, which is given in the following proposition.

PROPOSITION 2 *The expected number of periods until an escape towards collusion happens is inversely proportional to $\left(\bar{q}/\sqrt{\sigma_2^2}\right)^4$, the fourth power of the standardized Nash output, and inversely proportional to ϵ^2 , the square of the gain parameter:*

$$E[N^\epsilon] \sim \left(\frac{\bar{q}}{\sqrt{\sigma_2^2/\epsilon}}\right)^{-4}. \quad (23)$$

We find that the likelihood of cooperation (escape to cooperation is faster) is higher in industries with a higher Nash output \bar{q} and a lower volatility σ_2^2 of firm-specific shocks. This finding is rather intuitive. In industries with a higher Nash output \bar{q} , the difference between




FIGURE 5: Comparison of the analytically derived and simulated mean exit times from the NE for different specifications of the model parameters.

the competitive and cooperative profits is greater.¹³ The latter makes collusion more attractive, leading firms to quit the competitive Nash equilibrium more easily: the expected number of periods the competition could be supported becomes smaller. Since the firm-specific shocks are independent, it is obvious that the higher their volatility, the less correlated firms' actions become. With a lower firm-specific shock variance, the share of uncertainty in the data generating process directly attributed to idiosyncratic shocks decreases, therefore increasing the role of uncertainty related to belief updating which is common among firms. Increasing the share of the common factor raises the probability of firms' beliefs becoming correlated and thus the likelihood of escape from the NE; expected escape time drops.

Finally, the gain parameter ϵ measures the speed with which the firms incorporate the new information into their beliefs; lower ϵ means that the firms update their beliefs more slowly, and it takes a longer stretch of accidental correlation to lead to the firms' reaction.¹⁴ Therefore, lower ϵ makes the competitive phase in the industry last longer.

The above discussion considered an approximated value of the expected escape time that was derived using a one-dimensional approximation. Below we compare this analytical approxima-

¹³At NE, firm's average output is \bar{q} and the price $\frac{a}{3}$, giving an average profit of $\frac{a^2}{9b}$. At the CE, the average profit is $\frac{a^2}{8b}$. Therefore, the profit difference is proportional to $\frac{a^2}{9b} = \bar{q}^2 b$. (Here, for simplification, the marginal cost c is set to zero.)

¹⁴Alternatively, think about the firms which could alternatively use constant gain updating equations (17) with an infinite amount of data or run an OLS regression with T data points. In the OLS regression, the mean age of a data point is $T/2$. A constant gain learning algorithm is equivalent to a weighted Least Squares, where the weight of a point τ periods old is proportional to $(1 - \epsilon)^\tau$, giving $\frac{1-\epsilon}{\epsilon}$ as the mean age of a point when an infinite amount of data is used. Thus, $\epsilon \approx 2/T$ produces a similar average age of the data for the OLS and the constant gain RLS. Therefore, lower ϵ is equivalent to using more data in forming the beliefs by OLS or forgetting this data at a slower rate in constant gain RLS.

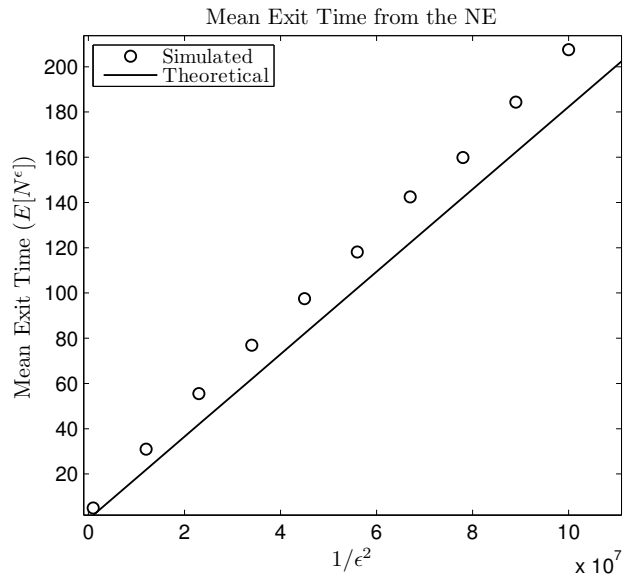


FIGURE 6: Linearity of simulated mean exit time from the NE in $1/\epsilon^2$, $a = 2$, $b = 0.1$, $\sigma_2^2 = 0.01$. It shows that the slopes of simulated and analytically derived mean exit in the space of $1/\epsilon^2$ are quite close to each other.

tion with the simulation evidence which demonstrates that the agreement is remarkably good. In Figure 5 we present the average escape time obtained by simulating the original discrete-time learning dynamics (17). One could see that the mean exit time from the Nash equilibrium is increasing in the firm-specific shock volatility σ_2^2 and decreasing in Nash output \bar{q} as well as in the constant gain parameter ϵ , exactly as predicted by the analytical expression (23). Additionally, analytically derived and simulated mean exit times are remarkably close.

To illustrate even better the validity of our analytical approximation (23), in Figure 6 we plot the expected escape time, expressed as a number of periods, against $1/\epsilon^2$. Proposition 2 states that this dependence is given by a straight line with a certain slope. The graph indicates, first, that average simulated escape times are indeed scaled as $1/\epsilon^2$. Second, the observed slope is remarkably close to the theoretically derived one. The figure unambiguously demonstrates that the escape dynamics could indeed be derived from the behavior of one-dimensional approximation with great precision.

5 Characterizing the Belief Trajectory

In this section we first characterize belief trajectories and define cooperative and price war phases. Next, we do numerical analysis of durations of these phases and their functional dependence on the model parameters.




FIGURE 7: Comparison of the mean belief trajectory and zero mean forecast error curve ($u = 0$), $a = 2$, $b = 0.1$, $\sigma_2^2 = 0.01$, $\epsilon = 0.005$.

5.1 From Competition to Cooperation and Back: Forecast Errors

We can notice from Figure 1 and Figure 2 that the beliefs are characterized by fast and slow dynamics. When the beliefs escape the Nash equilibrium, the firms quickly learn how to cooperate and rush towards the cooperative equilibrium: it can take less than ten periods from the time beliefs leave the SCE region of attraction to their arrival in the neighborhood of the CE. After the cooperative equilibrium is reached, the dynamics slows down dramatically. Below we give the reason for these fast and slow dynamics and associate them to the cooperative and price war phases defined subsequently in this section.

The belief updating procedure aims to minimize the forecast error, which implies that the mean dynamics trajectory should be near points in the belief space where $u = E[u_{in}] = 0$. The beliefs satisfying the level curve $u = 0$ are given as

$$\gamma_0 = \frac{2bc - (2a - c)\gamma_1}{2b - \gamma_1}. \quad (24)$$

In Figure 7, we plot the beliefs satisfying the above equation (24) together with the mean dynamics trajectory of beliefs derived from (19). From the figure we see that the mean forecast error stays positive during the escape until the cooperative equilibrium is reached, then its value remains for a while near zero and afterwards becomes negative when the belief trajectory crosses

the $u = 0$ level curve. Moreover, during the escape the mean forecast error is large in absolute value, which leads to fast belief revision. For the rest of the escape-cooperation-price war cycle, the forecast error is rather small, leading to very slow evolution of beliefs.

The beliefs escape from the SCE when they happen to be outside of the Nash equilibrium region of attraction. Intuitively, this may happen when a particular sequence of shocks induces both firms to contract the supply simultaneously, subsequently leading to an increase of the mean actual price $y = E[y_n]$. Since firms try to minimize the mean forecast error $u = E[u_{in}] = E[y_{in} - \hat{E}[y_{in}|y^{n-1}]]$, the increase of the mean actual price leads firms to update their beliefs accordingly and to increase the mean forecasted price by contracting the supply. The latter leads to a further increase in the mean expected price. This produces a period of upward price spiral and increasingly cooperative behavior.

After a period of cooperation the gap between the mean actual price and the mean forecasted price approaches zero and the beliefs are in a quasi-equilibrium for a while. Firms perceive this situation as an equilibrium state, as their beliefs are approximately validated by the observations. This period can be characterized as the *cooperative phase*.

However, the perception of the equilibrium state is destroyed when the mean forecast error becomes negative. At this moment the mean actual price falls below the mean forecasted price and the firms start to update their beliefs aggressively in order to reduce the mean price forecast. They achieve this by expanding supply. The expansion of supply implies further decrease of the average actual price, and the updating cycle continues to chase the data. This leads the industry into a period of downward price spiral which moves the beliefs back toward the SCE. Since at the SCE the mean forecast error is zero while during the price war it is negative, there is a point where the absolute value of u stops declining (\dot{u}/u becomes positive) and the belief system is pushed into the attraction region of the Nash equilibrium.

The periods of increasing and decreasing price spirals can be characterized as cooperative and the price war phases respectively. To be more specific, we define the *cooperative phase* as the period of increasing price spiral together with the time spent in the neighborhood of the cooperative equilibrium and the *price war phase* as the period of decreasing price spiral till the time when the belief system appears in the attraction region of the NE. Here, we can find some similarities to the Green and Porter model where it is the unanticipated demand shocks that cause price wars when the price falls below the predefined trigger price. In our context the mean forecast price serves as the trigger price, and the unanticipated fall of the mean actual price

below the mean forecasted price causes the price war. We thus find that the demand shock is not necessarily the sole source causing price wars, and we could generate the wars even without a demand shock.

5.2 Numerical Analysis of Durations

Our model allows us to investigate the dependence of the durations of cooperative and price war phases on the model parameters. We use mean dynamics ODE (19) to calculate belief thresholds that trigger respectively a price war and reversion to the NE. Therefore, for convenience durations from the real-time dynamical system (17) are expressed in continuous time. For the purposes of duration calculations we adopt the following definitions.

DEFINITION 2 *The duration of the cooperation is defined as $\tau^c = t_2 - t_1$ where $t_1 = \inf\{t \mid \|\gamma_t - \bar{\gamma}\| > \|\gamma^* - \bar{\gamma}\|\}$ and $t_2 = \inf\{t \mid u_t < 0\}$.*

The Cooperative phase is defined as the period between the time when beliefs escape the NE for the first time and the time the mean forecast error turns negative for the first time, triggering a price war.

DEFINITION 3 *The duration of the price war is defined as $\tau^{pw} = t_3 - t_2$ where $t_2 = \inf\{t \mid u_t < 0\}$ and $t_3 = \inf\{t \mid u_t = \min u_s\}$.*

The price war is a period when the forecast error is negative and decreasing, forcing the firms to increase output so that the actual price decreases, following the forecasts. The start of the price war phase coincides with the end of cooperation.

We turn now to the numerical results. We depict the duration of the cooperation and the price war phases in Figure 8. Both durations are increasing in the Nash equilibrium output \bar{q} and decreasing in volatility σ_2^2 of the firm-specific shock. A higher \bar{q} increases the losses that firms incur when the cooperation breaks down, as the profit difference between cooperative and Nash equilibria is proportional to \bar{q}^2 . Therefore, in industries with a higher \bar{q} , it is harder to break down a cartel. In industries with a higher σ_2^2 , the firms' actions are less correlated. The latter together with profit maximization implies a higher output and increases the likelihood of initiating the price war. Moreover, in industries with higher firm-specific shock volatility σ_2^2 , the cooperation could only be achieved for sufficiently high levels of the Nash equilibrium output \bar{q} . This is the other side of the difficulty of learning to cooperate as evidenced by mean expected exit time discussed above.

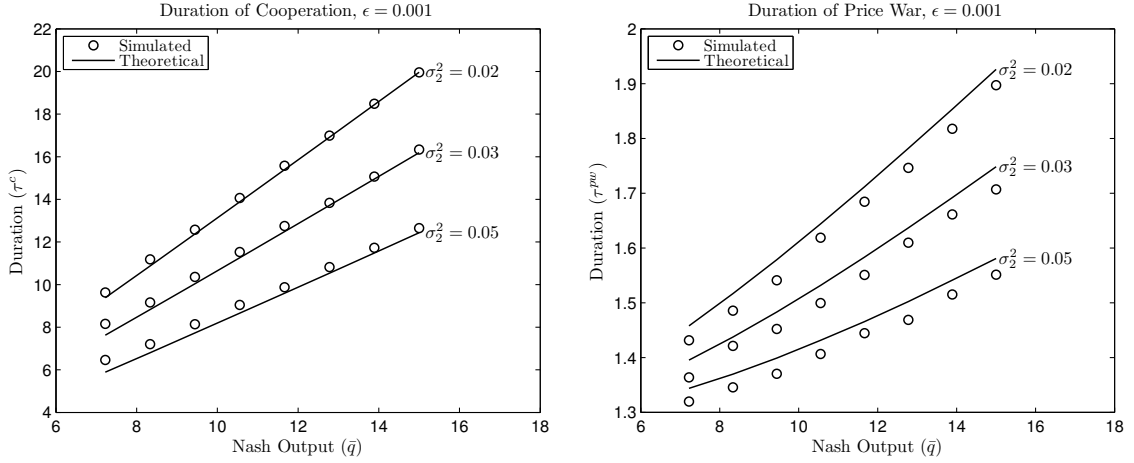


FIGURE 8: Durations of cooperation and price war phases as a function of Nash output \bar{q} for a different values of the firm-specific shock volatility σ_2^2




FIGURE 9: Durations of cooperation and price war phases as a function of normalized Nash output, $\bar{q}/\sqrt{\sigma_2^2}$.

In Figure 9 and Figure 10, we depict the same curves as in Figure 8 and Figure 5, respectively, but with normalized Nash output, $\bar{q}^N = \bar{q}/\sqrt{\sigma_2^2}$, on the horizontal axes. We see that the durations turn out to be a function of a single combination of parameters, $\bar{q}/\sqrt{\sigma_2^2}$. The durations of cooperative and price war phases, compared to the mean exit time from the SCE, do not depend on constant gain parameter ϵ since after beliefs escape the SCE, the model dynamics are mainly governed by the mean dynamics. From (23), we see that mean exit time depends on the expression $\bar{q}/\sqrt{\sigma_2^2/\epsilon} = \bar{q}^N \sqrt{\epsilon}$. To check this property we depict all the simulation results in both panels of Figure 5 in Figure 10 with $\bar{q}^N \sqrt{\epsilon}$ on the horizontal axis. Indeed, we see that all the average escape times from real-time dynamic simulations in Figure 5 are well approximated by the inverse of the 4th power of $\bar{q}^N \sqrt{\epsilon}$ (solid line).




FIGURE 10: Mean exit time from the NE as a function of $\bar{q}/\sqrt{\sigma_2^2}/\epsilon$.

The above findings are summarized in the following result. The result contains testable hypotheses about the dependence of durations of cooperative and price war phases on the structural parameters of the model.

PROPOSITION 3 *Durations of cooperative and price war phases are increasing functions of the normalized Nash output, $\bar{q}^N = \bar{q}/\sqrt{\sigma_2^2}$. These durations are thus increasing with the industry output \bar{q} and decreasing with the standard deviation of the idiosyncratic shock.*

The Duration of the competitive phase is inversely proportional to the 4th power of $\bar{q}^N\sqrt{\epsilon}$, thus decreasing with the industry output \bar{q} and increasing with the standard deviation of the idiosyncratic shock. The faster is the speed with which the firms update their beliefs, the lower is the industry ability to maintain a competitive regime.¹⁵

As shown in Appendix B, the largest eigenvalue of \bar{P} is approximately equal to $\bar{q}^2/\sigma_2^2 = (\bar{q}^N)^2$. Due to the approximate one-dimensionality of the model, this eigenvalue is the most important for both the real-time dynamics (17) and the mean dynamics (18). Therefore, mainly the evolution of our model is a function of two parameters, normalized Nash output and (for expected escape to cooperation time) the constant gain, which allows a very parsimonious description of the model dynamics.

¹⁵Sannikov and Skrzypacz (2007) show when information arrives continuously, under the same settings as in Green and Porter (1984), collusion is impossible. However, under speed of updating we do not mean the frequency of updating but the extent to which new information is incorporated into belief updating. In our model this is measured by ϵ , the constant gain parameter.

6 Conclusion

In the paper, we study the imperfect monitoring model of duopolistic industry in the sense of Green and Porter (1984) but with unknown structural parameters. Firms build their perceptions about the economy, form beliefs about unknown parameters, construct an estimation model and update their beliefs when new price signals are revealed by the market. We show that when firms take into account the possibility of structural changes in the estimation of parameters, the beliefs of firms may become correlated and this may generate fast escapes from the Nash equilibrium (NE) towards a collusive equilibrium (CE) and then a slow return back to the Nash equilibrium. We show that the movement towards the collusive equilibrium constitutes collusive behavior while movement back to the Nash constitutes a *price war*. The fast and slow dynamics are explained by the behavior of the mean forecast error.

Model predictions leads to testable hypotheses about the dependence of the likelihood of reversion to cooperation, and durations of the cooperative and price war phases on the structural parameters. We expect that industries with the same normalized output have the same durations of cooperative and price war phases as well as mean exit time from the NE. Moreover, industries with a higher normalized output possess longer durations of cooperative and price war phases but shorter mean exit times from the NE. Therefore, normalized output can serve as a measure to rank industries with respect to durations of cooperative and price war phases as well as mean exit times from the NE. In addition, we find that it is easier to form a cartel in industries where firms are induced to analyze more common information.

Our results show that a cartel can break down not only due to exogenous factors as in Green and Porter (1984) or Slade (1989) but also due to the endogenous dynamics involved with the uncertainty introduced through the unknown structural parameters. Finally, we show that firms do not have to *a priori* commit to some trigger strategies, and even though playing Nash equilibrium strategies all the time, firms can implicitly coordinate on the length of the duration of the punishment period and the price trigger that triggers cooperation when they update beliefs about the unknown parameters of the inverse demand function.

References

- Abreu, D., D. Pearce, and E. Stacchetti**, “Optimal cartel equilibria with imperfect monitoring,” *Journal of Economic Theory*, 1986, *39* (1), 251 – 269.
- Balvers, R. and T. Cosimano**, “Actively Learning about Demand and the Dynamics of Price Adjustment,” *Economic Journal*, September 1990, *100* (402), 882–98.
- Chakrabarti, S.**, “Collusive equilibrium in Cournot oligopolies with unknown costs,” *International Economic Review*, 2010, *51* (4), 1209–1238.
- Cho, I., N. Williams, and T. Sargent**, “Escaping Nash inflation,” *Review of Economic Studies*, 2002, *69* (1), 1–40.
- Ellison, G.**, “Theories of cartel stability and the joint executive committee,” *The Rand journal of economics*, 1994, *25* (1), 37–57.
- Ellison, M. and A. Scott**, “Learning and Price Volatility in Duopoly Models of Resource Depletion,” CEPR Discussion Papers 7378, C.E.P.R. Discussion Papers July 2009.
- Evans, G. and S. Honkapohja**, *Learning and Expectations in Macroeconomics*, Princeton, Princeton University Press, Princeton, NJ: Princeton University Press, 2001.
- , – , and **N. Williams**, “Generalized stochastic gradient learning,” *International Economic Review*, 2010, *51* (1), 237–262.
- Friedman, J.**, “A Non-cooperative Equilibrium for Supergames,” *Review of Economic Studies*, 1971, *38* (113), 1–12.
- , *Market Structure and Competition Policy*, Cambridge University Press,
- Fudenberg, D. and D. K. Levine**, “Self-Confirming Equilibrium,” *Econometrica*, 1993, *61* (3), 523–545.
- and – , “Self-confirming equilibrium and the Lucas critique,” *Journal of Economic Theory*, November 2009, *144* (6), 2354–2371.
- and **J. Tirole**, “A Signal-Jamming Theory of Predation,” *RAND Journal of Economics*, 1986, *17*, 366–376.

- Green, E. and R. Porter**, “Noncooperative collusion under imperfect price information,” *Econometrica: Journal of the Econometric Society*, 1984, 52 (1), 87–100.
- Grossman, Sanford J, Richard E Kihlstrom, and Leonard J Mirman**, “A Bayesian Approach to the Production of Information and Learning by Doing,” *Review of Economic Studies*, October 1977, 44 (3), 533–47.
- Kaneko, M.**, “Some remarks on the folk theorem in game theory,” *Mathematical Social Sciences*, October 1982, 3 (3), 281–290.
- Karatzas, I. and S. E. Sherve**, *Brownian Motion and Stochastic Calculus*, Springer-Verlag, 1991.
- Kasa, Kenneth**, “Learning, Large Deviations, And Recurrent Currency Crises,” *International Economic Review*, 2004, 45 (1), 141–173.
- Keller, G. and S. Rady**, “Optimal Experimentation in a Changing Environment,” *The Review of Economic Studies*, 1999, 66 (3), 475–507.
- Kolyuzhnov, D., A. Bogomolova, and S. Slobodyan**, “Escape Dynamics: A Continuous Time Approach,” 2006.
- Kreps, D.**, “Anticipated utility and dynamic choice,” *Frontiers of Research in Economic Theory*, 1998, pp. 1983–1977.
- Mirman, L., L. Samuelson, and A. Urbano**, “Duopoly signal jamming,” *Economic Theory*, 1993, 3, 129–149.
- , – , and – , “Monopoly Experimentation,” *International Economic Review*, August 1993, 34 (3), 549–63.
- , – , and **E. Schlee**, “Strategic Information Manipulation in Duopolies,” *Journal of Economics*, 1994, 62, 363–384.
- Riordan, M.**, “Imperfect Information and Dynamic Conjectural Variations,” *RAND Journal of Economics*, 1985, 16, 41–50.
- Rothschild, M.**, “A Two-armed Bandit Theory of Market pricing,” *Journal of Economic Theory*, 1974, 9, 185–202.

- Rustichini, A. and A. Wolinsky**, “Learning about Variable Demand in the Long Run,” *Journal of Economic Dynamics and Control*, 1995, *19*, 1283–1292.
- Sakai, Y.**, “The value of information in a simple duopoly model,” *Journal of Economic Theory*, June 1985, *36* (1), 36–54.
- Sannikov, Y. and A. Skrzypacz**, “Impossibility of Collusion under Imperfect Monitoring with Flexible Production,” *The American Economic Review*, 2007, *97*, 1794–1823.
- Sargent, T.**, *The Conquest of American Inflation*, Princeton: Princeton University Press, 1999.
- **and S. Williams**, “Impacts of Priors on Convergence and Escapes from Nash Inflation,” *American Economic Review*, 2005, *8*, 360–391.
- , **N. Williams, and T. Zha**, “The Conquest of South American Inflation,” *Journal of Political Economy*, 2009, *117* (2), 211–256.
- Slade, M.**, “Strategic pricing models and interpretation of price-war data,” *European Economic Review*, 1990, *34* (2-3), 524 – 537.
- Slade, Margaret E.**, “Price Wars in Price-Setting Supergames,” *Economica*, 1989, *56* (223), 295–310.
- Stigler, G.**, “A theory of oligopoly,” *The Journal of Political Economy*, 1964, *72* (1).
- Wieland, Volker**, “Monetary Policy, Parameter Uncertainty and Optimal Learning,” *Journal of Monetary Economics*, 2000, *46*, 199–228.
- Williams, N.**, “Escape dynamics in learning models.” PhD dissertation, University of Chicago 2001.
- , “Adaptive learning and business cycles,” 2003. Princeton University, Manuscript.

A E-stability of SCE

The SCE is E-stable when the Jacobian of the belief system (18) has all eigenvalues with negative real parts. Define $\theta = (\gamma'_i, \gamma'_j, \text{vec}(P_i)', \text{vec}(P_j)')$, then the Jacobian of the system (18) takes the following form:

$$J = \begin{bmatrix} \frac{\bar{P}_i \partial \bar{g}_i(\bar{\gamma})}{\partial \gamma_i} & \frac{\bar{P}_i \partial \bar{g}_i(\bar{\gamma})}{\partial \gamma_j} & 0 & 0 \\ \frac{\bar{P}_j \partial \bar{g}_j(\bar{\gamma})}{\partial \gamma_i} & \frac{\bar{P}_j \partial \bar{g}_j(\bar{\gamma})}{\partial \gamma_j} & 0 & 0 \\ -\text{vec} \left(\frac{\bar{P}_i \partial M(\bar{\gamma}_i) \bar{P}_i}{\partial \gamma_i} \right) & 0 & -\text{vec} \left(\frac{\partial \bar{P}_i M(\bar{\gamma}_i) \bar{P}_i}{\partial P_i} \right) & 0 \\ 0 & -\text{vec} \left(\frac{\bar{P}_j \partial M(\bar{\gamma}_j) \bar{P}_j}{\partial \gamma_j} \right) & 0 & -\text{vec} \left(\frac{\partial \bar{P}_j M(\bar{\gamma}_j) \bar{P}_j}{\partial P_j} \right) \end{bmatrix}.$$

Since PMP is a quadratic matrix in P , the sufficient condition for E-stability is that the following matrix has eigenvalues with negative real parts:

$$\bar{h}_\gamma = \begin{bmatrix} \frac{\bar{P}_i \partial \bar{g}_i(\bar{\gamma})}{\partial \gamma_i} & \frac{\bar{P}_i \partial \bar{g}_i(\bar{\gamma})}{\partial \gamma_j} \\ \frac{\bar{P}_j \partial \bar{g}_j(\bar{\gamma})}{\partial \gamma_i} & \frac{\bar{P}_j \partial \bar{g}_j(\bar{\gamma})}{\partial \gamma_j} \end{bmatrix}. \quad (25)$$

It is straightforward to show that the matrix is of the form

$$\bar{h}_\gamma = \begin{bmatrix} -1 & 0 & -1/2 & -\bar{q} \\ 0 & -1 & 0 & 0 \\ -1/2 & -\bar{q} & -1 & 0 \\ 0 & 0 & 0 & -1 \end{bmatrix}. \quad (26)$$

All the eigenvalues of the matrix (26) (-1.5, -1, -1, -0.5) have negative real parts. Thus, the SCE is E-Stable.

B Derivation of the Diffusion Equation for the Belief System

In addition to the mean dynamics approximation (18), it is possible to approximate the real-time dynamics under adaptive learning by a continuous time diffusion process as follows. Stacking together the system of stochastic difference equations (17) for both firms i and j , we can define the following system:

$$\theta_{n+1} = \theta_n + \epsilon \mathcal{H}(\theta_n, \zeta_n), \quad (27)$$

where $\theta_n = (\gamma_{in}^T, \gamma_{jn}^T, \text{vech}(P_{in})^T, \text{vech}(P_{jn})^T)^T$ is a vector in which we stack both firms' beliefs and current estimates of the variance-covariance matrices, and $\zeta_n = (x_{in}^T, x_{jn}^T, \omega_{in}, \omega_{jn})^T$ is the state vector.¹⁶

Evans and Honkapohja (2001) shows that for $\epsilon \rightarrow 0$, the time path of the system can be approximated by the following diffusion:

$$d\theta_t = h'_\theta(\theta(t, \theta_0))(\theta_t - \theta(t, \theta_0))dt + \sqrt{\epsilon}\Sigma(\theta(t, \theta_0))^{1/2}d\hat{W}_t, \quad (28)$$

where $\theta(t, \theta_0)$ is the solution of the mean dynamics ODE with the initial condition θ_0 , $h'_\theta(\cdot)$ is the Jacobian of $E[\mathcal{H}(\theta_n, \zeta_n)]$, \hat{W}_t is a multi-dimensional Wiener process with $\dim(\hat{W}_t) = \dim(\theta)$ and $\Sigma = \text{cov}[\mathcal{H}(\theta_n, \zeta_n)]$.

Since the SCE is E-stable (see Appendix A), following Evans and Honkapohja (2001) it suffices to analyze the time path of (28) around the SCE $\bar{\theta}$ which is a stable limit point of the mean dynamics. In addition, one can neglect the dynamics of elements in P_{in} and P_{jn} since they tend to change much more slowly than the beliefs γ and remain almost unchanged before the escape. Considering only the dynamics in beliefs γ , the approximating diffusion becomes¹⁷

$$d\varphi_t = \bar{h}_\gamma\varphi_t dt + \sqrt{\epsilon}\bar{\Sigma}^{1/2}dW_t, \quad (29)$$

where $\varphi_t = \gamma_t - \bar{\gamma}$, $\gamma_t = (\gamma_{in}^T, \gamma_{jn}^T)^T$, W_t is a multi-dimensional Wiener process with $\dim(W_t) = \dim(\varphi_t)$, $\bar{h}_\gamma = h'_\gamma(\bar{\theta})$, and $\bar{\Sigma} = \text{cov}[\mathcal{H}_\gamma(\bar{\theta}, \zeta_n)]$.

The diffusion (29) is four-dimensional. Following Kolyuzhnov et al. (2006), we can transform it into a one-dimensional diffusion. By pre-multiplying the above equation by v^T and denoting $\hat{\varphi}_t = v^T\varphi_t$, we obtain

$$d\hat{\varphi}_t = -\hat{\varphi}_t dt + \sqrt{\epsilon\lambda_{\bar{\Sigma}}}dz_t, \quad (30)$$

where $\lambda_{\bar{\Sigma}}$ and $v = \hat{v}/\|\hat{v}\|$ are respectively the dominant eigenvalue and the standardized dominant eigenvector of $\bar{\Sigma}$, $\hat{v} = (\bar{v}^T, \bar{v}^T)^T$; z_t is a one-dimensional Wiener process. A final transformation involves ignoring the term $-\hat{\varphi}_t$ and setting it to zero; thus, the original real-time learning dynamics is approximated by a one-dimensional Brownian motion.

This approximation is very useful because a very simple formula allows us to derive the

¹⁶Operator *vech* maps lower triangular part of a symmetric $n \times n$ matrix into a vector by fully stacking first column, then the second column from the second element down, and so on till a single (n, n) element of a matrix is added to the vector.

¹⁷See also Sargent and Williams (2005), pp. 376-377.

time until the one-dimensional process (30) leaves any interval of the real line. For the interval $[-d; +d]$, assuming that $z(0) = 0$, the expected time is given by the following formula: $E[\tau^\epsilon] = \frac{d^2}{\epsilon \lambda_{\bar{P}}}$ (see Karatzas and Sherve (1991, Eq. 5.62, p. 345)). The boundary of the interval d equals $\hat{\varphi}^* = \delta^* v^T \hat{v} = \sqrt{2} \delta^* \|\bar{v}\|$. Finally, notice that to translate continuous time units, in which τ^ϵ is expressed, into the number of periods, we need to divide τ^ϵ by ϵ .

Taking all of the above into account, the expected number of periods until an escape towards cooperation happens is given by the following formula:

$$\begin{aligned} E[N^\epsilon] &= \frac{\hat{\varphi}^{*2}}{\lambda_{\bar{P}}} \frac{1}{\epsilon^2} = \left(\frac{\delta^*}{b}\right)^2 \frac{1}{\epsilon^2} \\ &= 36 \frac{\sigma_2^4}{\bar{q}^4} \frac{1}{\epsilon^2}. \end{aligned}$$

The formula uses the value of $\lambda_{\bar{P}} = 2b^2 \sigma_2^2 \bar{\lambda}_1$, where $\bar{\lambda}_1 \approx \|\bar{v}\|^2 / \sigma_2^2$ is the dominant eigenvalue of \bar{P} .

Thus, the final dependence of $E[N^\epsilon]$ on the model parameters is given as

$$E[N^\epsilon] \sim \left(\frac{\bar{q}}{\sqrt{\sigma_2^2/\epsilon}}\right)^{-4}. \quad (31)$$

The inverse dependance of the mean exit time on $\bar{q}/\sqrt{\sigma_2^2}$ can be justified by the following observation:

$$\frac{\bar{q}}{\sqrt{\sigma_2^2}} = \sqrt{(\bar{\lambda}_1 - 1)(1 - \bar{\lambda}_2)}, \quad (32)$$

where $\bar{\lambda}_1$ and $\bar{\lambda}_2$ are respectively the largest and smallest eigenvalues of \bar{P} . Indeed, from (32) we see that $\bar{\lambda}_1 \geq 1$ and $0 \leq \bar{\lambda}_2 < 1$.

We can notice that an increase of $\bar{q}/\sqrt{\sigma_2^2}$ implies an increase of the ratio $\bar{\lambda}_1/\bar{\lambda}_2$. The latter implies acceleration of belief generation along the direction of the dominant vector of \bar{P} and therefore, on average, less time is needed to exit the attraction region. This explains why mean exit time depends inversely on $\bar{q}/\sqrt{\sigma_2^2}$.