Quantity Competition with Production Commitment: Theory and Evidence from the Auto Industry*

Guillermo Caruana
CEMFI
caruana@cemfi.es

Liran Einav
Stanford University
leinav@stanford.edu

June 30, 2004

Preliminary draft. Comments are very welcome.

Abstract

Economists frequently use simultaneous one-shot games to model strategic interactions. Quite often, however, the actions taken in such games require some preparation. If such preparations are costly to alter, strategic incentives may be an important consideration during this preparation stage, and may ultimately affect the equilibrium outcome. This paper analyzes a finite-horizon linear-quadratic differential game of quantity competition, where players may continuously adjust the actions they plan to take in the eventual interaction. Parties incur adjustment costs when changing previous plans. This allows them to use such production plans as a partial commitment device. We show that the equilibrium outcome is more competitive than its static analog. Moreover, depending on the initial production plans, the equilibrium path may exhibit one of two patterns: If initial plans are sufficiently high players continuously adjust their production plans downwards. In contrast, if the intended production is lower the equilibrium path exhibits a non-monotonic pattern. Players, involved in a Stackelberg warfare, start by increasing their intended production, and only later adjust it downwards. We use data about production plans of auto manufacturers in the U.S. to assess the predictions of our model. The data exhibit qualitatively similar patterns to the theoretical predictions.

Keywords: Differential games, Adjustment costs, Cournot, Quantity competition.

JEL classification: C72, C73, D43, L13, L62.

*We would like to thank Tim Bresnahan, Uli Doraszelski, Ken Judd, and especially Esteban Rossi-Hansberg for helpful conversations. We are extremely thankful to Maura Doyle and Chris Snyder for so kindly sharing their data with us. Address Correspondence to Liran Einav, Department of Economics, Stanford University, Stanford, CA 94305-6072; Tel: (650) 723-3704.

1 Introduction

Economists often model strategic interactions using simultaneous one-shot games. It is as if decisions were taken in the blink of an eye and realized instantaneously. This is, of course, a simplification. Complex decisions, such as those of entry, exit, or production are normally the result of a long preparation process towards the final decisions. If changing plans during this process is costly, incentives to behave strategically during the preparation stage should be explicitly considered, as they may be an important determinant of the final outcomes.

Consider, for example, the specific application of this paper. Auto manufacturers compete in producing cars.[1] Suppose that, ahead of time, a firm has a planned production level. In order to achieve its production target, it needs to take certain measures, such as hiring labor, canceling vacations, purchasing parts from suppliers, etc. If the firm then decided to change its desired production level, it would likely need to incur some costs adjusting the previous measures. To the extent that such preparation measures are either publicly observed or cannot perfectly be hidden from competitors, they can play a strategic role. Given the costly nature of these adjustments, the preparation stage acts a gradual commitment device. Firms realize that their planned production levels affect their rivals' production plans, and use this to their advantage, adjusting their own intentions strategically.

In our previous work (Caruana and Einav, 2004), we study the scope for endogenous commitment that may arise from the existence of these adjustment costs, and focus on discrete problems, such as those of entry and exit. It is hard to extend that framework for the analysis of games with continuous action spaces, and more importantly, with adjustment costs that depend on the magnitude of the adjustment. Therefore, in this paper we build a different model using differential games techniques. In particular, we use a linear-quadratic structure, which allows for a simple analysis without sacrificing the basic insights.[2]

Section 2 of the paper presents the basic model. At some specified date in the future two symmetric firms will engage in a Cournot competition. At date zero, each firm inherits a production structure that will lead to a certain production level. From that point on, the firm can make continuous adjustments to its production structure, but incurs quadratic costs every time it does so. When the initial production targets are not too high, both firms begin by gradually increasing their production plans. Firms use these intended plans as a commitment device; they want to commit to high production levels in order to obtain a Stackelberg leadership position in the industry. In equilibrium, however, both firms are provided with similar commitment opportunities, and thereby engage in a "Stackelberg warfare," each trying not to become a Stackelberg

[1] For simplicity, the arguments above abstract from other important factors, such as inventories or uncertainty. These are discussed later in the paper.

[2] While we believe that the qualitative nature of the results extend to other functional forms of costs and profits, solving the full dynamic model for such cases is in general not tractable. In this sense, linear-quadratic models are perceived as second-order approximations to more complicated models.

follower. As the horizon gets closer, however, both firms become sufficiently committed to producing high quantities. Thus, at a certain point before the final date, the (dynamic) commitment effect becomes less important, while the (static) incentive to best respond to the opponent's high production target increases and becomes dominant. Therefore, from that point on both firms start to gradually decrease their production plans in the direction of their static best-response levels. The eventual equilibrium outcome still remains more competitive than its static analog. The non-monotonic equilibrium path described above is a consequence of initial production targets being not too high. In the end of Section 2 we discuss different possibilities that can rationalize such initial positions. Nevertheless, the result that before the final date there would be some period in which firms gradually decrease their intended production is fully general; it does not depend on the initial production plans.

Section 3 extends the basic model along several dimensions. We investigate the case of asymmetric duopoly, an extension to more than two symmetric players, and time-varying adjustment costs, and discuss how the model would extend when the production game is repeated. We also analyze a case in which there is uncertainty (common to all firms) about demand or cost conditions. Uncertainties create a trade-off between the incentive to commit and the incentive to remain flexible and wait until the uncertainty is realized. The qualitative nature of the basic results remain unchanged, and the comparative statics work in the same expected directions.

Our model is related to other models of dynamic quantity competition analyzed in the literature.[3] Their focus, however, is on the stationary equilibrium of an infinite-horizon model (or on the limit of a finite-horizon one, as the horizon tends to infinity). They typically find that when actions are strategic substitutes, as in the quadratic Cournot competition we analyze, the stationary equilibrium is more competitive than its static analog, as players engage in a "Stackelberg warfare." Our finding that the eventual outcome is more competitive than its static analog has a similar spirit. Unlike this literature, however, our main focus lies on the non-stationary equilibrium path of the game. Another difference is that they consider a continuous flow of production payoffs as time goes by (production takes place at every instant), while in this paper the dynamics deal with the preparations for a single production period. Therefore, production payoffs are collected only in the end. These two ingredients result in an equilibrium path displaying a strong non-stationary pattern.

One advantage in studying non-stationary models is the qualitative testable implications they provide. It is hard to empirically test stationary dynamic models, as the static benchmark is typically not available (for example, marginal costs are typically not observed). In contrast, our non-stationary model provides a qualitative set of predictions, which can be testable even in the absence of information on, say, marginal costs.

Section 4 presents empirical evidence which is consistent with our model. We use the data collected by Doyle and Snyder (1999) about production plans made by U.S. auto manufacturers,

[3] See, among others, Cyert and DeGroot (1970), Hanig (1986), Maskin and Tirole (1987), Fershtman and Kamien (1987), Reynolds (1987, 1991), Lapham and Ware (1994), and Jun and Vives (2004).

as published in an industry newsletter. The focus of their study is quite different from ours. They are interested in the use of these production plans as a device for information sharing among manufacturers, and analyze the correlation among one manufacturer's production plan and its opponents' subsequent responses. In contrast, our focus lies on the evolution of the production targets over time.

The data contain multiple observations about future monthly production plans of the main auto manufacturers for the period 1965-1995. Starting about six months before production, the newsletter periodically reports its estimates of the production targets (measured as number of cars, at the manufacturer level) in a given future month by each of the biggest four auto manufacturers in the U.S. For each future production month, there are one to twelve (4.35 on average) different pre-production estimates. The main idea underlying the model in this paper is that actions taken in the preparation stage have a commitment value, and that this value stems from the costly nature of changing them. Thus, in order to link the data to the model, it is necessary to assume that the newsletter's estimates are based on real pre-production decisions being taken. This is an important assumption underlying our empirical analysis.

By pooling data from different production months, we estimate non-parametrically the relationship between production plans and eventual production levels, and how this relationship evolves over time. Across all manufacturers, the pattern of production targets consistently displays the non-monotonic trend predicted by the model. On average, targets start about 5% above eventual production levels. They gradually increase until they peak at 10% about 80 days before production. Then, they start a sharp decline until they reach actual production levels. When studied separately, different manufacturers exhibit different patterns. While Ford and Chrysler exhibit the strong non-monotonic pattern described above, General Motors and American Motors start much above their eventual production levels, and mainly exhibit a monotone decreasing pattern. Both of these patterns are consistent with the theoretical model.

At this point, the evidence presented is more suggestive than conclusive. There are other important dimensions of the industry, particularly the existence of inventories, which have not been considered in the study. We should note, however, that it seems hard to come with alternative models that could rationalize the consistent non-monotonic pattern observed in the data. Therefore, we are encouraged to think that the empirical patterns we report point to the strategic role that pre-production preparations may be playing in determining production decisions.

2 Basic model

2.1 Setup and solution

There are two players. Given initial production plans (q_i^0, q_j^0) at time $t = 0$, each player can continuously control the rate at which she changes her production target. Namely, if player i chooses a rate of x_i at time t then $q_i'(t) = x_i$. Note that x_i can be either positive or negative. By choosing x_i player i pays adjustment costs $c_i(x_i, t)$. At time $T > 0$ players engage in a Cournot competition and collect final payoffs of $\pi_i(q_i(T), q_j(T))$.

We will assume throughout a linear-quadratic structure. Thus, we assume that inverse demand is linear, given by $p = a - bQ$ and marginal costs are constant and given by c. Thus, we have that

$$\pi_i(q_i(T), q_j(T)) = q_i(T)(a - bq_i(T) - bq_j(T)) - cq_i(T) = (a - c)q_i(T) - bq_i^2(T) - bq_i(T)q_j(T) \quad (1)$$

We assume that adjustment costs are quadratic and take the form of

$$c_i(x_i, t) = \frac{\theta}{2}x_i^2 \quad (2)$$

Thus, adjustments costs are constant over time, symmetric across players, and symmetric for positive and negative rates. None of these properties is particularly important.

We solve for the Markov perfect equilibrium of the model. Thus, equilibrium strategies only depend on the state variables, q_i and q_j. Let $V_t^i(q_i, q_j)$ be the value function for player i at time t, with state variables q_i and q_j. Assume also that $V_t^i(q_i, q_j)$ exists and is continuous and continuously differentiable in its arguments. The value function must satisfy

$$V_t^i(q_i, q_j) = \max_{x_i}\left(-\frac{\theta}{2}x_i^2 + \frac{\partial V_t^i}{\partial q_i}x_i + \frac{\partial V_t^i}{\partial q_j}x_j + \frac{\partial V_t^i}{\partial t} + V_t^i(q_i, q_j)\right) \quad (3)$$

The first order condition for x_i is therefore

$$-\theta x_i + \frac{\partial V_t^i}{\partial q_i} = 0 \Longrightarrow x_i = \frac{1}{\theta}\frac{\partial V_t^i}{\partial q_i} \quad (4)$$

We can now substitute this back into equation (3), as well as the symmetric solution for x_j, rearrange, and obtain the following differential equation

$$0 = \frac{1}{2\theta}\left(\frac{\partial V_t^i}{\partial q_i}\right)^2 + \frac{1}{\theta}\left(\frac{\partial V_t^i}{\partial q_j}\right)\left(\frac{\partial V_t^j}{\partial q_j}\right) + \frac{\partial V_t^i}{\partial t} \quad (5)$$

The linear-quadratic structure is attractive. It is known that in this case, if one restricts the strategies to be analytic functions of the state variables, there exists a unique equilibrium of the game, which is also the limit of its discrete-time analog. Moreover, in such a case the unique value function is a quadratic function of the state variables.[4] Note that due to the inherent non-stationarity of the model, the parameters of this quadratic equation will depend on t in an unspecified way. We can express the value function as

$$V_t^i(q_i, q_j) = A_t + B_t q_i + C_t q_j + D_t q_i^2 + E_t q_j^2 + F_t q_i q_j \quad (6)$$

which also implies that

$$x_t^i(q_i, q_j) = \frac{1}{\theta}\frac{\partial V_t^i}{\partial q_i} = \frac{1}{\theta}\left(B_t + 2D_t q_i + F_t q_j\right) \quad (7)$$

[4] See Kydland (1975), who shows uniqueness for a discrete-time version, and Papavassilopoulos and Cruz (1979) and Papavassilopoulos and Olsder (1984) for analysis of existence and uniqueness in finite-horizon linear-quadratic differential games.

Substituting equations (6) and (7) into equation (5) gives us

$$0 = \frac{1}{2\theta}\left(B_t + 2D_t q_i + F_t q_j\right)^2 + \frac{1}{\theta}\left(C_t + 2E_t q_j + F_t q_i\right)\left(B_t + 2D_t q_j + F_t q_i\right) + \\ + \left(A_t' + B_t' q_i + C_t' q_j + D_t' q_i^2 + E_t' q_j^2 + F_t' q_i q_j\right) \quad (8)$$

This is a polynomial of q_i and q_j. Thus, all its coefficients have to be equal to zero. This gives the following set of six ODE's. To ease notation, we can just think about time going backwards, so all the derivatives with respect to time (A', B', etc.) reverse signs. This is convenient as our boundary condition is for $t = T$. The law of motion for the parameters is then given by

$$\begin{pmatrix} A' \\ B' \\ C' \\ D' \\ E' \\ F' \end{pmatrix} = \frac{1}{\theta}\begin{pmatrix} \frac{1}{2}B^2 + BC \\ 2BD + BF + CF \\ BF + 2BE + 2CD \\ 2D^2 + F^2 \\ \frac{1}{2}F^2 + 4DE \\ 4DF + 2EF \end{pmatrix} \quad (9)$$

The boundary condition (for $t = T$) is given by the profit function in equation (1), which implies

$$\begin{pmatrix} A_T \\ B_T \\ C_T \\ D_T \\ E_T \\ F_T \end{pmatrix} = \begin{pmatrix} 0 \\ a - c \\ 0 \\ -b \\ 0 \\ -b \end{pmatrix} \quad (10)$$

2.2 Illustration

The system of ordinary differential equations given by equation (9), with its boundary condition, defines the solution. It defines the value function at any point in time, which in turn allows us to compute the equilibrium strategies using equation (7). Unfortunately, the system does not have a closed-form solution, so we illustrate the results by approximating the equilibrium through the solution of the discrete-time analog of the game for very small time intervals.

Throughout this section, unless otherwise specified, we set $a = b = 1$, $c = 0$, $\theta = 1$, and $T = 10$. This implies that marginal costs are zero and that inverse demand is given by $p(Q) = 1 - Q$. Adjustment costs are $c_i(x_i, t) = \frac{1}{2}x_i^2$.[5] For later comparison, it is useful to observe that, for this choice of parameters, the static Nash equilibrium of the Cournot game involves each player producing $q = \frac{1}{3}$, while the Stackelberg leader and follower production levels are $q = \frac{1}{2}$ and $q = \frac{1}{4}$, respectively.

[5] One should note that some of these restriction are not important. The effect of a and c only enters through their difference $a - c$, so setting $c = 0$ is only a normalization. Similarly, optimal strategies are invariant to monotone transformations of the objective function, so, for example, setting $b = 1$ is a normalization.

Figure 1 presents the way the parameters of the (symmetric) value function for this game evolve over time. It shows the equilibrium values for the parameters of the value equation, as given in equation (6). One can see that as the horizon becomes longer all parameters, except for the constant A_t, approach zero. A_t converges to approximately 0.0925. Thus, for games with long horizon the values converge to 0.0925, which are about 20% lower than the static Cournot profits of $\frac{1}{9}$. This is the first illustration of how the dynamic interaction leads to a reduction in profits. If they could, the two parties would have liked to avoid the "preparation race" and commit to the static Cournot outcome throughout.

Figure 2 presents the symmetric equilibrium path for different initial production plans. As long as these production targets are not too high, the two parties begin by increasing their targets, each trying to become a Stackelberg leader, or at least not to fall behind and become a Stackelberg follower. As the deadline gets closer, both firms are sufficiently committed to high output, so the static best response functions begin to dominate. Both parties realize that they are much above their static best responses, and optimally decide to gradually adjust towards it. Due to adjustment costs, the parties will always end up at higher levels of output than those implied by their static Nash equilibrium. In the particular example, the equilibrium outcome is about 0.37, compared to the static outcome which is $\frac{1}{3}$. The reason that $\frac{1}{3}$ cannot be reached is that the profit function is flat at the static best response level. Thus, with any positive adjustment costs, no party will ever fully adjust all the way to her static best response level. The optimal strategy will always lead to only partial adjustment.

If initial production plans are sufficiently high (greater than about 0.44 in this particular example), however, both parties are sufficiently committed to high production at date zero, so there is no need to engage in further increases of production targets. The rate at which they decrease their production targets over time is not constant, however, due to the commitment effect. They first decrease quantities slowly, so they remain committed to high quantities, and only later on they speed up this rate in the direction of their static best response levels.

Figures 3 and 4 present comparative statics with respect to the length of the horizon and with respect to the size of the adjustment cost parameter. To keep intuition simple, in both figures we keep initial production plans fixed at the static Cournot level. A quick inspection of equation (9) reveals that these two exercises are somewhat similar. One can think about a proportional increase in adjustment costs as slowing down the evolution of the value function. Loosely speaking, this can be thought of as a horizontal stretch of Figure 1, so changes in the adjustment cost parameter are similar to a rescaling of time. Figure 3 shows how the length of the horizon affects the equilibrium path. As the horizon gets longer, there is more time and room for dynamic effects, so production targets increase to higher levels, but ultimately decrease faster as both firms are further away from their static best response functions. Therefore, the equilibrium outcome is not affected "by much." Figure 4 shows that this effect is not monotone. When adjustment costs are decreased, there are two effects in play. The direct effect makes it cheaper to increase production targets, creating an incentive for more competitive targets. In contrast, lower adjustment costs also make the commitment value of higher production targets lower, reducing the incentive to increase production targets. In Figure 4 one can see that once adjustment costs are low enough ($\theta = 0.2$) firms do not increase their production targets as much.

The non-monotonic effect of the adjustment costs is clearly seen by looking at the two extreme cases. As θ goes to zero, firms lose the ability to commit and the game reduces to a cheap talk. Thus, in equilibrium players just play their static Cournot levels throughout. As θ goes to infinity, firms stick to their initial plans and, again, a flat equilibrium path will arise. Thus, only intermediate values of θ give rise to the patterns shown in the graphs. As the effect of θ is smooth, it has to be case that the effect of the size of the adjustment costs is non-monotonic.

2.3 Intuition from a two-period model

The key qualitative prediction of the model is that firms start by exaggerating their production intentions during some time interval before the production date. The intuition for this result is straightforward and can be achieved by exploring a simple two-period analog. Suppose that at period $t = 0$ firms start by having production targets of y_i and y_j. At period $t = 1$ firms can revise their plans to x_i and x_j, but pay quadratic adjustment costs when they do so. Finally, in period $t = 2$ firms have a final opportunity to revise the quantities they want to produce and set them to q_i and q_j, paying the corresponding adjustment costs. Given these production levels, market price is given by $p(q_i + q_j)$, where $p(\cdot)$ satisfies the standard assumptions about inverse demand function. There is no discounting. Payoffs are the final Cournot profits (with zero marginal costs) minus any adjustment costs incurred in the process.

We can now solve for the Markov perfect equilibrium of the game using backward induction. In period $t = 2$ each player i chooses q_i to solve

$$\max_{q_i} q_i p(q_i + q_j) - \frac{\theta}{2}(q_i - x_i)^2 \tag{11}$$

so for each player i, q_i has to satisfy

$$p + q_i p' - \theta(q_i - x_i) = 0 \tag{12}$$

One can easily observe that if x_i, x_j are the static Nash equilibrium quantities, then setting $q_i = x_i$ for each i is an equilibrium. In general, the first order conditions define a best-response function which is a rotation of the static best-response at the previously targeted production level. Because of the adjustment costs, each player's response to a change in her opponent's quantity is not as strong as it would otherwise be. With downward sloping best responses, it can be easily shown that if $x = x_i = x_j$ is greater (less) than the symmetric static Nash equilibrium quantities, the players will end up adjusting in the direction of their static best responses, but not fully, thereby ending up at a more (less) competitive equilibrium. To see this, consider the symmetric solution to the first order condition, denoted by q^*. By taking derivative of the first order condition with respect to x and rearranging, we get that

$$\frac{\partial q^*(x)}{\partial x} = \frac{-\theta}{3p' + 2q^*(x)p'' - \theta} \tag{13}$$

which is positive provided that the inverse demand function is not too convex (a similar condition to the one that guarantees downward sloping best response curves). As already mentioned, at x equal to the static Nash equilibrium, we have that $q^*(x) = x$. Thus, we obtain our result. Figure 5 illustrates this situation for a linear-quadratic framework. One can see how the best response curves rotate at the planned quantity x_i, and become flatter. In the figure, we show how higher production plans lead to a reduction in equilibrium quantities (compared to target levels), but still to a more competitive outcome than the static Nash equilibrium.

In period $t = 1$ firms choose x_i and x_j, accounting for the equilibrium strategies at $t = 2$. Let $q_i(x_i, x_j)$ and $q_j(x_i, x_j)$ be the solution to the system defined by equation (12), then at $t = 1$ each player i chooses x_i to solve

$$\max_{x_i} q_i(x_i, x_j)p(q_i(x_i, x_j) + q_j(x_i, x_j)) - \frac{\theta}{2}(q_i(x_i, x_j) - x_i)^2 - \frac{\theta}{2}(x_i - y_i)^2 \tag{14}$$

implying the following first order condition for each player:

$$\frac{\partial q_i}{\partial x_i}p + \left(\frac{\partial q_i}{\partial x_i} + \frac{\partial q_j}{\partial x_i}\right) q_i p' - \theta(q_i - x_i)\left(\frac{\partial q_i}{\partial x_i} - 1\right) - \theta(x_i - y_i) = 0 \tag{15}$$

For simplicity, we can now assume a linear demand function, i.e. $p(Q) = a - bQ$. Using equation (12), best response functions (at $t = 2$) are given by

$$q_i = \frac{a - bq_j + \theta x_i}{2b + \theta} \tag{16}$$

and the $t = 2$ equilibrium outcome is given by

$$q_i(x_i, x_j) = \frac{ab + \theta\left(a + (2b + \theta)x_i - bx_j\right)}{3b^2 + 4b\theta + \theta^2} \tag{17}$$

which is linear in x_i and x_j. It is now easy to see that $\frac{\partial q_i}{\partial x_i} = \frac{2b\theta + \theta^2}{3b^2 + 4b\theta + \theta^2} = k \in (0, 1)$ and $\frac{\partial q_i}{\partial x_j} = \frac{-b\theta}{3b^2 + 4b\theta + \theta^2} = -l \in (-\frac{1}{4}, 0)$. It is also easy to see that $l = \left|\frac{\partial q_i}{\partial x_j}\right| < \left|\frac{\partial q_i}{\partial x_i}\right| = k$. Substituting this into equation (15) one could verify that if initial production targets and x_j are at the static Cournot level, the first order condition is positive, so player i would optimize by setting production target which is higher than her static Cournot level. Simplifying more, by setting $a = b = \theta = 1$, one can get $t = 1$ equilibrium strategies to be $x_i = \frac{1}{104\frac{16}{21}}(17\frac{1}{7} + 64y_i - \frac{64}{21}y_j)$. If we start at the Cournot level, i.e. $y_i = y_j = \frac{1}{3}$ we will get an equilibrium with $x_i = x_j \approx 0.357$, which is higher than the Cournot level. It is easy to see that at $t = 2$ we will then set $q_i = q_j \approx 0.339$. Thus, the qualitative conclusions are the same as in the continuous time case. Planned production levels increase first, and decrease later.

2.4 Initial actions

Our analysis so far took the "inherited" initial production levels as given. Clearly, one would like to endogenize them. In this section we shed some light on how to do so.

One natural way would be to allow players to freely decide on their initial plans simultaneously at date zero. In the continuous linear-quadratic case, this simultaneous-move game at date zero has a unique equilibrium, as best response functions are linear. For example, in the two player case the equilibrium is given by each player solving

$$\max_{q_i} \left(A_0 + B_0 q_i + C_0 q_j + D_0 q_i^2 + E_0 q_j^2 + F_0 q_i q_j \right) \tag{18}$$

which leads to a first order condition of

$$B_0 + 2D_0 q_i + F_0 q_j = 0 \tag{19}$$

In the symmetric case, the equilibrium is

$$q_j = q_i = \frac{-B_0}{2D_0 + F_0} \tag{20}$$

which gives rise to a "flat" equilibrium path until a certain point, and then to the regular decline in production plans towards the end of the game. To understand why this path is initially flat, all we need to see is that optimal choices in the Nash equilibrium of the initial production plans game satisfy the first order condition $\frac{\partial V_0^i(q_i,q_j)}{\partial q_i} = 0$. At the same time, the dynamic equilibrium strategies determining the adjustment rates are given by equation (7), namely $x_0^i(q_i, q_j) = \frac{1}{\theta} \frac{\partial V_0^i(q_i,q_j)}{\partial q_i}$, which implies $x_0^i(q_i, q_j) = 0$.

There are other aspects that can be introduced into the model, affect the initial decisions, and give rise to the various paths described in Section 2, and in particular to the non-monotonic equilibrium path. Note that, as shown in Figure 1, the value function at date zero becomes constant as the horizon goes to infinity. Thus, while there is a unique best response of the initial production plan for any finite horizon, players are almost indifferent among any action taken at that point. This implies that if there are other effects which may play a role in setting initial actions, these effects are likely to dominate the strategic effects arising from the dynamic game.

What other effects may be important? Let us emphasize two possible sources. The first is the introduction of uncertainty, which is discussed in more detail in Section 3. Suppose that market conditions (e.g. the constant of the inverse demand function, a) at the deadline can be good or bad, and that players get to observe the realization only after setting their initial production plans. In such a case, players would tend to initially set some weighted average between the optimal initial plan for the good realization and the optimal initial plan for the bad realization. If conditions turn out to be low, we should then observe a downwards equilibrium path. If conditions are high, however, we will observe the non-monotonic path arising.

A second potential effect that may affect initial production plans has to do with the existence of an already installed technology at date zero. This is very reasonable if one thinks that these production encounters may happen repeatedly. Think of the application of this model to auto manufacturers. They do not build a factory from scratch every month. For instance, it may be reasonable to assume that they start with an inherited initial structure that would lead to a final production similar to last period's production. Now, one can imagine firms paying adjustment costs if their first production plan is far away from the last actual production level. This, coupled with the fact that the objective function is quite flat, would imply that in a stationary equilibrium such adjustment costs would dominate the initial decision, and would make firms initially set production targets that will give rise to the non-monotonic path. We discuss this extension more in the next section.

3 Extensions to the basic model

The linear-quadratic model is attractive as it is quite simple to generalize the basic model along several dimensions. In this section we briefly illustrate how the results extend. Most extensions would retain the qualitative predictions of the model. This is true if we introduce time-varying adjustment costs, asymmetries among players, symmetric uncertainty, or more than two players. We also discuss the extension of the model to include repeated interaction. There are two cases in which the qualitative predictions are different. This happens when the final strategic interaction is of strategic complements, or when there is flow payoffs rather than payoffs that are collected only in the end.

3.1 Asymmetric players

One could relax the symmetry assumption. Asymmetry can be introduced either through the final payoff function (for example, firms may vary in their marginal costs) or through asymmetry in the adjustment costs (for example, the labor of one firm is unionized while that of the other is not). In what follows, we allow both, but will later do comparative statics on each dimension separately.

We keep notation as before, with the addition of superscripts to denote the identity of the player. Thus, player i's adjustment costs function is now $c_i(x_i, t) = \frac{\theta^i}{2} x_i^2$, and her (constant) marginal costs are c^i. The same goes for the time-varying parameters of each player's value function.

One can follow exactly the same steps as in Section 2 until the first point in which we imposed the symmetry assumption, i.e. equation (8). Now, equation (8) can be written as

$$0 = \frac{1}{2\theta^i}\left(\frac{\partial V_t^i}{\partial q_i}\right)^2 + \frac{1}{\theta^j}\left(\frac{\partial V_t^i}{\partial q_j}\right)\left(\frac{\partial V_t^j}{\partial q_j}\right) + \frac{\partial V_t^i}{\partial t} \tag{21}$$

The value function for each player is now different, and we write it as

$$V_t^i(q_i, q_j) = A_t^i + B_t^i q_i + C_t^i q_j + D_t^i q_i^2 + E_t^i q_j^2 + F_t^i q_i q_j \tag{22}$$

Substituting it into equation (21) gives us

$$\begin{aligned} 0 \;=\; & \frac{1}{2\theta^i}\left(B_t^i + 2D_t^i q_i + F_t^i q_j\right)^2 + \frac{1}{2\theta^j}\left(C_t^i + 2E_t^i q_j + F_t^i q_i\right)\left(B_t^j + 2D_t^j q_j + F_t^j q_i\right) + \\ & + \left(A_t^{i\prime} + B_t^{i\prime} q_i + C_t^{i\prime} q_j + D_t^{i\prime} q_i^2 + E_t^{i\prime} q_j^2 + F_t^{i\prime} q_i q_j\right) \end{aligned} \tag{23}$$

By collecting terms, we obtain the following law of motion for the parameters in player i's value function (symmetrically for player j):

$$\begin{pmatrix} A^{i\prime} \\ B^{i\prime} \\ C^{i\prime} \\ D^{i\prime} \\ E^{i\prime} \\ F^{i\prime} \end{pmatrix} = \begin{pmatrix} \frac{1}{2\theta^i} B^{i2} + \frac{1}{\theta^j} B^j C^i \\ \frac{2}{\theta^i} B^i D^i + \frac{1}{\theta^j} B^j F^i + \frac{1}{\theta^j} C^i F^j \\ \frac{1}{\theta^i} B^i F^i + \frac{2}{\theta^j} B^j E^i + \frac{2}{\theta^j} C^i D^j \\ \frac{2}{\theta^i} D^{i2} + \frac{1}{\theta^j} F^i F^j \\ \frac{1}{2\theta^i} F^{i2} + \frac{4}{\theta^j} D^j E^i \\ \frac{2}{\theta^i} D^i F^i + \frac{2}{\theta^j} D^j F^i + \frac{2}{\theta^j} E^i F^j \end{pmatrix} \tag{24}$$

with the boundary condition given, for example, by

$$\begin{pmatrix} A_T^i \\ B_T^i \\ C_T^i \\ D_T^i \\ E_T^i \\ F_T^i \end{pmatrix} = \begin{pmatrix} 0 \\ a - c^i \\ 0 \\ -b \\ 0 \\ -b \end{pmatrix} \tag{25}$$

It may be interesting to look at comparative statics. Figure 6 illustrates the case of asymmetric final payoffs. In particular, it uses the same example as the one used in Section 2, but introduces a (constant) marginal cost of 0.2 for player 2. Static Nash equilibrium is now given by $q_1 = 0.4$ and $q_2 = 0.2$. Figure 6 presents this case (with identical adjustment costs) for different initial conditions. The general pattern is quite similar to the symmetric case, with the more efficient player always producing more than her opponent, and more than her static Nash equilibrium quantity. There are two interesting points to note. First, not surprisingly, this setup can give rise to cases in which one of the players has the non-monotonic equilibrium path while the other only goes down. This happens when the initial plans for both players are intermediate (0.3 in the example plotted in Figure 6). This is analogous to the symmetric setup, in which players vary in their initial production plans. In the asymmetric case, such variation may seem more realistic. For example, one can imagine that both firms set their initial plans first, and only later learn which one of them enjoys the cost advantage. The second interesting observation is that with asymmetries, it can happen that the less efficient player eventually ends up producing less than her static Nash quantity. This is shown in the thin solid line. This happens because once we introduce asymmetries in payoffs, it also implies asymmetries in commitment opportunities. Although both players have identical adjustment costs, the more efficient player is producing more, so, in absolute terms, her static payoff function is steeper around the equilibrium. This allows her to have better commitment opportunity, thereby enjoying a Stackelberg advantage. This makes the less efficient firm take a role of a Stackelberg follower, thereby producing less than its static Nash quantity. In all cases, however, overall quantity is higher (more competitive) than the static equilibrium level of 0.6. The fact that the more efficient firm enjoys, *ceteris paribus*, a commitment advantage also implies some allocative efficiency, so welfare is higher due to both higher consumer surplus and more efficient allocation of resources among the firms.

Figure 7 presents the case of symmetric payoff functions (identical to the example of Section 2), but asymmetric adjustment costs. In the limit of such a case, when one player has zero adjustment costs while the other has infinite adjustment costs, we are approaching the simple Stackelberg case, with the zero adjustment costs player acting as a Stackelberg follower, as her opponent can commit to any level of output. Figure 7 presents more intermediate cases, which illustrate this. As the cost asymmetry is higher, the commitment opportunity for the high adjustment cost firm increases, and we get closer to the Stackelberg outcome of 0.5 and 0.25. One should note, however, that with finite time, it is not only that the relative adjustment costs matter, but the absolute ones as well. For example, keeping the relative adjustment costs the same, the thin solid line can be compared to the thick solid line. As can be expected, lower level of adjustment costs (the thin line) makes commitment less important, and the eventual quantities produced are somewhat closer to the static best response. It is somewhat interesting to note that this case (the thin line) shows that one could generate a slight reversed non-monotonic equilibrium path for the Stackelberg follower in this model.

3.2 Uncertainty

Adjustment costs introduce a mechanism for commitment in this model. Firms would like to commit to enjoy a first-mover advantage, which is the main reason for the non-monotonic equilibrium path. With uncertainty, however, commitment comes at a cost, as firms who make a commitment are not as flexible to react to unexpected events later on. To understand the trade-off, we introduce the simplest source of uncertainty into the linear-quadratic model. Consider the two-player symmetric game, but assume that final demand can take one of two values, high (H) or low (L). We assume that firms have no private information, and that at any given point in time their (common) beliefs are such that either H or L will be realized. These beliefs follow a symmetric Markov process: with probability λ per unit time the beliefs will change, and with probability $(1-\lambda)$ they will remain the same.[6]

[6] One could also think about firms getting continuous signals about the state of the world in time T, and using Bayesian updating to revise their beliefs. This somewhat more natural modeling is likely to generate similar

We can now follow the same steps as in Section 2 with few modifications. We now have two value functions, as the additional state variable is another indicator. Let these two value functions be V_{Lt} and V_{Ht}. Thus, the equation for V_L is given by:

$$V_{Lt}^i(q_i, q_j) = \max_{x_i} \left(-\frac{\theta}{2}x_i^2 + \frac{\partial V_{Lt}^i}{\partial q_i}x_i + \frac{\partial V_{Lt}^i}{\partial q_j}x_j + \frac{\partial V_{Lt}^i}{\partial t} + V_{Lt}^i(q_i, q_j) + \lambda \left(V_{Ht}^i(q_i, q_j) - V_{Lt}^i(q_i, q_j) \right) \right) \tag{26}$$

and symmetrically for V_H. The structure of uncertainty is attractive as the continuous choice of adjustment rate does not directly depend on λ, as at a given point in time there is probability zero for a change in beliefs. Optimal adjustment rate is given by

$$-\theta x_i + \frac{\partial V_{Lt}^i}{\partial q_i} = 0 \Longrightarrow x_i = \frac{1}{\theta}\frac{\partial V_{Lt}^i}{\partial q_i} \tag{27}$$

and symmetrically for H. It clearly depends on λ indirectly, through the dependence of the value functions on λ.

This structure results in an identical system of ordinary differential equation, with only one modification. In each equation, with probability λ we switch to the other value function. Let $s = L, H$ be the state of the world, and let $r = H, L$ be the other state. The system of 12 differential equations, which defines the law of motion for the parameters, is given by:

$$\begin{pmatrix} A_s' \\ B_s' \\ C_s' \\ D_s' \\ E_s' \\ F_s' \end{pmatrix} = \lambda \begin{pmatrix} A_r - A_s \\ B_r - B_s \\ C_r - C_s \\ D_r - D_s \\ E_r - E_s \\ F_r - F_s \end{pmatrix} + \frac{1}{\theta} \begin{pmatrix} \frac{1}{2}B_s^2 + B_sC_s \\ 2B_sD_s + B_sF_s + C_sF_s \\ B_sF_s + 2B_sE_s + 2C_sD_s \\ 2D_s^2 + F_s^2 \\ \frac{1}{2}F_s^2 + 4D_sE_s \\ 4D_sF_s + 2E_sF_s \end{pmatrix} \tag{28}$$

with the boundary condition given by the different profit functions at each state.

In the equilibrium of this game, firms will be somewhat more reluctant to commit. After all, firms care about the current belief about demand only to the extent that it affects the eventual realization of demand. With the horizon far enough into the future, the current state is not particularly informative about the final state of demand, so the two value functions will be similar to each other. Only as the horizon draws near, firms will be more willing to act upon their beliefs. Thus, one may think about this in a similar way to the time-varying adjustment costs case described below: firms will be more flexible towards the end, as early on they will be reluctant to take an action. This is somewhat similar to a rescaling of time.

The equilibrium path of such a game would clearly depend on the realization of firms' beliefs. It may be worth mentioning the two extremes, when the beliefs never change throughout the game. In this case, firms will initially set some intermediate level of initial production target. As

qualitative results. It is more complicated as it introduces a continuous state variable, while the setup we use introduces only a binary state variable.

the horizon draws near, they will either have the non-monotonic equilibrium path, in the case of beliefs being always H, or the monotone (downwards) path in the case of beliefs being always L.

3.3 Repeated Interaction

As our application described below is focused on a case where repeated interaction among players is present, it is important to show that our qualitative results can still hold. We sketch below the simplest possible extension of our framework for a repeated game. This extension also provides another rationale for the choice of initial production plans, which will give rise to the non-monotonic equilibrium path.

Consider two symmetric players playing an infinitely repeated game. Time is continuous, but production takes place only at a discrete set of points in time, at which payoffs are collected (except for adjustment costs, which, as before, are spent continuously). Players discount profits with a common discount factor of δ. Between two consecutive production points, players engage in a production preparation stage. This stage involves quadratic adjustment costs, just as in our basic model. As we think about production plans as involving real actions, it seems plausible that (quadratic) adjustment costs are also paid when initial production plans for a subsequent production period are different from last period actual production levels. These adjustment costs are of somewhat different nature from adjusting production plans, so they may have a different parameter. We solve for the Markov Perfect Equilibrium of the model. As each production period is identical, we search for the symmetric stationary equilibrium.

To summarize, given last period production of (y_i, y_j), each period lasts T units of time and has the following structure:

- At $t = 0$ each player i sets her initial production plan for next period, $q_i(0)$. When doing so she pays production adjustment costs of $\frac{\mu}{2}(q_i(0) - y_i)^2$.
- At $t \in (0, T]$ each player i continuously adjusts her production target, paying an adjustment cost of $\frac{\theta}{2}(q_i'(t))^2$ in the process.
- At $t = T$ each player i produces her target quantity, $q_i(T)$, and collects payoffs of $(a - b(q_i(T) + q_j(T)) - c)q_i(T)$.

This process is then being infinitely repeated.

Given that the game still has a linear-quadratic structure, the value function remains quadratic in the state variables. Thus, the solution to the value function within each period will be the same as in our benchmark model, but with different boundary conditions. The boundary conditions will now be determined endogenously, as part of the steady state. The qualitative nature of the results, however will remain unchanged. Initial production plans will be chosen optimally, but because of the production adjustment costs will typically be lower than the unconstrained choice.

We sketch the qualitative pattern of the equilibrium path in Figure 8. Note that the level of production in such an equilibrium will be even higher than the production level in our basic model. This is because in addition to the commitment effect, which we already described, there is an additional dynamic effect of commitment through the production adjustment costs. This second effect is the same as in the standard infinitely repeated linear-quadratic models analyzed in the literature (Maskin and Tirole, 1986; Reynolds, 1987 and 1991; Jun and Vives, 2004).

3.4 Other Extensions

N **players** The basic results remain unchanged with more than two players. We consider only the fully symmetric case, as things become quite messy if we consider many asymmetric players. The derivation is provided in Appendix A. The system of differential equations gives rise to qualitatively identical equilibrium patterns as that obtained from the basic two player model. The only difference from the two-player case is that the value function has one additional element. Player i's value now depends also on the distribution of target quantities among her opponents. In particular, $\sum_{j\neq i}\sum_{k\neq i,j} q_j q_k$ enters player i's value. As may be expected, the parameter on this additional element evolves quite similarly to the parameter on $\sum_{j\neq i} q_j^2$, namely E_t, presented in Figure 1.

Time-varying adjustment costs One may argue that adjustment costs may vary over time. One reason for this may be discounting. While payoffs are collected only in the end, adjustment costs are spent over time, so discounting would make future adjustment costs relatively lower. A different reason for time-varying costs may lead to the opposite effect. It may be reasonable to believe that adjustments are more difficult as the production date gets closer. As an example, hiring more temporary labor three months before production may be cheap, while labor availability one day before production is scarce, and will require higher wages or higher search costs on the employer part. The second case would make this model a special case of the general framework proposed in Caruana and Einav (2004).

Fortunately, it is straightforward to incorporate such effects into the model. All one needs to do is to specify the adjustment cost function as

$$c_i(x_i, t) = \frac{\theta(t)}{2} x_i^2$$

imposing no restrictions on $\theta(t)$. Discounting would simply mean that $\theta(t) = \beta^t \theta$, while Caruana and Einav (2004) framework would mean that $\theta(0) = 0$ and $\theta(T) = \infty$. Since the model is not stationary by construction, this imposes no additional difficulties. The derivation of the system of ordinary differential equations remains the same as in equation (9), with the only difference being that θ is replaced by $\theta(t)$. As one can notice, however, θ enters to the system just in a proportional way. Therefore, replacing it by $\theta(t)$ is somewhat similar to a rescaling of time. When $\theta(t)$ is low the coefficients on the value function change fast, and when $\theta(t)$ is high the coefficients change slow. All other qualitative predictions remain unchanged.

Flow payoffs One can modify our model so payoffs are collected continuously rather than only in the end. This will make the structure of the model quite similar to Cyert and DeGroot (1970), with the main difference being the commitment technology. Cyert and DeGroot (1970) introduce commitment by an alternating move game, so a player is committed for one additional period in a discrete time case. We introduce commitment, as many other papers in the literature do in an infinite horizon context,[7] through adjustment costs.

Solving our differential game with flow payoffs is quite similar. The only difference is that it "moves" the boundary condition to the law of motion. The boundary condition in our case reflects the final payoffs. In the flow payoffs case, this is just the way payoffs are accumulated over time, thereby affecting the slope of the value function's parameters. In the basic model, this would mean a boundary condition of zeros for all parameters, and a law of motion of

$$\begin{pmatrix} A' \\ B' \\ C' \\ D' \\ E' \\ F' \end{pmatrix} = \begin{pmatrix} 0 \\ a \\ 0 \\ -b \\ 0 \\ -b \end{pmatrix} + \frac{1}{\theta} \begin{pmatrix} \frac{1}{2}B^2 + BC \\ 2BD + BF + CF \\ BF + 2BE + 2CD \\ 2D^2 + F^2 \\ \frac{1}{2}F^2 + 4DE \\ 4DF + 2EF \end{pmatrix} \tag{29}$$

The important point to note is that despite the fact that this flow payoff game is not stationary, it behaves very close to a stationary game. After a quick adjustment to the stationary path in the beginning and a quick adjustment in the end, for most of the time things are stable. The reason is that the incentive to commit do not change over time as much as they do in our original game. This is because much of the payoffs from committing are collected along the way, reducing future incentives to commit. This feature is the main reason why, indeed, Cyert and DeGroot (1970) mainly focus on the stationary part of their finite-horizon game, and that the literature has so far found it somewhat unappealing to look at finite horizon flow-payoff games. As we show in this paper, the finite horizon has much stronger effects when payoffs are only collected in the end, as in the basic game presented in this paper.

Strategic complements One can notice that in our derivation of the law of motion for the game, the specific assumption about quantity competition was not important. It only affected the solution through the boundary condition. In that sense, one could easily apply the same structure for other linear-quadratic settings, such as games of strategic complements. For example, suppose firms compete by setting prices, and pay adjustment costs any time they change the price. Indeed, Jun and Vives (2004) analyze a flow-payoff infinite-horizon linear-quadratic games of both strategic substitutes, strategic complements, and a mixture of the two.

[7] See, among others, Masking and Tirole (1987), Fershtman and Kamien (1987), Lapham and Ware (1994), Reynolds (1987, 1991), and Jun and Vives (2004).

There are two reasons why we focus all our attention on quantity competition and strategic substitutes. The first is because of the data. In our data auto manufacturers set production plans, making quantity to be the natural strategic variable in this setup. Second, coordination rather than commitment is central to games with strategic complements. Therefore, the strategic interaction obtained is not that different from that of standard infinitely repeated games. Loosely speaking, the preparation phase only allow firms to coordinate on higher prices, but does not lead to interesting time dependencies.

4 Evidence from the auto industry

We analyze data about production targets of the major auto manufacturers in the U.S. These decisions are also studied by Doyle and Snyder (1999).[8] Prior to each production month, the major U.S. auto manufacturers – General Motors (GM), Ford, Chrysler, and American Motors (AMC) – decide about their production targets for future months. Estimates of these targets are published in *Ward's Automative Reports* for all manufacturers, as early as six months prior to the actual production date.[9] Targets are summarized by the number of cars to be produced by each manufacturer, aggregated over all brands and models. *Ward's Automative Reports* is a weekly industry newsletter, specializing in industry data and statistics.[10] The data set has a panel structure and covers the years 1965 to 1995, for a total of 372 production months.[11] While production target estimates are typically published on a monthly basis, the number of published estimates vary across production months. Overall, for each production month, we observe 1 to 12 production targets (for each manufacturer), with a mean of 4.35. The data include 1,620 target levels for GM, Ford, and Chrysler, and 1,114 for AMC (whose data are only available through mid 1987). These production targets are later matched with actual production figures. Figure 9 presents the total number of published estimates made at each 10 day interval prior to actual production. As one can see, production plans are typically published once a month, typically on the last week of the month, although one can see some density between the monthly peaks. One can also observe that the number of observations is quite stable over the 3-4 months before production. There are significantly fewer earlier observations.

It is important to discuss the fit of the automobile application to the theoretical framework.

[8] We are extremely grateful to Maura Doyle and Chris Snyder for the willingness to share their data with us. We will be brief in describing the data here, as one can refer to Doyle and Snyder (1999) for more details.

[9] These targets are being described by various synonyms: "assemply targets," "assembly schedules," "production plans," "production forecasts," etc.

[10] Potential readers are encouraged to buy subscription to the newsletter by the following quote posted on *Ward's* web site: "News and numbers you can't do without. Auto analysts and decision-makers must get the latest, vital statistics on the industry's health, plus updated news, analysis and projections that impact their companys' futures." (*http://wardsauto.com/war/index.htm*)

[11] Some of the observations in the data include post production revisions. We omit these observations. We only focus on estimates made *before* production.

In the model, production plans have a commitment value as they are costly to change; otherwise, they would be pure cheap talk. In order for this to be case, it may be useful to discuss what the production targets in the data may stand for. We think about the main auto manufacturers continuously taking actions that affect their future production capabilities (contracting to hire more labor, canceling vacations, contracting more parts from suppliers, etc.). These actions are observable to their competitors, just as they are observable to the publisher of *Ward's.* The newsletter reports these actions to third parties (suppliers, dealers, analysts). *Ward's* only reports snapshots of those actions on a monthly basis. Thus, we think about these discrete observations as reflecting an underlying continuous decision process, such as the one described by our theoretical model. It seems natural to assume that such production-related decisions are costly to change. Moreover, if it were purely cheap talk, it seems unreasonable that *Ward's* would have found it profitable to publish it. The only reason to publish such information is if it had some commercial value to third parties. In addition, one should note that if such third parties act upon this information, and if these actions affect manufacturers' profits, this by itself creates adjustments costs of the type analyzed by the theoretical model.

Consequently, we will focus on two key variables. The first is the time until the deadline (in days). The second is the production target. In order to make targets comparable over time and across manufacturers, we normalize all targets by eventual production. A data point in our analysis is (d_{it}, q_{it}) for manufacturer i and production month t. d_{it} is the number of days between the day the estimated production target was published and the last day of the production month, for which it was made. q_{it} is the normalized target, i.e.

$$q_{it} \equiv 100 \left(\frac{P_{it} - Q_{it}}{Q_{it}} \right) \tag{30}$$

where Q_{it} is actual production by manufacturer i at month t, and P_{it} is manufacturer i's production plan. This transformation of the data is similar to the PPE measure used in Doyle and Snyder (1999). Our measure has the opposite sign and uses a slightly different normalization in order to more closely relate the variables to the theoretical predictions.[12] Thus, q_{it} is positive (negative) when a manufacturer plans higher (lower) production than she eventually produces. Our key theoretical prediction is that manufacturers will typically exaggerate, and that towards the production date they will gradually reduce their production plans.

Our basic evidence is based on pooling observations from multiple production months. The basic assumption that justifies it is that, up to the normalization discussed above, the same game is being played repeatedly over time. It enables treating different production targets in different games as if they are made in the same context. We then use quartic (biweight) kernel regressions of q_{it} on d_{it} to describe nonparametrically the evolution of production plans over time. In all figures, we use a bandwidth of 30 days.

[12] This normalization requires us to drop two production targets by GM for October 1970, a month in which GM produced (almost) no cars, i.e. $Q_{it} = 0$. All our qualitative results carry through if, instead, we use the PPE measure of Doyle and Snyder (1999), which would allow us to keep these two observations.

Figure 10 presents the pattern for the major three manufacturers, and for their average. One can observe a sharp decline in production plans towards the actual production date. This is true for the average, as well as for each of the manufacturers separately. This is consistent with the predictions of the model. The average shows a non-monotonic pattern: it peaks approximately 2-3 months before production at production targets which are about 10 percent above actual production, and then declines. One can see, however, that this pattern is not uniform across manufacturers. While Ford and Chrysler, the two smaller firms, follow similar non-monotonic pattern, GM exhibits a very different behavior. GM's average initial production target is about 15 percent higher than its eventual production, and it gradually declines as the deadline gets closer. This is, of course, not inconsistent with the model: if initial production targets are high, the model predicts a gradual decline over time. It would be interesting to explain why GM's (relative) initial production plans are consistently much higher than those of Ford and Chrysler. As mentioned in Section 3, uncertainty may help to rationalize it. For example, if uncertainty is not about aggregate demand but about an event that has different effects on different firms, the theoretical model can predict differential patterns for different manufacturers. As an example, if firms are uncertain whether GM's workers will go on a strike or not, such uncertainty may result in a different behavior for GM and for its rivals.[13]

Figure 11 repeats the same exercise for AMC, as well as for the average of all four manufacturers. The solid line in Figure 11 is the same as the thick solid line in Figure 10 to facilitate comparison between the two figures. As can be easily seen, AMC exhibits a similar trend to that of GM, but its magnitude is much higher. AMC begins with an average initial production target of almost 80 percent higher than its eventual production. We are not completely sure about the interpretation of this. One should remember, however, that AMC's data cover a shorter observation period, and account for an average market share of about 2%, compared to much higher market shares of the other three manufacturers (42%, 21%, and 11% for GM, Ford, and Chrysler, respectively). Thus, this makes the underlying strategic effect of AMC negligible for two reasons. First, if adjustment costs are related to the absolute (not relative) magnitude of adjustments, adjusting downwards by AMC is pretty cheap, giving its exaggerated production plan little commitment value. Second, due to AMC's tiny market share, 80% increase in its anticipated production still has little effect on its competitors' profits. Despite this qualification, one should note that this pattern is still consistent with the qualitative predictions of the model.

Figure 12 reports the above kernel estimates with 95 percent confidence intervals. We compute the confidence interval by bootstrapping the data, and running the same kernel regression on each bootstrapped sample. The dashed lines in each figure report the point-by-point 2.5 and 97.5 percentiles, while the solid line reports the 50th percentile, which is approximately the same as the estimates reported above. It shows that the observed decline in planned production towards the production deadline is quite precisely estimated. It also shows that the confidence intervals

[13] The important effects of labor strikes on production during the observation period are discussed in Doyle and Snyder (1999).

shrink as the production deadline gets closer. This happens for two reasons. First, as may be expected, the variance in the estimates is lower close to the day of production. This may be due to various factors related to uncertainty, which are outside of our model. Such factors are likely to be more pronounced when the production deadline is further away in the future. The second reason is because the number of early production target estimates observed is significantly smaller than the number of observations closer to the production deadline.[14]

One should note that none of the empirical findings here are inconsistent with the empirical findings of Doyle and Snyder (1999). Doyle and Snyder focus on the positive correlation among revisions to production plans by different manufacturers, which is interpreted as evidence for information sharing. Our theoretical model also predicts such positive correlation, but due to strategic considerations. Doyle and Snyder (1999) also point out that production plans are, on average, higher than actual production. They do not analyze this pattern as their main theoretical framework of information sharing does not provide any restrictions on this dimension. Finally, we should also note that our findings do not imply that information-sharing has no role in this setting. The observed pattern of production plans may well be driven by both information-sharing motives as well as strategic commitment considerations. In fact, we pool observation from different periods in order to average out the period-specific noise. The period-specific patterns vary quite substantially, and may be driven by different realizations of uncertainties. Our framework is therefore more relevant for the average pattern rather than for the period-by-period pattern, while information-sharing motives are more likely to be important and observed *within* production periods.

One important gap between the model and the evidence is that the model analyzes sales, while the available data is about production.[15] For these to be similar, inventories (and, to a lesser extent, quantity produced abroad) should be roughly stable over time.[16] Modeling inventories is beyond the scope of this paper. In our view, however, it seems difficult to construct a model that, by considering inventory fluctuations *per se*, could generate the pattern in production plans that we observe. The fact that our results rely on a panel rather than on a cross-section implies that inventory fluctuations should be integrated out once we average over all monthly production periods. Still, of course, inventories change the strategic environment, and modeling it in a more structural way may be useful. We leave this for future research.

Finally, our basic theoretical framework is focused on a one-time production interaction, which follows a dynamic phase of gradual commitment. In contrast, the data is generated by a set of

[14]These two reason are consistent with each other. If early production targets are subject to a great deal of uncertainty, they have less value to industry decision makers, making them less likely to be reported by the newsletter.

[15]We should note that we are not the first to use a Cournot framework to approximate the strategic interaction in this market. Doyle and Snyder (1999) do the same, as well as Berndt, Friedlander, and Wang Chiang (1990), who cannot reject the Cournot model in this context.

[16]See Kahn (1992) and Bresnahan and Ramey (1994) for theory and evidence about the relationship between sales and production. See also Judd (1996) for a dynamic model of inventories in a framework similar to ours.

repeated production games. One could consider the interaction in the auto industry as a nested repeated game. The main auto manufacturers repeatedly play the production game, which could be approximated by a standard infinitely repeated game. Within each stage game, however, manufacturers engage in a dynamic preparation phase. As we discuss in Section 3, extending our framework in this direction is straightforward. In doing so, the qualitative predictions from the basic model carry through, thereby making our empirical analysis still valid.

5 Concluding remarks

In this paper we presented a dynamic commitment model in a Cournot framework. Pre-production preparations provide a commitment device, as changing them is costly. Thus, in a full information environment firms use such preparations strategically. We illustrate this point using a finite-horizon linear-quadratic differential game. We show that under these conditions firms have an incentive to exaggerate in their production targets in an attempt to achieve a Stackelberg leadership position. As a consequence, the final production levels are higher than in a static framework. More precisely, the model predicts that firms will first increase their intended production levels over time, and only later on, as the deadline gets closer, they will start lowering their production targets.

The finite horizon nature of the problem is an interesting feature of the model. It provides rich qualitative predictions that can be supported (or falsified) by the data. We use data on production plans of auto manufacturers to investigate the model's implications. The evidence show that, on average, auto manufacturers increase their production targets over time, until about 2-3 months before production, when they start decreasing them. This pattern is consistent with the theoretical prediction.

At this stage, we interpret these evidence as suggestive only. It seems to us, however, difficult to come up with alternative models which can generate the same qualitative predictions. Thus, we are encouraged to view these findings as empirical support for the relevance of the strategic role of pre-production preparations in determining final production decisions.

On a more methodological level, we think that this exercise illustrates the empirical potential of finite-horizon non-stationary models. When they are applicable, such models may provide sharper qualitative predictions, which have the potential to be empirically verified or falsified without the need for more precise structural assumptions.

References

[1] Berndt, Ernst R., Ann F. Friedlander, and Judy S. Wang Chiang (1990), "Interdependent Pricing and Markup Behavior: An Empirical Analysis of GM, Ford, and Chrysler," NBER Working Paper no. 3396.

[2] Bresnahan, Timothy F., and Valerie A. Ramey (1994), "Output Fluctuations at the Plant Level," *Quarterly Journal of Economics*, 109, 593-624.

[3] Caruana, Guillermo, and Liran Einav (2004), "A Theory of Endogenous Commitment," mimeo, Stanford University.

[4] Cyert, R. M., and M. H. DeGroot (1970), "Multiperiod Decision Models with Alternating Choice as the Solution to the Duopoly Problem," *Quarterly Journal of Economics*, 84, 410-429.

[5] Doyle, Maura P., and Christopher M. Snyder, "Information Sharing and Competition in the Motor Vehicle Industry," *Journal of Political Economy*, 107(6), 1326-1364.

[6] Fershtman, Chaim, and Morton I. Kamien (1987), "Dynamic Duopolistic Competition with Sticky Prices," *Econometrica*, 55(5), 1151-1164.

[7] Hanig, Marco (1986), "Differential Gaming Models of Oligopoly," Ph.D. Dissertation, Massachusetts Institute of Technology.

[8] Judd, Kenneth L. (1996), "Cournot versus Bertrand: A Dynamic Resolution," mimeo, Hoover Institution.

[9] Jun, Byoung, and Xavier Vives (2004), "Strategic Incentives in Dynamic Duopoly," *Journal of Economic Theory*, 116(2), 249-281.

[10] Kahn, James A. (1992), "Why is Production More Volatile than Sales? Theory and Evidence on the Stockout-Avoidance Motive for Inventory-Holding," *Quarterly Journal of Economics*, 107, 481-510.

[11] Kydland, Finn (1975), "Noncooperative and Dominant Player Solutions in Discrete Dynamic Games," *International Economic Review*, 16(2), 329-341.

[12] Lapham, Beverly, and Roger Ware (1994), "Markov Puppy Dogs and Related Animals," *International Journal of Industrial Organization*, 12(4), 569-593.

[13] Maskin, Eric, and Jean Tirole (1987), "A Theory of Dynamic Duopoly III: Cournot Competition," *European Economic Review*, 31, 947-968.

[14] Papavassilopoulos, G. P., and J. B. Cruz (1979), "On the Uniqueness of Nash Strategies for a Class of Analytic Differential Games," *Journal of Optimization Theory and Applications,* 27, 309-314.

[15] Papavassilopoulos, G. P., and G. J. Olsder (1984), "On the Linear-Quadratic, Closed-Loop, No-Memory Nash Game," *Journal of Optimization Theory and Applications, 4*2, 551-560.

[16] Reynolds, Stanley S., (1987), "Capacity Investment, Preemption and Commitment in an Infinite Horizon Model," *International Economic Review*, 28(1), 69-88.

[17] Reynolds, Stanley S., (1991), "Dynamic Oligopoly with Capacity Adjustment Costs," *Journal of Economic Dynamics and Control*, 15(3), 491-514.

Appendix

A N players

Consider $N > 2$ symmetric players. We can write the value function as

$$V_t^i(q_i, q_{-i}) = \max_{x_i} \left(-\frac{\theta}{2} x_i^2 + \frac{\partial V_t^i}{\partial q_i} x_i + \sum_{j \neq i} \frac{\partial V_t^i}{\partial q_j} x_j + \frac{\partial V_t^i}{\partial t} + V_t^i(q_i, q_{-i}) \right) \tag{31}$$

The first order condition for x_i is, as before

$$-\theta x_i + \frac{\partial V_t^i}{\partial q_i} = 0 \Longrightarrow x_i = \frac{1}{\theta} \frac{\partial V_t^i}{\partial q_i} \tag{32}$$

We can now plug this back into equation (31), as well as the symmetric solution for x_j, rearrange, and get the following differential equation

$$0 = \frac{1}{2\theta} \left(\frac{\partial V_t^i}{\partial q_i} \right)^2 + \frac{1}{\theta} \sum_{j \neq i} \left(\frac{\partial V_t^j}{\partial q_j} \right) \left(\frac{\partial V_t^i}{\partial q_j} \right) + \frac{\partial V_t^i}{\partial t} \tag{33}$$

We guess that the value function would be symmetric in the opponents' state variables, so it will only have one additional element compared to the basic model. The quadratic value function can be written as

$$\begin{aligned} V_t^i(q_i, q_j) &= A_t + B_t q_i + \sum_{j \neq i} C_t q_j + D_t q_i^2 + \sum_{j \neq i} E_t q_j^2 + \sum_{j \neq i} F_t q_i q_j + \sum_{j \neq i} \sum_{k \neq i,j} G_t q_j q_k = \\ &= A_t + B_t q_i + C_t Q_{-i} + D_t q_i^2 + E_t R_{-i} + F_t q_i Q_{-i} + G_t S_{-i} \end{aligned} \tag{34}$$

where $Q_{-i} = \sum_{j \neq i} q_j$, $R_{-i} = \sum_{j \neq i} q_j^2$, and $S_{-i} = \sum_{j \neq i} \sum_{k \neq i,j} q_j q_k$. Note that $Q_{-i}^2 = R_{-i} + S_{-i}$. This also implies that

$$x_t^i(q_i, q_j) = \frac{\partial V_t^i}{\partial q_i} = B_t + 2D_t q_i + F_t Q_{-i} \tag{35}$$

Thus, we can write equation (33) again to be

$$\begin{aligned} 0 \;=\; & \frac{1}{2\theta} \left(B_t + 2D_t q_i + F_t Q_{-i} \right)^2 + \frac{1}{\theta} \sum_{j \neq i} \left(C_t + 2E_t q_j + F_t q_i + 2G_t (Q_{-j} - q_i) \right) \left(B_t + 2D_t q_j + F_t Q_{-j} \right) + \\ & + \left(A_t' + B_t' q_i + C_t' Q_{-i} + D_t' q_i^2 + E_t' R_{-i} + F_t' q_i Q_{-i} + G_t' S_{-i} \right) \end{aligned} \tag{36}$$

After collecting terms (and reversing signs for A', B', etc. as before), we get the following law of motion:

$$\begin{pmatrix} A' \\ B' \\ C' \\ D' \\ E' \\ F' \\ G' \end{pmatrix} = \frac{1}{\theta} \begin{pmatrix} \frac{1}{2}B^2 + (N-1)BC \\ 2BD + BF(N-1) + CF(N-1) \\ BF + 2BE + 2CD + CF(N-2) + 2BG(N-2) \\ 2D^2 + F^2(N-1) \\ \frac{1}{2}F^2 + 4DE + 2FG(N-2) \\ 4DF + 2EF + F^2(N-2) + 2FG(N-2) \\ \frac{1}{2}F^2 + 2EF + 4GD + 4FG(N-3) \end{pmatrix} \tag{37}$$

with the boundary condition (for $t = T$) given by

$$\begin{pmatrix} A_T \\ B_T \\ C_T \\ D_T \\ E_T \\ F_T \\ G_T \end{pmatrix} = \begin{pmatrix} 0 \\ a-c \\ 0 \\ -b \\ 0 \\ -b \\ 0 \end{pmatrix} \tag{38}$$

Figure 1: Parameters of the value function

This figure plots the parameters of the value function in the basic model, when parameters are set to $a = b = 1$, $c = 0$, and $\theta = 1$. The value function is given by equation (6):

$$V_t^i(q_i, q_j) = A_t + B_t q_i + C_t q_j + D_t q_i^2 + E_t q_j^2 + F_t q_i q_j$$

and the figure below shows how each of its parameters change over time.

The parameters can be thought of either as initial value functions for games with different horizons (so T is on the horizontal axis), or as continuation values within a particular game (so t is on the horizontal axis). Due to the Markov structure, these two interpretations are identical.

One can see that as the horizon becomes longer all parameters, except for the constant A_t, approach zero. A_t converges to approximately 0.0925. Thus, for games with long horizon the values converge to 0.0925, which are about 20% lower than the static Cournot profits of $\frac{1}{9}$.

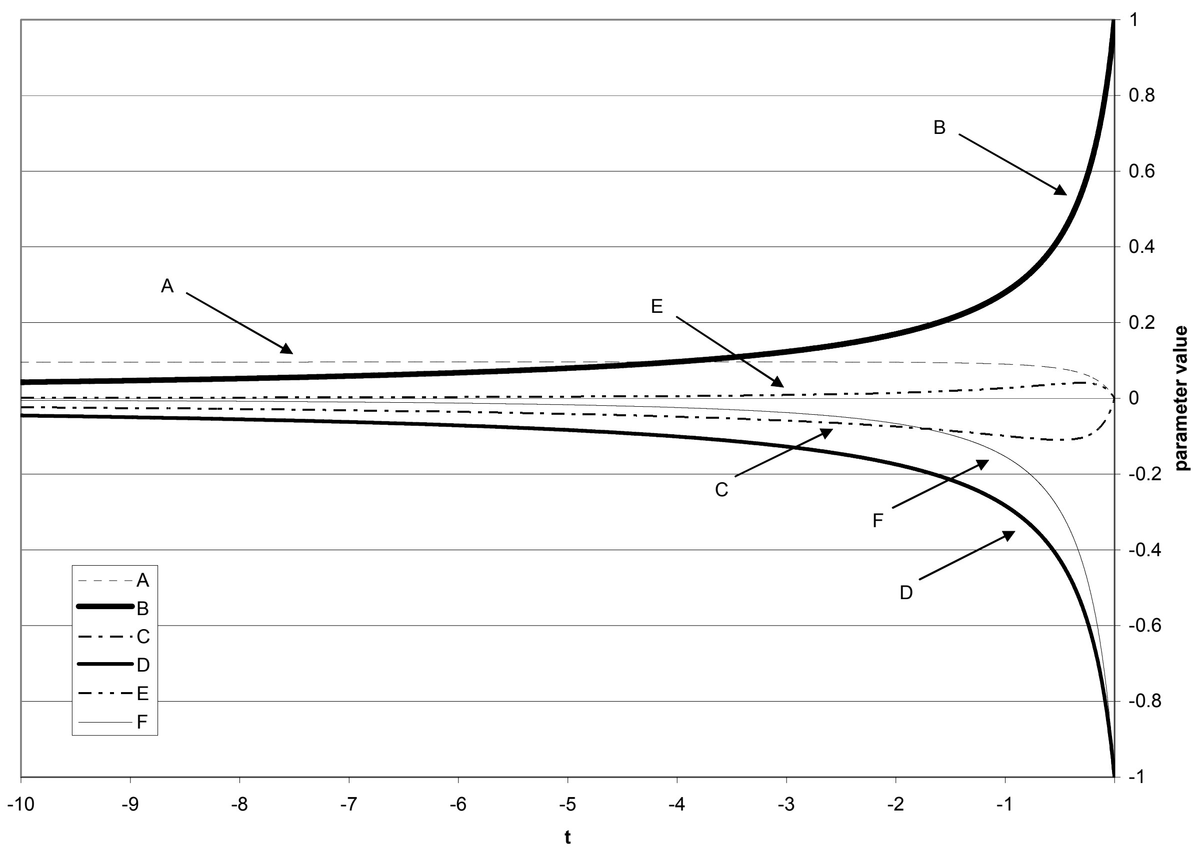

Figure 1: Parameters of the value function

Figure 2: Equilibrium path with different (symmetric) initial actions

This figure plots the equilibrium path in the basic model, when parameters are set to $a = b = 1$, $c = 0$, and $\theta = 1$. It does so for different values of initial production plans: $\frac{1}{3}$ (the Cournot level), 0.37, 0.4, 0.43, 0.46, and 0.5 (the Stackelberg level). All paths are using a symmetric case, in parameters and in initial actions, so equilibrium path is identical for both players.

Clearly, equilibrium paths of the different cases do not cross each other. Note, however, the final production levels are much closer (around 0.37 in all cases) to each other compared to the initial production plans. Note also that the equilibrium path is non-monotonic when initial actions are sufficiently low (less than about 0.44 in this case), with the peak being closer to the deadline as the initial actions are lower. When initial actions are higher, equilibrium path is monotone, but the rate of decrease in production targets is much higher towards the deadline, due to the commitment effect.

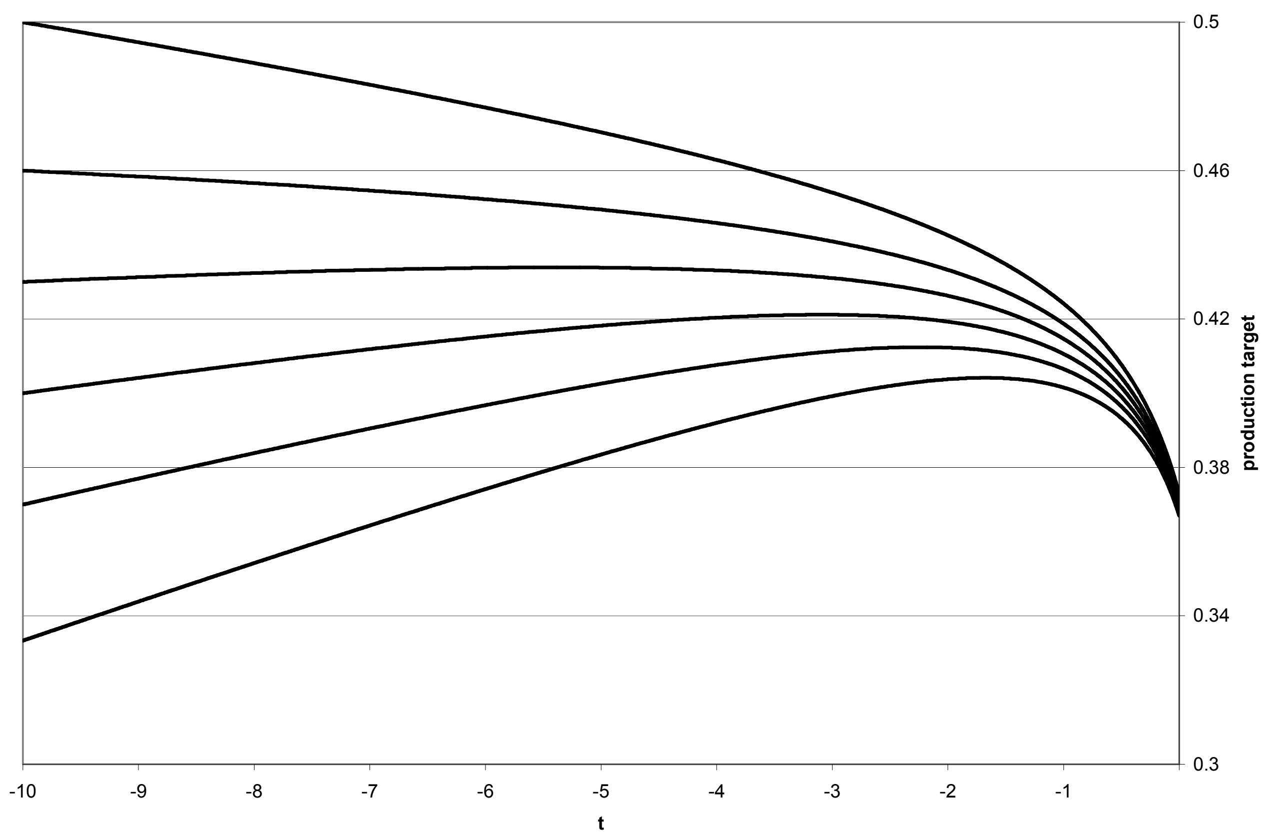

Figure 2: Equilibrium path with different (symmetric) intial actions

Figure 3: Equilibrium path with different horizons

This figure plots the equilibrium path in the basic model, when parameters are set to $a = b = 1$, $c = 0$, and $\theta = 1$, and initial production plans are $\frac{1}{3}$ (the Cournot level). It does so for different horizons: 100, 50, 10, and 1.

As can be seen, as the horizons gets longer, players have more time to smooth out their production target increase, therefore peaking at a higher level. Once the deadline gets closer, however, this higher build-up declines faster, leading to increase in cost. Final production levels do not change by much, unless the horizon is very short (as is the case when $T = 1$).

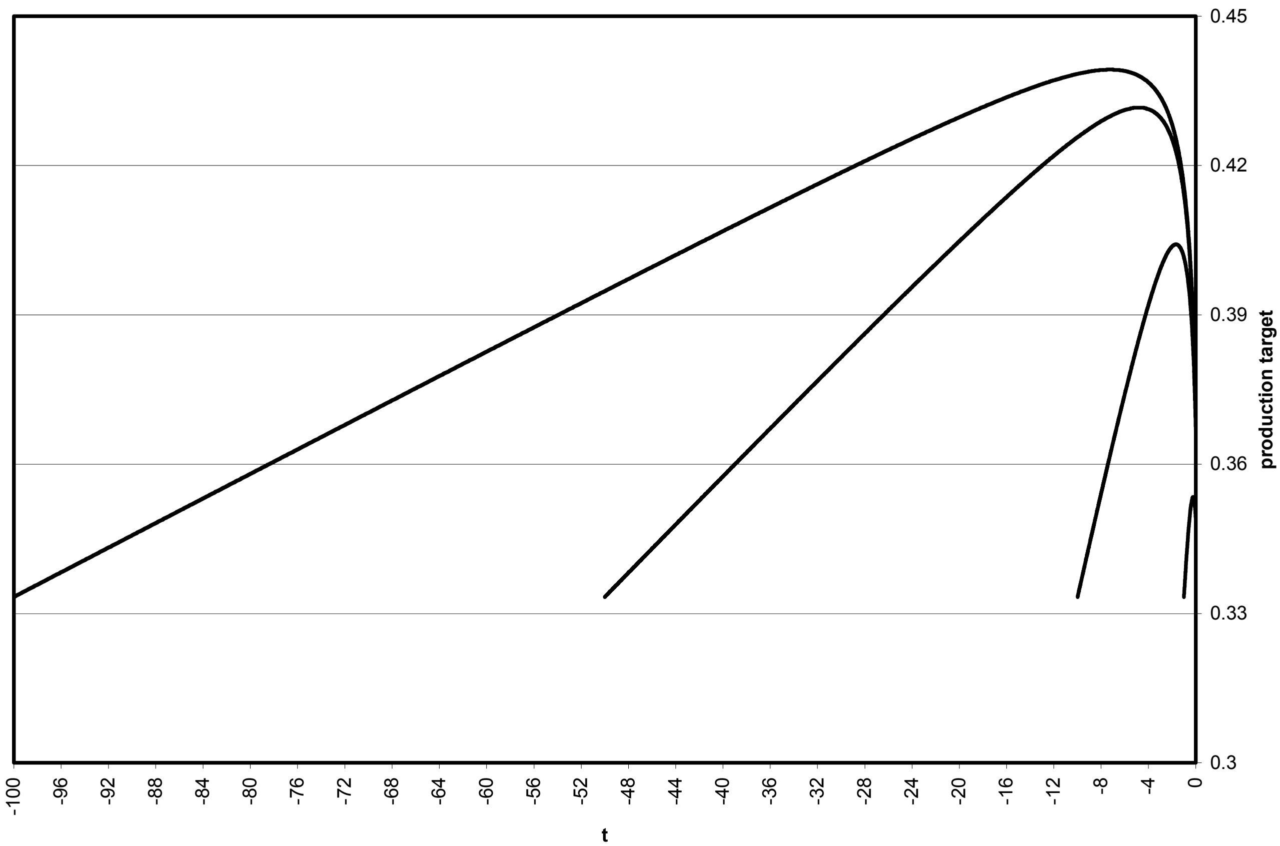

Figure 3: Equilibrium path for different horizons

Figure 4: Equilibrium path with different adjustment cost parameters

This figure plots the equilibrium path in the basic model, when parameters are set to $a = b = 1$, $c = 0$, and initial production plans are $\frac{1}{3}$ (the Cournot level). It plots different cases for the adjustment cost parameters, θ (0.2, 0.5, 1, and 5).

From the figure, one can a get sense of the non-monotonicity of the commitment effect. When adjustments are costly, the commitment effect is greater, but increasing production targets to higher levels are more costly. Depending on the size of adjustment costs, one effect or the other is more dominant.

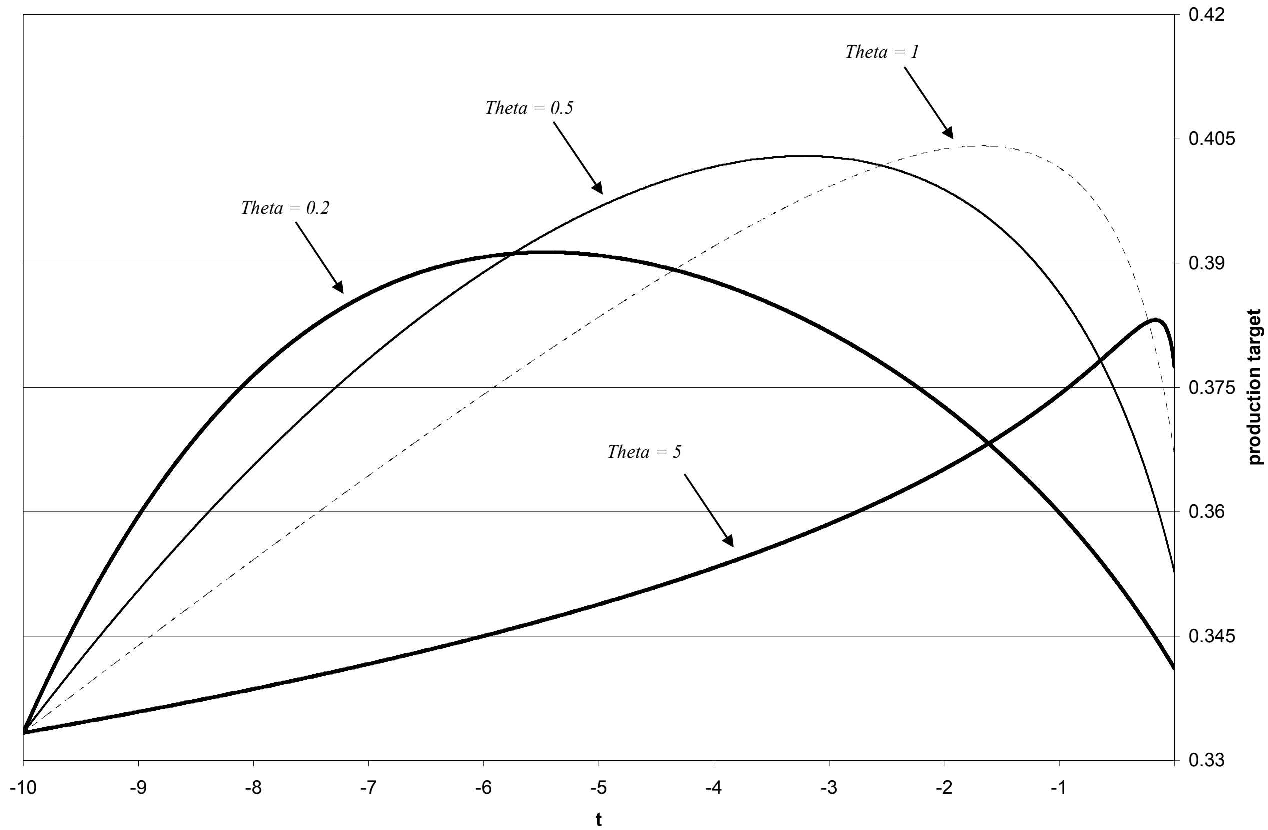

Figure 4: Equilibrium path with different adjustment cost parameters

Figure 5: Illustration of how the best response functions change as a result of adjustment costs

This figure sketches the dynamic effect of adjustment costs in the context of a two-period model. The solid lines are the static best response functions. The dashed lines are the best response functions when production targets are higher than the Cournot level. Due to adjustment costs, the best response function rotates at the level of the production target, and becomes less responsive to the opponent's action. The new equilibrium is therefore given by the intersection of the two dashed lines, giving rise to production levels which are more competitive, namely higher than the Cournot level.

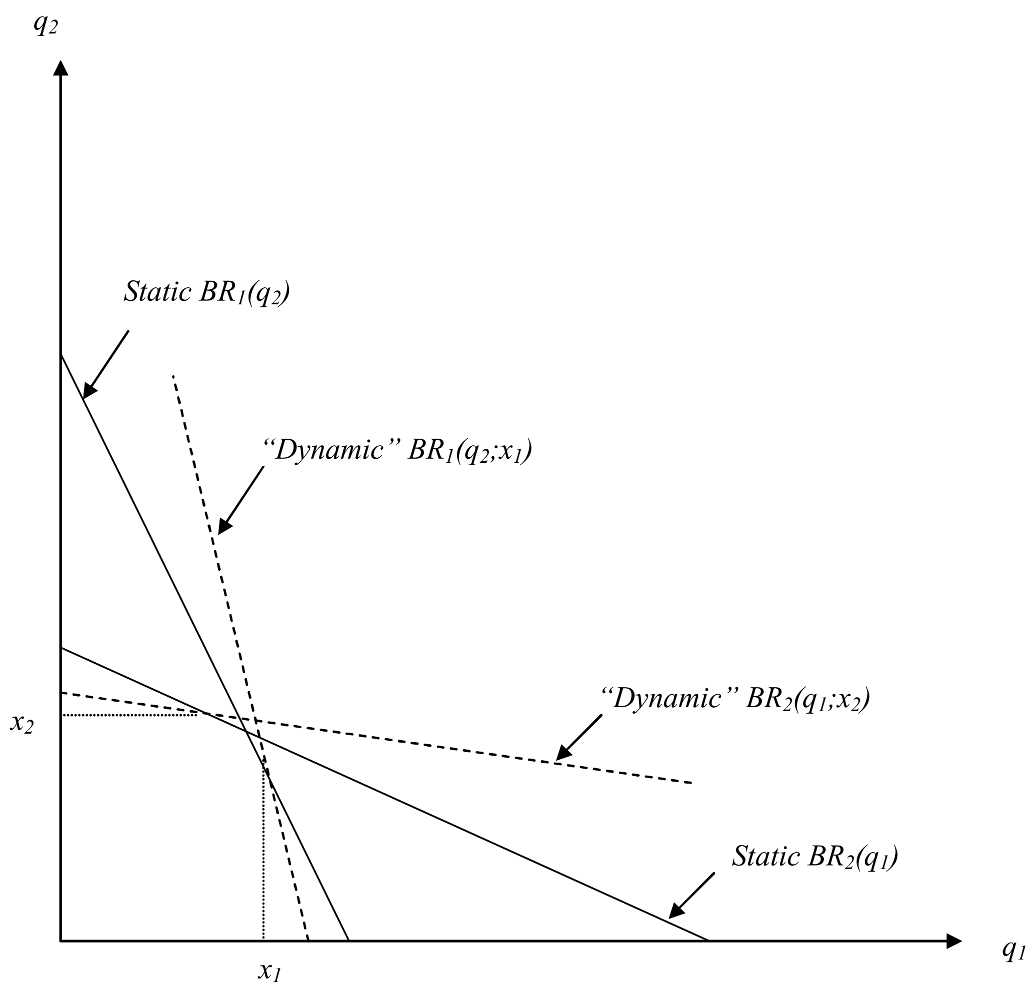

Figure 5: Illustration of how the best response functions change as a result of adjustment costs

Figure 6: Asymmetric players with different marginal costs

This figure plots the equilibrium path in a two-player model with asymmetric players. Parameters are set to $a = b = 1$, and $\theta = 1$. One player has zero marginal costs ($c_1 = 0$), while the other has positive marginal costs ($c_2 = 0.2$).

The figure plots three different cases, for different initial production plans. As players are asymmetric, each case has two paths, one for each player. The thin solid lines present the case where initial production plans are at the Cournot level ($q_1 = 0.4$, $q_2 = 0.2$). The dashed lines present the case of a reversed initial production plans ($q_1 = 0.2$, $q_2 = 0.4$), and the thick solid lines present the case of identical initial plans ($q_1 = q_2 = 0.3$).

As the horizon is reasonably long, in all cases the lower marginal cost player (player 1) eventually gain higher market share. Her market share is higher the higher is her initial production plan. It is somewhat interesting to note that player 2 ends up producing (slightly) less than her Cournot level in one of the cases (0.195 compared to her Cournot level of 0.2).

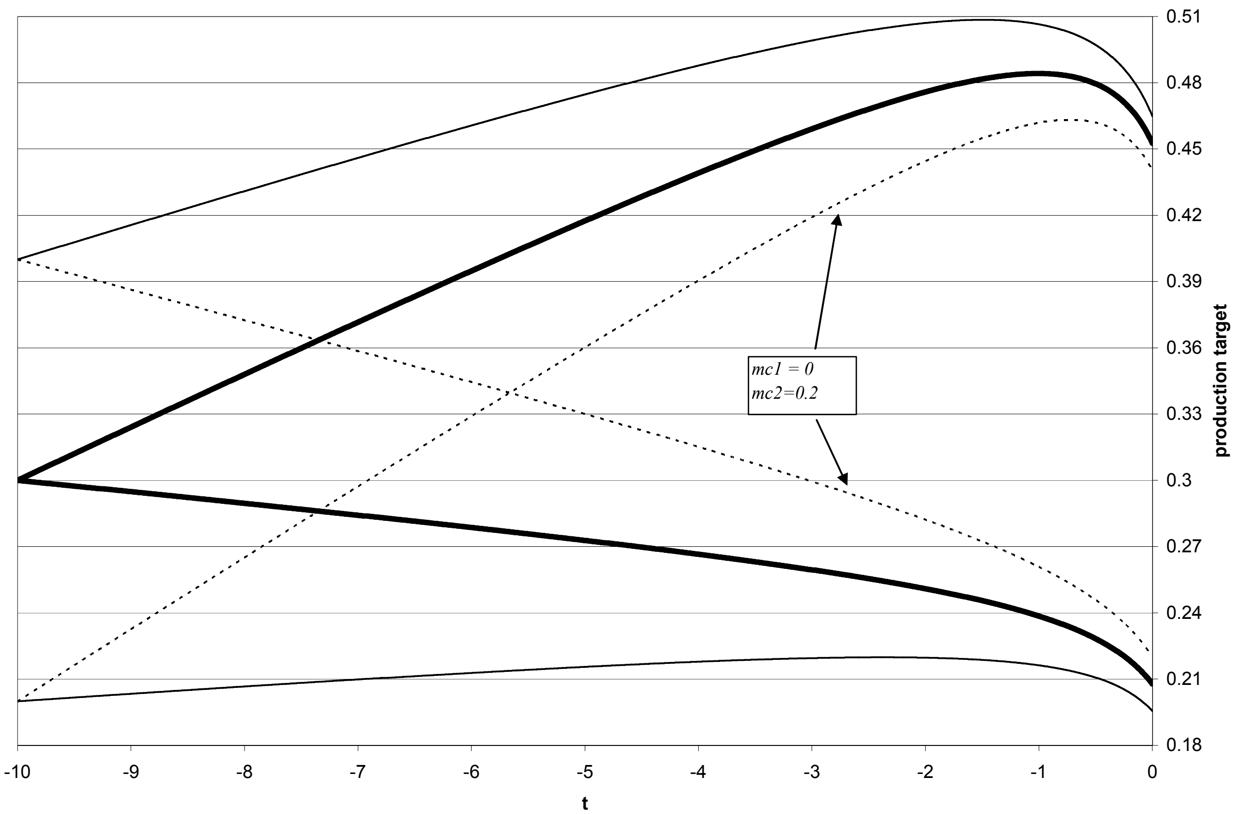

Figure 6: Asymmetric players with different marginal costs

Figure 7: Asymmetric players with different adjustment costs

This figure plots the equilibrium path in a two-player model with asymmetric players. Parameters are set to $a = b = 1$, $c = 0$, and $\theta = 1$. One player (player 1), however, has higher adjustment cost parameter than her opponent. This allows her to more credibly commit to higher production levels. The figure presents three different cases, all of them with initial production plan set at the Cournot level.

As can be seen, in all cases, the higher adjustment costs of player 1 gives her a commitment advantage, so she takes a Stackelberg leadership position in the market. In the limit, if she could fully commit, she would choose the Stackelberg level of 0.5. The second player reacts by reducing her production target, making her look more like a Stackelberg follower. It is somewhat interesting that in this case once could get non-monotonic path that goes in the reverse direction for player 2.

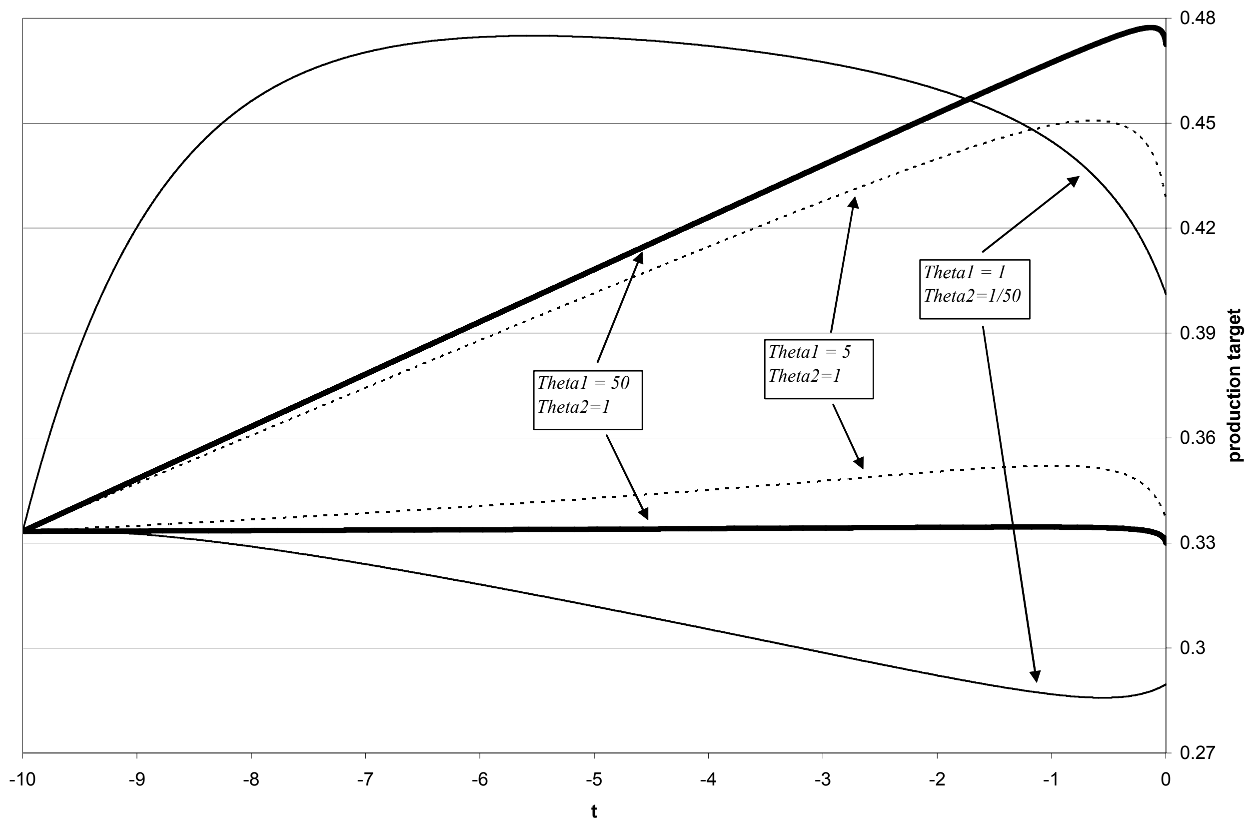

Figure 7: Asymmetric players with different adjustment costs

Figure 8: Sketch of an equilibrium path of a "nested" repeated game

This figure sketches the equilibrium path of a "nested" linear-quadratic repeated game, which we describe in Section 3. Given some stationary equilibrium of the infinitely repeated game, the one-period value function will serve as the boundary condition for the preparation stage. This will give rise to the pattern of production targets described in the figure. With adjustment costs in setting initial production plans, players will never fully adjust to the level of the peak production target, thereby giving rise to the non-monotonic path within each preparation stage.

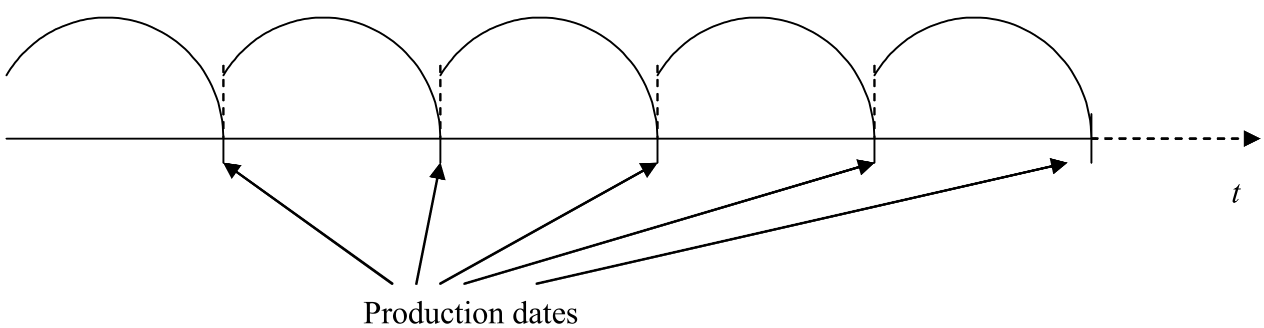

Figure 8: Sketch of an equilibrium path of a "nested" repeated game

Figure 9: Frequency and timing of production target observations

This figure provides information about the timing of the observations available. Recall, there are 372 production months in the data. Thus, one can see that starting at around four months before production, observations are available at least on a monthly basis, typically at the last week of the month. Earlier observations about production targets are not as regular, with very few coming more than six months prior to actual production.

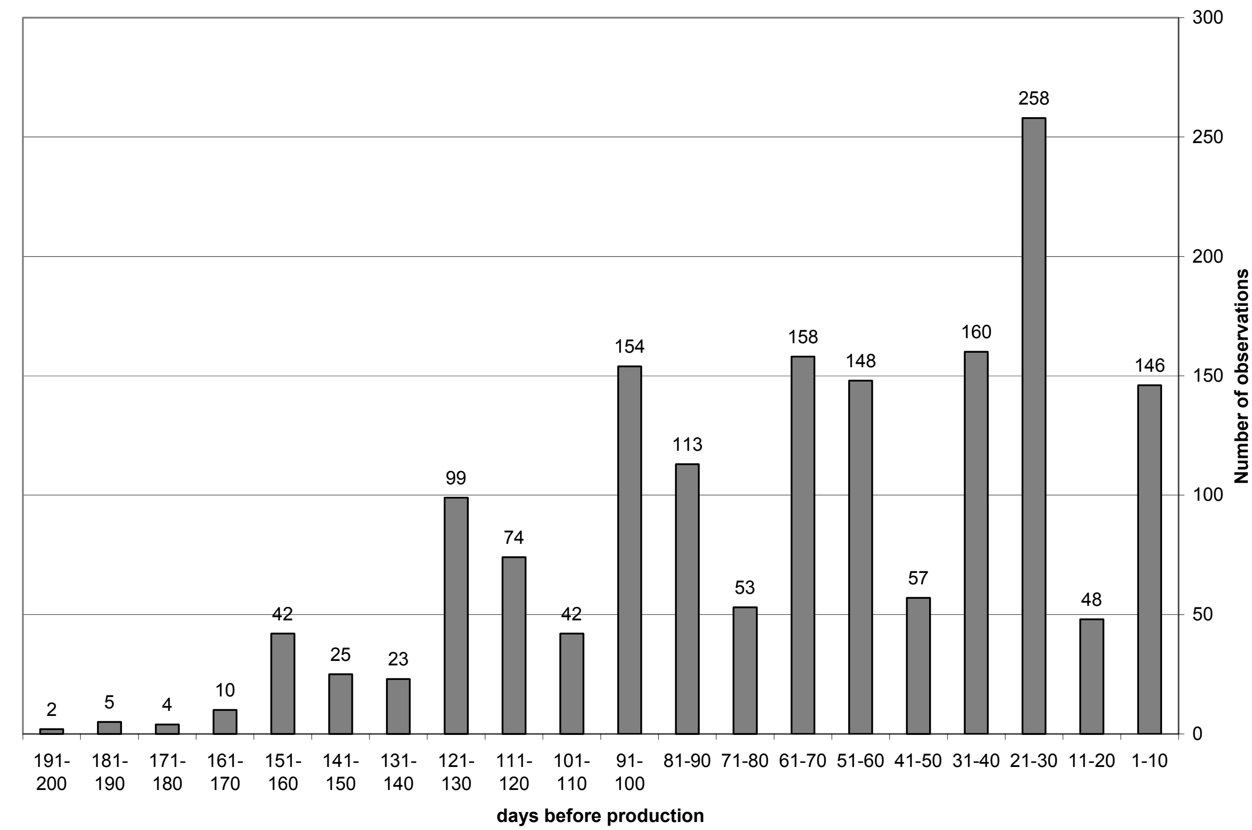

Figure 9: Frequency and timing of production target observations

Figure 10: Kernel regressions for the major three manufacturers (1965-1995)

This figure presents quartic (biweight) kernel regressions of production targets, measured according to equation (30), as a function of the number of days before production. It does so for each of the major three manufacturers (GM, Ford, and Chrysler), as well as for the (unweighted) average ("Big3"). Each series is based on 1,620 observations. Standard errors for these estimates are reported in Figure 12. All estimates use bandwidth of 30 days.

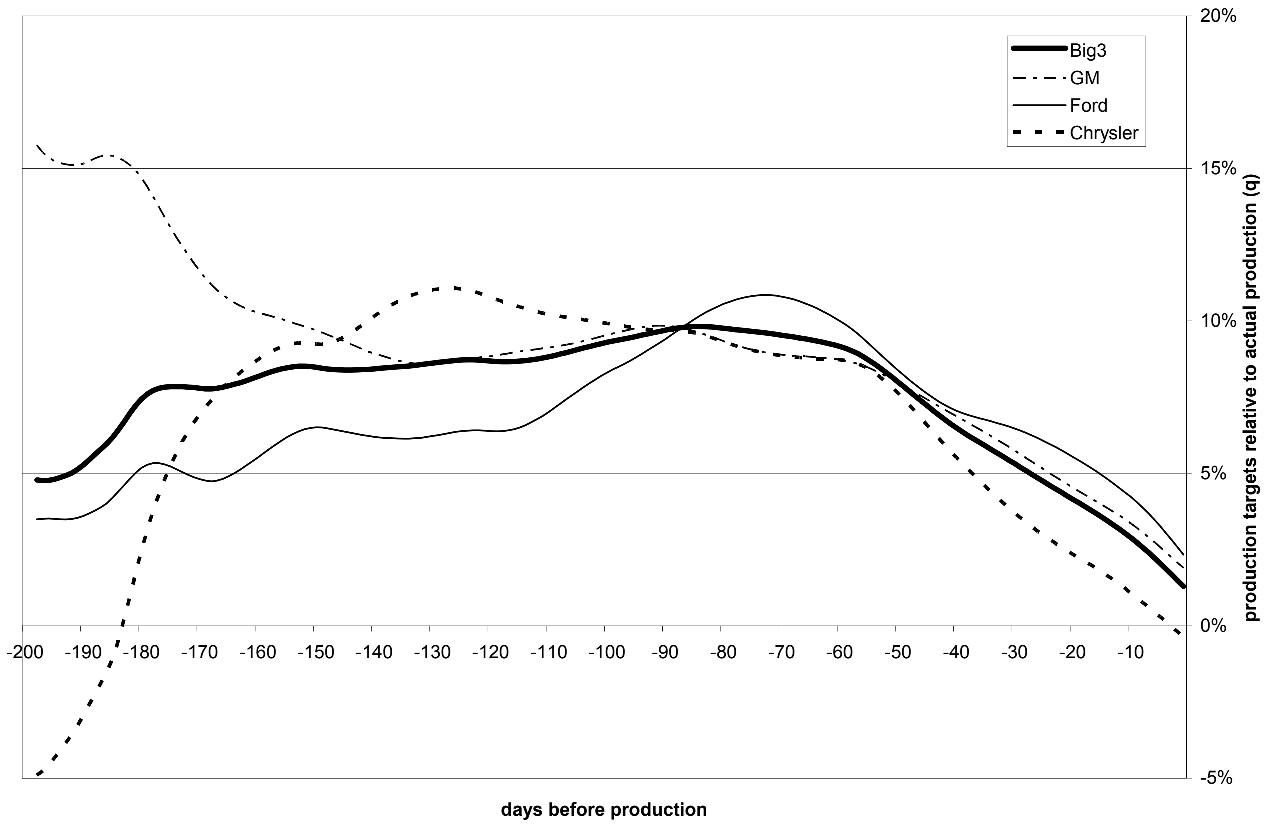

Figure 10: Kernel regressions for the major three manufacturers (1965-1995)

Figure 11: Kernel regressions for the major four manufacturers (1965-1987)

This figure, just as the previous one, presents quartic (biweight) kernel regressions of production targets, measured according to equation (30), as a function of the number of days before production. Since the data about AMC production targets span shorter observation period, this figure also reports the (unweighted) average for all four manufacturers ("Big4"), as well as for the major three ("Big3*"), for the *same* observation period. The other "Big3" series covers the whole sample, is identical to the one presented in the previous figure, and is shown for comparison. The rest of the series are based on 1,114 observations (compared to 1,620) and 270 production months (compared to 372). Standard errors for these estimates are reported in Figure 12. All estimates use bandwidth of 30 days.

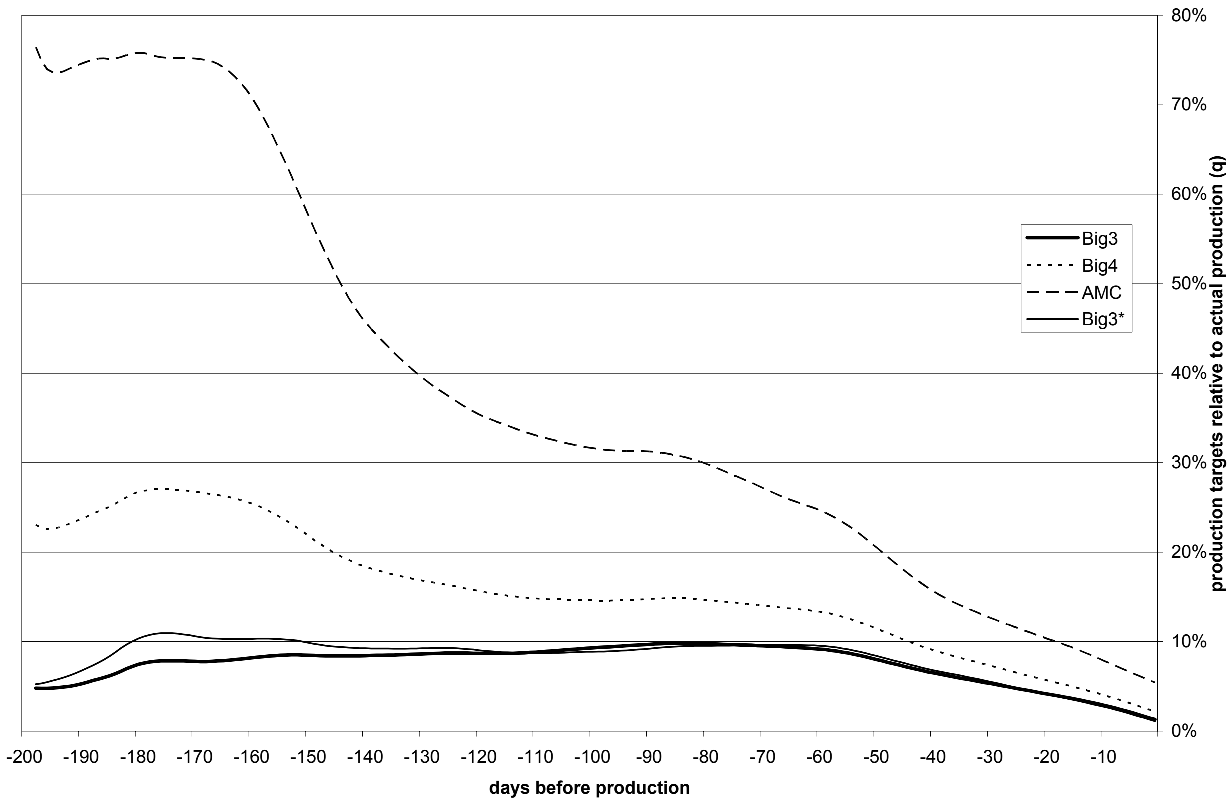

Figure 11: Kernel regressions for the major four manufacturers (1965-1987)

Figure 12: Kernel estimates with confidence intervals

This figure provides 95 percent confidence intervals for the estimates reported in Figure 10 and Figure 11. Confidence intervals are computed by bootstrapping the data, and running the same kernel regression on each bootstrapped sample. The dashed lines in each figure report the point-by-point 2.5 and 97.5 percentiles, while the solid line reports the 50th percentile, which is approximately the same as the estimates reported above in Figure 10 and Figure 11.

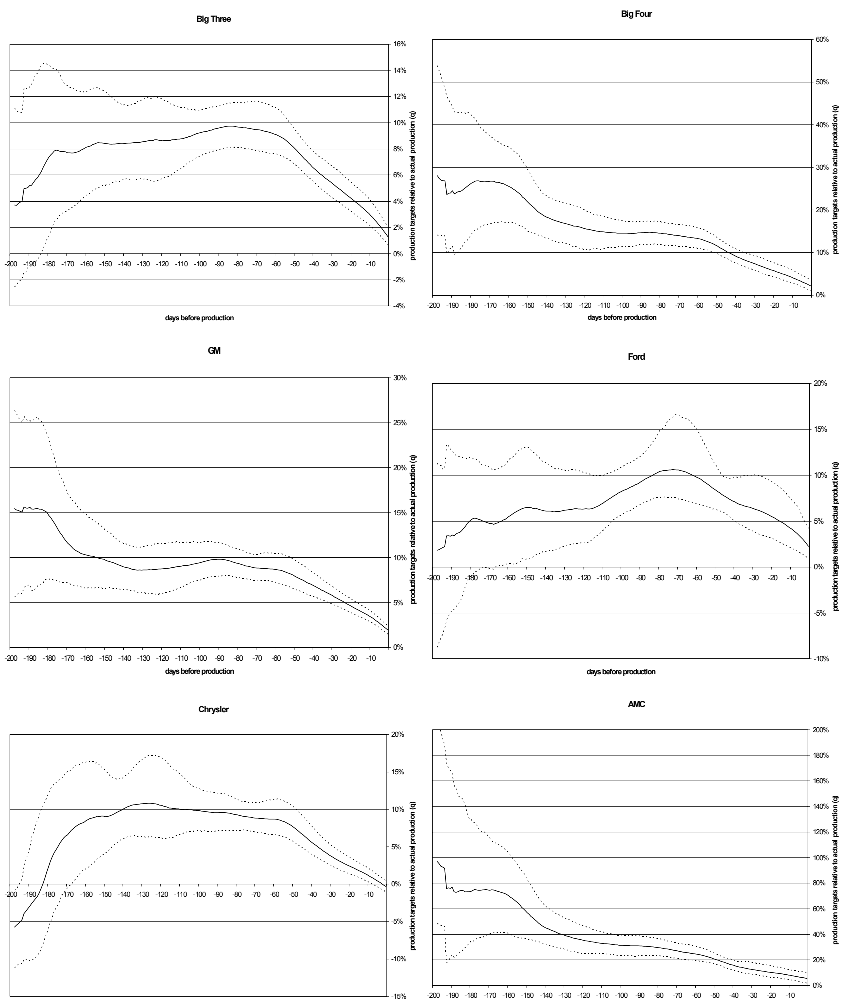

Figure 12: Kernel estiamtes with confidence intervals