

Production Targets*

Guillermo Caruana
CEMFI
caruana@cemfi.es

Liran Einav
Stanford University and NBER
leinav@stanford.edu

June 20, 2005

Abstract

We present a dynamic quantity setting game, where players may continuously adjust their quantity targets, but incur convex adjustment costs when they do so. These costs allow players to use quantity targets as a partial commitment device. We show that the equilibrium path of such a game is hump-shaped and that the final equilibrium outcome is more competitive than its static analog. We then test the theory using monthly production targets of the Big Three U.S. auto manufacturers during 1965-1995 and show that the hump-shaped dynamic pattern is present in the data. Initially, production targets steadily increase until they peak about 2-3 months before production. Then, they gradually decline to eventual production levels. This qualitative pattern is fairly robust across a range of similar exercises. We conclude that strategic considerations play a role in the planning phase in the auto industry, and that static models may therefore under-estimate the industry's competitiveness.

KEYWORDS: Differential games, adjustment costs, Cournot, quantity competition, dynamic oligopoly games.

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

*An earlier version of this paper had the title "Quantity Competition with Production Commitment: Theory and Evidence from the Auto Industry." We thank Susan Athey, Tim Bresnahan, Uli Doraszelski, Ken Judd, Ignacio Palacios-Huerta, Chris Snyder, and especially Esteban Rossi-Hansberg for helpful conversations, and seminar participants at IO Fest 2004, Bilbao, Carlos III, CEMFI, ECARES, and Olin School of Business for comments. We are extremely thankful to Maura Doyle and Chris Snyder for so kindly sharing their data with us. All remaining errors are ours. 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 entry, exit, or production are normally the result of a long preparation process. If plans cannot be hidden from competitors and changing them is costly, incentives to behave strategically during the preparation stage should be explicitly considered, as they may be an important determinant of the final equilibrium outcomes.

Consider, for example, the automobile industry. Suppose that, ahead of time, an auto manufacturer has planned a certain production target. In order to achieve it, the firm needs to take certain actions, 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 actions. To the extent that such preparations are not or cannot be fully hidden from competitors, they may play a strategic role. Given the costly nature of these adjustments, the preparation stage acts as 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.

The main goal of the paper is to develop this argument in the context of a quantity setting game, and to establish its empirical relevance using data from the U.S. auto industry. The first part of the paper constructs a dynamic quantity setting game with a planning phase and adjustment costs. In the second part, we use data on monthly production targets by the Big Three auto manufacturers – General Motors, Ford, and Chrysler – and show that the empirical pattern is consistent with the theoretical prediction.

The paper makes three separate contributions. First, we present new theoretical predictions for quantity setting games regarding the non-monotone evolution of production targeting. Since the framework is fairly simple and general, these predictions may be relevant in a wide range of strategic interactions. Second, we present empirical evidence that shows a similar non-monotonic pattern of production targets in the U.S. auto industry. Since this is one of the largest industries in the U.S., we think that documenting this pattern is of interest, even in the absence of the underlying theoretical framework. Finally, the match between the theory and the data suggests two important implications for the auto industry: (i) adjustment costs and strategic considerations may play an important role in the planning phase of production; and (ii) static models may underestimate the competitiveness of the industry.

Section 2 contains the theoretical part of the paper. We first present a benchmark model. At some specified date in the future two symmetric firms engage in Cournot competition. At date zero, each firm inherits a production structure, which serves as its initial production target. From that point onwards, each firm can make continuous adjustments to its future production structure, but incurs convex adjustment costs every time it does so. When inherited 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 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 rest of Section 2 extends the benchmark model along several dimensions and shows that all these extensions retain the same qualitative predictions. We allow for more than two players, various forms of asymmetries between players, time-varying adjustment costs, and uncertainty (common across players). We then nest the benchmark model as the stage game of an infinitely repeated game. We solve for the Markov Perfect Equilibrium of this game, and show that its stationary equilibrium path exhibits the same non-monotonic pattern. Moreover, the repeated game provides a natural way to endogenize the initial production plans, which are taken as given in the benchmark model. It also takes the model one step closer to the reality of the empirical application we study later in the paper.

There are three key assumptions that are important for our results. First, control variables are strategic substitutes, leading to a commitment incentive. Second, adjustment costs are convex, so commitment advantage monotonically increases with planned production levels. Third, all the payoffs (net of adjustment costs) are collected in the end, leading to strong competitive effects once the production date is sufficiently close. Other assumptions, we believe, are less important. For example, all the results are obtained using a linear-quadratic structure. Namely, with linear demand, constant marginal costs, and quadratic adjustment costs. This is done for tractability, as solving for the equilibrium outside of the linear-quadratic framework is not feasible. Moreover, linear-quadratic games can be viewed as second-order approximations to more general games. We could also accommodate asymmetric costs, upwards and downwards, without affecting the results, but this again would take us out of the linear-quadratic framework.¹

The model we present is a model of endogenous commitment and is therefore related to Caruana and Einav (2005), in which we mainly focus on discrete decisions, such as entry and exit. The current work is also close to the dynamic quantity competition literature (Cyert and DeGroot, 1970; Hanig, 1986; Fershtman and Kamien, 1987; Maskin and Tirole, 1987; Reynolds, 1987 and 1991; Lapham and Ware, 1994; and Jun and Vives, 2004). These papers focus on the stationary equilibrium of an infinite-horizon model (or on the limit of a finite-horizon one, as

¹Saloner (1987) and Romano and Yildirim (2005) study an extreme two-period version of such a model, in which adjustment costs upwards are free while adjustment costs downwards are infinitely costly. Unfortunately, this extreme version gives rise to a wide range of equilibria, and therefore does not provide sharp predictions.

the horizon tends to infinity); they typically find that (when actions are strategic substitutes) the stationary equilibrium is more competitive than its static analog, as players engage in a “Stackelberg warfare.”² Our model shares this feature, but unlike this literature our main focus is on the non-stationary dynamic pattern of the planning phase. One advantage in studying the dynamics of the planning phase is its strong non-stationarity; it provides clear testable prediction with respect to an observed and exogenous state variable, namely time. Stationary dynamic models are much harder to test, as the static benchmark is typically not available (for example, marginal costs are typically not observed).

Section 3 tests the predictions of the model using data on monthly production targets by the Big Three auto manufacturers in the U.S. during 1965-1995. These production targets are published in a trade journal approximately every month starting as early as six months before production. We normalize production targets by subsequent production, pool production targets from different production months, and estimate a kernel regression in order to describe the evolution of these targets as the production date gets closer. The results show that, on average, production targets exhibit a non-monotonic pattern, which is consistent with the theoretical prediction. Early targets, about six months prior to production, overstate eventual production by about five percent. Then they start to slowly increase, until they peak at ten percent about 2-3 months before production. At this point, they start to gradually decline towards the eventual production levels. This result is robust to alternative measurements and across different subsamples.

The end of Section 3 is devoted to a careful discussion of the relationship between the data analyzed and the theory previously developed. First, we discuss potential sources of adjustment costs in the production planning phase of the industry. In particular, we emphasize the nature and timing of contracts with suppliers of parts. Second, we discuss the link between the real production plans held by firms and the published figures in the study. We argue that these are likely to be very related. Finally, we discuss some relevant differences between the stylized theoretical model and the nature of competition in the industry (e.g. inventories and product differentiation), and argue that these gaps are unlikely to change the qualitative results. Thus, establishing the relationship between the empirical pattern and the theoretical predictions allow us to conclude that adjustment costs and strategic considerations play an important role in the planning phase of production and that static models may therefore under-estimate the competitiveness of the industry.

At some general level, this work can be classified within the recent empirical studies of dynamic oligopolies (e.g. Benkard, 2004; and Ryan, 2004). In contrast to these studies, which primarily focus on estimating the parameters associated with a given theoretical framework, which is assumed, our theoretical framework provides testable implications. Therefore, the primary objective here is testing the qualitative prediction of the theoretical framework. Once validated, the next obvious step, which is outside of the scope of this paper, is to parameterize the model and estimate

²This can also be viewed as a dynamic extension of a “top dog” strategy within the Fudenberg and Tirole (1984) taxonomy of strategic behavior.

structural parameters.

The data we use in this work is also used in Doyle and Snyder (1999), who investigate the role of the published production targets as an information sharing device by focusing on the positive correlation among manufacturers in the revisions to their production targets. Our results are consistent with their theoretical framework, which provides no restrictions on the way production targets evolve over time. Their results are also consistent with ours, as the model of this paper predicts that manufacturers would follow similar patterns over time, thereby creating positive correlation in revisions of production targets. Therefore, we view the two studies as complementary; 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 observations 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. We believe that any attempt to quantify either effect, by, for example, estimating structural parameters, should take both effects of strategic considerations and uncertainty into account.

2 Theory

2.1 The benchmark model

There are two players. At time $t = 0$, they start with exogenously inherited initial production plans of $(q_1(0), q_2(0))$. At all points $t \in [0, T]$ each player i chooses $x_i^t \in \mathbb{R}$, which controls the rate at which she changes her production plan, i.e. $q_i'(t) = x_i^t$. Note that x_i^t can be either positive or negative. If a player changes her plans at a rate of x_i , she has to pay adjustment costs of $c_i(x_i, t)$. At time T , and given their final plans, $q_1(T)$ and $q_2(T)$, players compete in quantities and collect final payoffs of $\pi_i(q_i(T), q_j(T))$.

In order to make the model more tractable, we use a linear-quadratic structure. Thus, we assume that inverse demand is linear, given by $p = a - b(q_1 + q_2)$, and marginal costs are constant and given by c . Thus, we have that

$$\begin{aligned} \pi_i(q_i(T), q_j(T)) &= (a - bq_i(T) - bq_j(T))q_i(T) - cq_i(T) = \\ &= (a - c)q_i(T) - bq_i^2(T) - bq_i(T)q_j(T) \end{aligned} \tag{1}$$

In addition, we assume that adjustment costs are quadratic and take the form of

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

Note that adjustment costs are constant over time,³ symmetric across players, and symmetric for

³For simplicity, there is no time discounting. Time discounting is a special case of the extension of the model to time-varying adjustment costs, which we analyze later.

positive and negative rates. None of these properties is important for the main results.

We solve for the Markov Perfect Equilibrium of the model. Thus, strategies only depend on the state variables, q_1 and q_2 and time t . Let $V_i^t(q_i, q_j)$ be the value function for player i at time t , with state variables q_i and q_j . If $V_i^t(q_i, q_j)$ exists and is continuous and continuously differentiable in its arguments, then it satisfies the following Bellman equation

$$\max_{x_i^t} \left(-\frac{\theta}{2} (x_i^t)^2 + \frac{\partial V_i^t}{\partial q_i} x_i^t + \frac{\partial V_i^t}{\partial q_j} x_j^t + \frac{\partial V_i^t}{\partial t} \right) = 0 \quad (3)$$

The first order condition for x_i^t implies that

$$x_i^t = \frac{1}{\theta} \frac{\partial V_i^t}{\partial q_i} \quad (4)$$

We can now substitute this back into equation (3), and obtain the following differential equation

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

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

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

which, using equation (4), implies that

$$x_i^t(q_i, q_j) = \frac{1}{\theta} (B_t + 2D_t q_i + F_t q_j) \quad (7)$$

Given that players are symmetric, we can substitute equations (6) and (7) into equation (5) and obtain

$$\begin{aligned} 0 = & \frac{1}{2\theta} (B_t + 2D_t q_i + F_t q_j)^2 + \frac{1}{\theta} (C_t + 2E_t q_j + F_t q_i) (B_t + 2D_t q_j + F_t q_i) + \\ & + (A'_t + B'_t q_i + C'_t q_j + D'_t q_i^2 + E'_t q_j^2 + F'_t q_i q_j) \end{aligned} \quad (8)$$

This is a polynomial in q_i and q_j . Since it has to be satisfied for all values of q_i and q_j , all its six coefficients (which are functions of t) have to be equal to zero. This gives the following set of ordinary differential equations. To ease notation, we can just think of time as going backwards.

⁴See Kydland (1975), who shows uniqueness for a discrete-time version, and Lukes (1971), Papavassilopoulos and Cruz (1979), and Papavassilopoulos and Olsder (1984) for analysis of existence and uniqueness in finite-horizon linear-quadratic differential games.

This is convenient as our boundary condition is for $t = T$. Thus, all derivatives with respect to time (A' , B' , etc.) reverse signs, and the law of motion for the parameters is given by

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

with boundary condition (for $t = T$)

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

which is provided by the profit function in equation (1).

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 cannot be solved analytically, so we approximate 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 $p = 1 - q_1 - q_2$. 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 this game involves each player producing her Cournot quantity of $q = \frac{1}{3}$, while the Stackelberg leader and follower production levels are $q = \frac{1}{2}$ and $q = \frac{1}{4}$, respectively.

Figure 1 shows how the parameters of the (symmetric) value function, as given in equation (6), evolve over time. As the horizon becomes longer (i.e. as $T \rightarrow \infty$) A_0 converges to approximately 0.0925 and all other parameters approach zero. Thus, for games with long horizon the equilibrium profits converge to 0.0925, which are approximately 17% lower than the static Cournot profits of $\frac{1}{9}$ (14% is due to higher production and lower equilibrium prices, while 3% is due to adjustment

⁵One 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.

costs). 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 the game in which both players inherit an initial production plan at the static Cournot level. 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 realize that they are sufficiently committed to high output, but that they are much above their static best responses, and optimally decide to gradually adjust towards it. Given that adjusting is costly, the parties do not adjust all the way to the static Nash equilibrium.⁶ In this particular example, the equilibrium outcome is about 0.37, compared to the static outcome of $\frac{1}{3}$. Finally, we also depict one off-equilibrium-path strategy for each player. Suppose that player i receives an unexpected shock to her intended plan at $t = T - 4$ and has her plan reverted to the Cournot level. Both players realize that player j has achieved a leader position in the market. Player j capitalizes on this advantage by increasing her own plans even further. Meanwhile, player i 's best response is to rebuild its size. Nevertheless, the advantageous position acquired by player j never fully diminishes and is kept until the production date.

Figure 3 presents the symmetric equilibrium path for different initial production plans. If these are not too high, one observes the same pattern as in the previous figure. If initial production plans are sufficiently high (greater than about 0.44 in this particular example), both parties are sufficiently committed to high production from date zero and do not 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 they speed up adjustments in the direction of their static best response levels.⁷

Figures 4 and 5 present comparative statics with respect to the length of the horizon and with respect to the size of the adjustment cost parameter. An inspection of equation (9) reveals that these two exercises are similar. A proportional increase in the adjustment cost can be viewed as a slowdown in the evolution of the value function. Loosely speaking, it is a horizontal stretch of Figure 1. Thus, changes in the adjustment cost parameter are similar to a rescaling of time.⁸ Figure 4 shows how the length of the horizon affects the equilibrium path. As the horizon

⁶With convex adjustment costs, the optimal strategy always leads to partial adjustments. This is because the static profit function is flat at the static best response level. Thus, the marginal cost of adjustment is zero for small adjustments and higher for greater ones, while the marginal benefit is strictly positive for small adjustments but zero for full adjustments.

⁷Note that if the initial targets were very low and the adjustment parameters high, one could also see a fully increasing equilibrium path.

⁸It is similar but not identical. Think of the game in discrete time. A lower θ is similar to increasing the length of a period, without changing the number of periods. Increasing T is similar to increasing the number of periods, without changing their length. Thus, loosely speaking, stretching of time allows for more opportunities to adjust

gets longer, there is more time to build up commitment. Similarly, Figure 5 shows that as the adjustment costs decrease, building commitment becomes cheaper. In both cases this leads to higher targets and an ultimate faster decline.

2.3 Intuition from a two-period model

The key qualitative prediction of the model, namely that players have an incentive to exaggerate their production intentions as a way to achieve commitment, can be obtained within the context of a simple two-period model. Suppose that firms start with inherited production targets of y . At $t = 1$ they can revise their plans to z_1 and z_2 , but pay a quadratic adjustment cost when they do so. Then, in period $t = 2$ firms have a final opportunity to revise the quantities they want to produce and set them to q_1 and q_2 , paying the corresponding adjustment costs. Given these production levels, market price is given by $p = 1 - q_1 - q_2$. There is no discounting, so payoffs are the final Cournot profits (with zero marginal costs) minus any adjustment costs incurred in the process.

We can solve for the Subgame Perfect Equilibrium of the game using backward induction. In period $t = 2$ each player i chooses q_i to solve

$$\max_{q_i} (1 - (q_i + q_j))q_i - \frac{\theta}{2}(q_i - z_i)^2 \quad (11)$$

Best response functions are

$$q_i = \frac{1 - q_j + \theta z_i}{2 + \theta} \quad (12)$$

and the second period equilibrium strategies are

$$q_i(z_i, z_j) = \frac{1 + \theta(1 + (2 + \theta)z_i - z_j)}{(\theta + 3)(\theta + 1)} \quad (13)$$

One can easily observe that if firms target the Cournot quantities, $z_i = z_j = \frac{1}{3}$, then setting $q_i = z_i$ 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 (see Figure 6). Each player's response to a change in her opponent's quantity is not as strong as in the absence of adjustment costs. Thus, if $z_i = z_j$ are greater (less) than $\frac{1}{3}$ the players end up adjusting in the direction of their static best responses, but not fully, thereby ending up in a more (less) competitive equilibrium.

In period $t = 1$ firms choose z_i and z_j , accounting for the equilibrium strategies at $t = 2$. Thus, each player i chooses z_i to solve

$$\max_{z_i} (1 - q_i(z_i, z_j) - q_j(z_i, z_j))q_i(z_i, z_j) - \frac{\theta}{2}(q_i(z_i, z_j) - z_i)^2 - \frac{\theta}{2}(z_i - y)^2 \quad (14)$$

implying the following first order condition for each player:

$$\frac{\partial q_i}{\partial z_i}(1 - q_i - q_j) - q_i \left(\frac{\partial q_i}{\partial z_i} + \frac{\partial q_j}{\partial z_i} \right) - \theta(q_i - z_i) \left(\frac{\partial q_i}{\partial z_i} - 1 \right) - \theta(z_i - y) = 0 \quad (15)$$

behavior.

This yields a solution $z(y, \theta)$ and $q(y, \theta)$.⁹ For example, if $y = \frac{1}{3}$, i.e. firms' inherited targets are at the Cournot level, their final productions would be

$$q\left(\frac{1}{3}, \theta\right) = \frac{1}{3} + \frac{\theta}{3\theta^3 + 30\theta^2 + 78\theta + 54} \quad (16)$$

which are always above $\frac{1}{3}$ for any $\theta > 0$. When $\theta = 1$, for example, equilibrium targets at $t = 1$ are $z \approx 0.357$ and final productions are $q \approx 0.339$. Thus, the qualitative conclusions are the same as in the continuous time case: planned production levels increase first, and decrease later.

2.4 Extensions to the benchmark model

Here we present some of the most natural extensions to the benchmark model. The main message is that all of them retain the same qualitative predictions of the model. The derivations are provided in the appendix.

N players: The benchmark model is constructed for two players only for convenience. Results remain unchanged with more than two players. The value function has one additional element, $\sum_{j \neq i} \sum_{k \neq i, j} q_j q_k$, which results in an additional equation in the system of differential equations. We computed the equilibrium for different sets of parameters and the equilibrium patterns are qualitatively identical to those obtained for the two-player model.

Asymmetric players: Asymmetries among firms can be introduced either through the final payoff function (for example, firms may vary in their marginal costs) or through the adjustment costs (for example, labor may be more unionized in one firm than the other). In the appendix we treat them jointly, but we do comparative statics on each dimension separately.

Figure 7 illustrates the case of asymmetric marginal costs. In particular, it uses the same parameter values as in Section 2.2, but introduces a (constant) marginal cost of 0.2 for player 2. The figure presents the equilibrium paths for different (but symmetric) initial conditions. The general pattern is similar to the benchmark case. Now the more efficient player produces more than her opponent, and more than her static Nash equilibrium quantity ($q_1 = 0.4$ and $q_2 = 0.2$). In this case the less efficient player may produce less than her static Nash quantity. This is shown in the thin solid line. The reason for this is that asymmetric marginal costs introduce asymmetries in the commitment opportunities. Given that the more efficient player is producing more, her static payoff function is steeper around the equilibrium. This allows her to enjoy higher levels of commitment and attain a Stackelberg advantage. In all cases, however, overall quantity is higher (more competitive) than the static equilibrium level of 0.6. This might hint a welfare improvement, due to both higher consumer surplus and more efficient allocation of resources among the firms, but one has to include the adjustment costs in the analysis to obtain a definitive answer.

⁹The solution is $z(y, \theta) = \frac{4+4\theta+\theta^2+y(\theta+1)(\theta+3)^2}{(26\theta+10\theta^2+\theta^3+18)}$ and $q(y, \theta) = \frac{(y\theta+2)(\theta+1)(\theta+3)}{(26\theta+10\theta^2+\theta^3+18)}$.

Figure 8 presents the case of asymmetric adjustment costs for different values of the θ coefficients. The shape of the equilibria is the same as before. It is interesting to notice that it is the more flexible player who is able to end up producing more. When adjustment costs are high ($\theta_1 = 1$ and $\theta_2 = 5$) this is simply because player 2 cannot afford to increase her plans so rapidly (recall that initial plans and the length of the horizon are fixed in this exercise). When the costs are lower the leadership position is achieved through the higher ability of the flexible player to increase her plans further as a way to commit to high output.¹⁰

Time-varying adjustment costs: One may argue that adjustment costs may vary over time. One reason may be discounting, which would result in declining adjustment costs. It is also reasonable to consider that adjustments become more expensive as the production date gets closer. As an example, hiring 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.¹¹

It is straightforward to incorporate such effects into the benchmark model. The adjustment cost function would be

$$c_i(x_i, t) = \frac{\theta(t)}{2} x_i^2 \quad (17)$$

where no restrictions are imposed on $\theta(t)$. The derivation of the system of ordinary differential equations is the same as in equation (9), with θ replaced by $\theta(t)$. Notice that θ enters into the system in a proportional way. Therefore, replacing it by $\theta(t)$ is 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. Qualitatively, the predictions of the model remain unchanged.

Uncertainty: In the presence of uncertainty, there is a general trade-off between commitment and flexibility, as remaining flexible would allow firms to adjust to unexpected events. The precise impact of considering uncertainty within the context of this work will depend on the type of uncertainty explored. In the appendix we consider a model with a natural source of common uncertainty within the linear-quadratic framework. Suppose that final demand can be high or low depending on whether the state of the economy is either high (H) or low (L). The economy (symmetrically) fluctuates between the two states following a Poisson process: at each point, at hazard rate λ the state changes.

Initially, with the horizon far enough in the future, the current state is not particularly informative about the final state of demand. Given that firms only care about the eventual realization of demand, on equilibrium they start by having a similar behavior independently of the actual state. As the production date draws near, however, firms become more responsive to changes in

¹⁰Note that if the initial inherited positions were higher, say $q_0 = 0.4$, and the adjustment costs high as well, the previous result could be reversed. In this case the non-flexible player would be at a credible position not to change her plans far away from 0.4, which would force the flexible player to adjust downwards.

¹¹This second case is closer to the framework studied in Caruana and Einav (2005).

the state of the economy. This typically results in upwards (downwards) adjustments to production targets in response to changes into the high (low) state. As firms foresee this happening, they are more reluctant to adjust early, compared to the benchmark model, and therefore build up commitment more slowly. While the equilibrium path is random as it depends on the realization of uncertainty, the expected equilibrium path (computed numerically) exhibits a non-monotonic pattern as in the benchmark model.

2.5 Repeated interaction

Many real-world situations, like the monthly production decisions in the auto industry we study later, are repeated in nature. Here we consider an infinitely repeated game in which the benchmark model is the stage game and there are adjustment costs between stages. These costs between stages capture the fact that firms are constrained in their future plans by their actual production infrastructure.

Formally, each stage of the game is played as follows. Given last period production of (y_1, y_2) , players first decide simultaneously on their initial production plans $q_1(0)$ and $q_2(0)$ for next period, but pay a cost of $\frac{\varphi}{2} (q_i(0) - y_i)^2$ when they do so. For the next T units of time they play the benchmark model with inherited initial plans of $(q_i(0), q_j(0))$ and quadratic adjustment costs with parameter θ . That is, they can continuously adjust their production targets, paying an adjustment cost of $\frac{\theta}{2} (q'_i(t))^2$ if they do so (where t is the time elapsed since the beginning of the period). At the end of each stage, production takes place and the stage payoffs are collected. Players discount profits with a common discount factor β per period. For simplicity we assume that players do not discount payoffs within a period.

We solve for a symmetric Markov Perfect Equilibrium (MPE). Thus, the state variables are the most recent production plans and the elapsed time t . Given that the game has a linear-quadratic structure, we guess that the value function is quadratic in the state variables. We search for an equilibrium satisfying this assumption and find one, justifying the initial guess. The solution to the value function within each stage follows the same law of motion as in the benchmark model and thus satisfies equation (9). The boundary condition is different: in this case, it is determined endogenously as part of the equilibrium. In particular, there is a relationship between the value function at the beginning of the stage game and the value function at the end of it. We establish this relationship below.

In equilibrium, players set initial production plans to satisfy

$$\max_{q_i} (A_0 + B_0 q_i + C_0 q_j + D_0 q_i^2 + E_0 q_j^2 + F_0 q_i q_j) - \frac{\varphi}{2} (q_i - q_i(T))^2 \quad (18)$$

which leads to the following first order condition:

$$B_0 + 2D_0 q_i + F_0 q_j - \varphi (q_i - q_i(T)) = 0 \quad (19)$$

Equation (19), together with its analog for q_j , provides a closed-form relationship between $(q_1(0),$

$q_2(0)$) and $(q_1(T), q_2(T))$. Since, by construction

$$V_i^T(q_i(T), q_j(T)) = \pi_i(q_i(T), q_j(T)) - \beta \frac{\varphi}{2} (q_i(0) - q_i(T))^2 + \beta V_i^0(q_i(0), q_j(0)) \quad (20)$$

we can substitute the relationship between $(q_1(0), q_2(0))$ and $(q_1(T), q_2(T))$ into equation (20). As this has to be satisfied for any $q_i(T)$ and $q_j(T)$ we can equate coefficients, and obtain a system of six equations that provides a closed-form relationship between A_0, \dots, F_0 and A_T, \dots, F_T . This is the boundary condition that substitutes equation (10) of the benchmark game. The solution to equation (9) and this new boundary condition constitutes the MPE of the repeated game. Finally, we focus on the steady state of the equilibrium, in which the production decisions (but not production plans) are constant at every stage.

The equilibrium is computed by numerically searching for a solution. One starts with a guess for A_T, \dots, F_T , and then iterates the law of motion in equation (9) to obtain A_0, \dots, F_0 . Then, using the boundary condition one obtains new values for A_T, \dots, F_T . We iterate this procedure until convergence. Although, in general, one cannot establish uniqueness (or even existence) for this game, the problem seems to be well behaved. The procedure converges extremely rapidly to the same values for a wide range of initial conditions. Thus, on numerical grounds, we believe that the repeated interaction game has a unique symmetric MPE, or at least a unique symmetric linear-quadratic MPE.

In Figure 9 we show the equilibrium path for the usual benchmark parameter values ($a = b = 1$, $c = 0, \theta = 1, T = 10$), a discount factor of $\beta = 0.9$, and $\varphi = 0.1$. As one can see, the equilibrium stage pattern exhibits the same hump shape as in the benchmark model. The production levels are now higher than what would be produced in the benchmark model if the inherited plans were the ones from the steady state equilibrium. This is because, in addition to the commitment effect already described, there is a dynamic effect of commitment through the adjustment costs between stages. This second effect is the same that is present in all dynamic quantity games with sticky controls analyzed in the literature (Maskin and Tirole, 1986; Reynolds, 1987 and 1991; Jun and Vives, 2004). Its importance is diminished in this model by the fact that the planning phase provides an additional opportunity to revise production levels. Naturally, this additional dynamic effect increases with θ and decreases with T . Figure 10 provides some comparative statics with respect to the relative importance of the two types of adjustment costs by varying φ and θ . As one can observe, φ primarily affects the size of the jump between production levels and initial plans for the subsequent production period, with high values of φ implying small jumps. In contrast, θ primarily affects the shape of the production plan adjustments and final equilibrium production levels.

One important special case of the repeated game is the one in which $\varphi = 0$. In such a case, there is no link between consecutive production periods and the model collapses to the benchmark model with free initially chosen plans. That is, at $t = 0$ players decide simultaneously and costlessly on their initial plans $(q_1(0), q_2(0))$ and then continue playing as in the benchmark model. In the simultaneous-move game played at date zero players solve equation (18) (with

$\varphi = 0$), implying a unique equilibrium of

$$q_j = q_i = \frac{-B_0}{2D_0 + F_0} \quad (21)$$

These initial plans give rise to an equilibrium path, in which production plans are flat at $t = 0$ and gradually decline thereafter (see also Figure 10).¹² For any $\varphi > 0$, however, the equilibrium path presents the hump-shaped pattern emphasized throughout.

3 Evidence

3.1 Data

We use data on domestic production targets of the major auto manufacturers in the U.S. These are the same data used by Doyle and Snyder (1999).¹³ Therefore, we focus only on the dimensions of the data that are relevant for our empirical analysis; Doyle and Snyder (1999) provide descriptive statistics and further details of the data.

The unit of analysis is a production month. Prior to each production month, the Big Three U.S. auto manufacturers – General Motors (GM), Ford, and Chrysler – decide about their production targets for future months.¹⁴ These targets are posted in a weekly industry trade journal, *Ward’s Automotive Reports*, which specializes in industry data and statistics. Targets are posted approximately every month, starting as early as six months prior to actual production.

Production targets are summarized by the number of cars to be produced by each manufacturer, aggregated over all models. Thus, variation across models or the introduction of new models cannot be directly used. The data set has a panel structure and covers the years 1965 to 1995, for a total of 372 production months.¹⁵ Every time a production target is published, it includes production targets for *all* three manufacturer. Thus, manufacturers do not decide when to post their targets, as this is requested by *Ward’s*. Overall, we observe 1,621 production targets for each manufacturer.¹⁶ This amounts to an average of 4.42 production targets per production

¹²This path is initially flat because, in equilibrium, initial production plans (q_i^*, q_j^*) satisfy $\frac{\partial V_i^0(q_i^*, q_j^*)}{\partial q_i} = 0$. From equation (4), the rate of adjustment at $t = 0$ is given by $x_i^0(q_i, q_j) = \frac{1}{\theta} \frac{\partial V_i^0(q_i, q_j)}{\partial q_i}$, implying $x_i^0(q_i^*, q_j^*) = 0$.

¹³We are extremely grateful to Maura Doyle and Chris Snyder for the willingness to share their data with us.

¹⁴These targets are being described by various synonyms: “assembly targets,” “assembly schedules,” “production plans,” “production forecasts,” etc.

¹⁵Some of the observations in the data include post-production revisions. We discard these observations. We only focus on targets posted *before* production. Five production months have no pre-production targets, and are therefore omitted from the analysis.

¹⁶The data also include production targets for American Motors (AMC) until its exit from the market in 1987. We do not use these data for the reported results. AMC has a small market share (2.3% on average) and it exhibits a similar pattern to that of the Big Three, with the exception of its last three years of operation, during which AMC’s market share, production, and production targets rapidly declined. The qualitative results of the paper remain unchanged if we use pre-1984 AMC data.

month, ranging from some cases with a single production target to others with up to 12 associated targets.¹⁷

Figure 11 presents the total number of published targets made at each 10 day interval prior to actual production.¹⁸ It shows that production targets are published approximately 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.

3.2 Empirical Analysis

Let us first introduce some notation. Denote by Q_{it} the actual quantity produced by manufacturer i during month t . Denote by A_{it}^d the production target made by manufacturer i for production month t , with $-d$ representing the number of days between the date of the production target and the target date. Namely, if a production target A_{it}^d is made at date t' then $d = t' - t$. The focus of the analysis is on the way in which A_{it}^d evolves with d .

In order to make targets comparable over time and across manufacturers, we normalize all targets by eventual production. Namely, a (normalized) production target is defined as

$$a_{it}^d \equiv \frac{A_{it}^d - Q_{it}}{Q_{it}} \quad (22)$$

Thus, a_{it}^d is the percentage deviation of the target from the eventual production; it is positive (negative) when a production target is higher (lower) than eventual production.^{19,20} Our key theoretical prediction concerns the change of a_{it}^d with respect to d . We expect a_{it}^d to gradually increase early on, when d is high (in absolute value), and decrease later, as it gets closer towards the production date.

Our analysis is based on pooling observations from multiple production months. The underlying assumption is that, up to the normalization, the same game is played repeatedly over time. This enables us to treat different production targets in different games as if they are made in the

¹⁷The frequency of posted production targets significantly increased in the 1970s. The average number of production targets per production month was 2.13 during 1965-1975, compared to 5.94 and 5.32 during 1976-1985 and 1986-1995, respectively.

¹⁸Since production decisions reflect total production for the month, we follow Doyle and Snyder (1999) and use the last day of the production month as the relevant “date” of production.

¹⁹This transformation of the data is similar to the *PPE* measure used in Doyle and Snyder (1999). Our measure uses a slightly different normalization to relate it more closely to the theoretical predictions. All the qualitative results are robust to alternative normalization choices, including the *PPE* measure of Doyle and Snyder.

²⁰There are six instances of extreme outliers. Five of them are due to unexpected low Q_{it} 's, which generate high a_{it}^d 's, more than three times eventual production ($a_{it}^d > 2$). The sixth instance is of zero announcements by Chrysler. While these cases do not affect the general pattern in any important way, we drop them to reduce noise. We take a conservative approach and also drop all other production targets (at different times and by other manufacturers) associated with the same production month. This leaves us with 361 production months and 1,598 targets by each manufacturer for the empirical analysis.

same context. We then use quartic (biweight) kernel regressions of a_{it}^d on d to non-parametrically describe the evolution of production targets over time. In all figures, we use a bandwidth of 30 days. We repeat this exercise for each manufacturer i separately, for the Big Three average, $a_{Big3,t}^d = \frac{1}{3} \left(a_{GM,t}^d + a_{Ford,t}^d + a_{Chrysler,t}^d \right)$, and for different subsamples of the data. In this section we describe our findings; we defer to the next section the discussion of the link between the empirical exercise and the theoretical assumptions.

The key evidence is presented in Figure 12, which pools all production months in the data. The qualitative picture is of a non-monotonic pattern. On average, production targets start about 5 percent above eventual production levels and gradually increase. They peak 2-3 months before production at about 10 percent, and then gradually decline towards actual production levels. This pattern is not uniform across manufacturers. While Ford and Chrysler, the two smaller firms, follow a similar non-monotonic pattern of production targets, GM exhibits a different behavior. GM's average initial production target is about 15 percent above its eventual production level, and it gradually declines as the deadline gets closer. This is 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 higher than those of Ford and Chrysler. In the repeated game model, for example, such variation could arise if the φ parameter for GM is sufficiently close to zero.

The dashed lines in Figure 12 report 95 percent confidence intervals. These 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. These show that the observed decline in planned production towards the production deadline is quite precisely estimated. This is a pattern that is extremely consistent across manufacturers and for different subsamples. Figure 12 also shows that the confidence intervals significantly 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 information shocks, which are likely to be more pronounced when the production deadline is further away in the future. The second reason is apparent from Figure 11: the number of observed production targets 3-6 months before production is significantly smaller than the number of observations 0-3 months before production.

Our theoretical prediction concerns a non-monotonic pattern of production targets with respect to the same production month. A potential concern may be that while the average pattern shown is qualitatively consistent with the theoretical prediction, it may be generated by aggregation over periods, but is not present in individual patterns.²¹ To address this concern, we repeat the same exercise for different subsamples of the data. Figure 13 divides the sample into three decades. Figure 14 performs the analysis for each calendar month separately to account for potential

²¹For example, one could imagine an extreme case in which half of the patterns are monotonically increasing and concave and half are monotonically decreasing and concave. In such a case, the average pattern may show non-monotonicity even though none of the individual patterns is such.

seasonal variation (due, for example, to model-year product-life-cycle effects; see Copeland, Dunn, and Hall, 2005). Figure 15 repeats the exercise separately for months in which production growth is positive and months in which production growth is negative. All these exercises show similar qualitative patterns. First, the declining production targets during the last 2-3 months before production are present in every single regression. Second, in the majority of the cases one can observe the increase in production targets early on. This second observation does not hold in every regression. This may be expected because, as already mentioned, the data are more noisy for early targets.

As already discussed, the non-monotonic pattern predicts a positive slope of a_{it}^d with respect to d early on, and a negative slope towards production. In order to test directly for the change in slopes, we perform two final exercises. First, we divide production targets into three categories – Early, Middle, and Late – according to how far in advance these targets are made. Table 1 reports the frequencies in which (i) early targets are lower than intermediate targets, (ii) intermediate targets are higher than late targets, and (iii) late targets are higher than eventual production. We report this for each manufacturer, as well as for the Big Three average. All these 12 frequencies except one are greater than 50%. None of them is significantly lower than 50% and the majority of them are significantly higher. This is all consistent with the theory, and gives support to the non-monotonic pattern.

Second, we define the percentage change, per day, in production targets by

$$s_{it}^d \equiv \frac{A_{it}^d - A_{it}^{d'}}{(d - d')A_{it}^{d'}} \quad (23)$$

where $A_{it}^{d'}$ and A_{it}^d are two consecutive production targets associated with the same production month. We then run similar kernel regressions of s_{it}^d with respect to d . Figure 16 reports these regressions. One can observe that in all cases the slope of production targets is positive between 130 days and 80 days before production, and that the confidence interval for the slope estimates lies entirely or almost entirely, depending on the manufacturer, in the positive region. Later on, the slope is significantly negative in all regressions, establishing the non-monotonic pattern.

3.3 Discussion

The theoretical model presented abstracts from certain important aspects of the empirical application, such as inventories, product differentiation, and multi-product manufacturers. While any quantitative analysis ought to account for these effects explicitly, we argue that the qualitative predictions should still hold. The key for the theoretical results is that control variables are strategic substitutes. Thus, as long as production decisions, rather than sales, operate as strategic substitutes, the existence of inventories should not have a qualitative effect on the theory, and therefore on the interpretation of the empirical exercise. Moreover, as long as inventories (and, to a lesser extent, quantity produced abroad) are roughly stable over time, it seems difficult for

inventory fluctuations *per se* to generate the pattern of production targets we observe.²² Similarly, product differentiation and multi-product firms are also unlikely to change the maintained assumption that control variables are strategic substitutes. These assumptions are also consistent with earlier works, which use a Cournot framework to model competition among the Big Three. Berndt et al. (1990) cannot reject the Cournot model in this context, and Doyle and Snyder (1999) use it to test for information-sharing.

In the previous section we show that the pattern of production targets in the U.S. auto industry is consistent with the theoretical framework. To complete the analysis, it is important to discuss two key aspects. First, we identify sources of adjustment costs in the auto industry. Second, we question the manufacturers' incentives to reveal their production targets truthfully. We discuss each aspect in turn.

First, the model assumes that production targets are associated with some real actions, which cannot be costlessly reversed. Auto manufacturers are continuously taking actions that affect their future production capabilities. They contract parts from suppliers, hire temporary labor, cancel vacations, etc. It seems natural to assume that such production-related decisions are costly to change. A late order of parts may be more expensive, revising previously signed contracts may involve penalties, firing workers results in compensation payments, and changing promises may have reputational costs. Moreover, auto manufacturers deal with many third parties, both on the supply and the retail level. If these parties also organize their plans according to the manufacturers' publicly posted targets, a change in these targets may cause them some adjustment costs which may later feed back to the manufacturers' profits.

The contracting channel is one of the main sources of adjustment costs. Given the magnitude and timing of the processes involved in the industry, forward contracts are widespread. In principle, every change in plans would involve renegotiating these contracts. In reality this does not happen so often, as contracts often stipulate clauses that deal with these instances. Typical part contracts in the industry explicitly specify minimum and maximum monthly orders, assigning financial penalties to deviations from this contracted range. Even if these contracts are never renegotiated, implicit adjustment costs arise when contracting with different parties does not simultaneously take place. Since parts are complements in production (consider, for example, an O-ring production function), once contracts are signed sequentially each new contract represents a stronger commitment to a certain production level. Therefore, signing new contracts, which are not fully consistent with earlier signed contracts, carries an implicit adjustment cost, as it would have been cheaper if previous contracts had been set differently.

Second, in the analysis we assume that the manufacturers' reported production targets to *Ward's* truthfully represent their real production plans. This is an important assumption. If these announcements were not anchored to any real decision, they would constitute pure cheap talk. Our view, which is consistent with conversations with manufacturers and *Ward's* publishers,

²²See 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.

is that real actions taken by these firms are difficult to hide from others. Hiring processes for extra shifts, orders of big amounts of windshields, or the construction of a new plant can be easily monitored, not only by competitors, but also by *Ward's*, by the press, and by other external analysts. The main task of the trade journal is to report these actions to third parties (suppliers, dealers, etc.), which cannot perform the monitoring so easily.²³ As we argue above, third parties' plans crucially depend on this information. Moreover, firms are aware that both *Ward's* and other external parties monitor the information they provide to *Ward's*. Finally, one should note that strategic considerations may also work towards providing incentives for truthful reporting. If commitment is achieved by credible higher production targets, then credibility can only be achieved by a reputation for truthful reporting. Given all the above, we consider the production targets published at *Ward's* a good proxy for real decisions being taken by firms. They may represent monthly snapshots of the real underlying continuous decision processes taking place, like the one described by our theoretical model.

4 Concluding remarks

This paper studies the dynamics of pre-production preparation as a commitment device in a quantity setting framework with adjustment costs. We show that firms have a strategic incentive to exaggerate their production targets in an attempt to achieve a Stackelberg leadership position. More precisely, firms start by first steadily increasing their intended production levels and only as production gets closer, do production targets gradually decline. As a result, final production levels are higher than in a static framework.

We test the main predictions of the theory using data on production targets in the U.S. auto industry. The observed production targets exhibit a non-monotonic pattern similar to the one predicted by the theory. While one can write down a variety of theoretical models that could explain why production targets are higher than eventual production,²⁴ it is harder to find alternative explanations for the non-monotonic pattern. This encourages us to view these findings as empirical support for the existence of adjustment costs and for the relevance of the strategic role of pre-production preparations in determining final production decisions.

This study has intentionally abstracted from informational issues as a way to focus on the strategic aspects. Our view is that in reality both components are important and should be accounted for. Given that the model can be easily extended to accommodate (common) uncertainty, as well as multiple players, asymmetries, and a repeated interaction, one could seriously pursue a

²³One could argue that the only reason to publish such information would be its commercial value to third parties. Potential readers are encouraged to subscribe to *Ward's* with the following quote: "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>)

²⁴For example, in the presence of uncertainty, if adjusting quantity downwards is cheaper than adjusting it upwards, then over-targetting would have an option value.

more structural estimation approach. This would be interesting for policy purposes, as one could, for example, quantify the intensity of competition (estimating how far the equilibrium is from the Cournot levels) or perform welfare analysis.²⁵ We leave this exercise for future work.

On a methodological level, we think that our exercise illustrates the empirical potential of non-stationary predictions. As they exhibit rich interesting dynamics, they provide sharp qualitative predictions which have the potential to be empirically verified or falsified. All they require is exogenous variation in time, which is typically satisfied, but do not require further exogeneity assumptions.

References

- [1] Benkard, C. Lanier (2004), “A Dynamic Analysis of the Market for Wide-Bodied Commercial Aircraft,” *Review of Economic Studies*, 71(3), 581-611.
- [2] 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.
- [3] Bresnahan, Timothy F., and Valerie A. Ramey (1994), “Output Fluctuations at the Plant Level,” *Quarterly Journal of Economics*, 109, 593-624.
- [4] Caruana, Guillermo, and Liran Einav (2005), “A Theory of Endogenous Commitment,” mimeo, Stanford University.
- [5] Copeland, Adam, Wendy Dunn, and George Hall (2005), “Prices, Production, and Inventories over the Automotive Model Year,” mimeo, Yale University.
- [6] 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.
- [7] Doyle, Maura P., and Christopher M. Snyder (1999), “Information Sharing and Competition in the Motor Vehicle Industry,” *Journal of Political Economy*, 107(6), 1326-1364.
- [8] Fershtman, Chaim, and Morton I. Kamien (1987), “Dynamic Duopolistic Competition with Sticky Prices,” *Econometrica*, 55(5), 1151-1164.
- [9] Fudenberg, Drew, and Jean Tirole (1984), “The Fat-Cat Effect, the Puppy-Dog Ploy, and the Lean and Hungry Look,” *American Economic Review*, 74, 361-366.
- [10] Hanig, Marco (1986), “Differential Gaming Models of Oligopoly,” Ph.D. Dissertation, Massachusetts Institute of Technology.
- [11] Judd, Kenneth L. (1996), “Cournot versus Bertrand: A Dynamic Resolution,” mimeo, Hoover Institution.

²⁵The welfare implications of the model are ambiguous. Compared with static models, welfare increases as a result of higher production, but decreases as a result of “lost” adjustment costs during the planning phase.

- [12] Jun, Byoung, and Xavier Vives (2004), “Strategic Incentives in Dynamic Duopoly,” *Journal of Economic Theory*, 116(2), 249-281.
- [13] 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.
- [14] Kydland, Finn (1975), “Noncooperative and Dominant Player Solutions in Discrete Dynamic Games,” *International Economic Review*, 16(2), 329-341.
- [15] Lapham, Beverly, and Roger Ware (1994), “Markov Puppy Dogs and Related Animals,” *International Journal of Industrial Organization*, 12(4), 569-593.
- [16] Lukes, D. L. (1971) “Equilibrium Feedback Control in Linear Games with Quadratic Costs,” *SIAM Journal on Control*, 9, 234-252.
- [17] Maskin, Eric, and Jean Tirole (1987), “A Theory of Dynamic Duopoly III: Cournot Competition,” *European Economic Review*, 31, 947-968.
- [18] 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.
- [19] Papavassilopoulos, G. P., and G. J. Olsder (1984), “On the Linear-Quadratic, Closed-Loop, No-Memory Nash Game,” *Journal of Optimization Theory and Applications*, 42, 551-560.
- [20] Reynolds, Stanley S., (1987), “Capacity Investment, Preemption and Commitment in an Infinite Horizon Model,” *International Economic Review*, 28(1), 69-88.
- [21] Reynolds, Stanley S., (1991), “Dynamic Oligopoly with Capacity Adjustment Costs,” *Journal of Economic Dynamics and Control*, 15(3), 491-514.
- [22] Romano, Richard, and Huseyin Yildirim (2005), “On the Endogeneity of Cournot-Nash and Stackelberg Equilibria: Games of Accumulation,” *Journal of Economic Theory*, 120(1), 73-107.
- [23] Ryan, Stephen P. (2004), “The Costs of Environmental Regulation in a Concentrated Industry,” mimeo, Duke University.
- [24] Saloner, Garth (1987), “Cournot Duopoly with Two Production Periods,” *Journal of Economic Theory*, 42, 183-187.

Appendix

The appendix derives the equations that describe the solutions to three extension of the benchmark model, as discussed in Section 2.4.

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

$$\max_{x_i^t} \left(-\frac{\theta}{2} (x_i^t)^2 + \frac{\partial V_i^t}{\partial q_i} x_i^t + \sum_{j \neq i} \frac{\partial V_i^t}{\partial q_j} x_j^t + \frac{\partial V_i^t}{\partial t} \right) = 0 \quad (24)$$

The first order condition for x_i^t implies

$$x_i^t = \frac{1}{\theta} \frac{\partial V_i^t}{\partial q_i} \quad (25)$$

We can now substitute this back into equation (24), as well as the symmetric solution for all other x_j^t 's, rearrange, and obtain the following differential equation

$$0 = \frac{1}{2\theta} \left(\frac{\partial V_i^t}{\partial q_i} \right)^2 + \frac{1}{\theta} \sum_{j \neq i} \left(\frac{\partial V_j^t}{\partial q_j} \right) \left(\frac{\partial V_i^t}{\partial q_j} \right) + \frac{\partial V_i^t}{\partial t} \quad (26)$$

We guess that the value function will be symmetric in the opponents' state variables, so that the quadratic value function can be written as

$$\begin{aligned} V_i^t(q_i, q_j) &= A_t + B_t q_i + \sum_{j \neq i} C_t q_j + D_t q_i^2 + \sum_{j \neq i} E_t q_j^2 + \sum_{j \neq i} F_t q_i q_j + \sum_{j \neq i} \sum_{k \neq i, j} G_t q_j q_k = (27) \\ &= A_t + B_t q_i + C_t Q_{-i} + D_t q_i^2 + E_t R_{-i} + F_t q_i Q_{-i} + G_t S_{-i} \end{aligned}$$

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

$$x_i^t(q_i, q_j) = \frac{1}{\theta} (B_t + 2D_t q_i + F_t Q_{-i}) \quad (28)$$

Thus, we can rewrite equation (26) as

$$\begin{aligned} 0 &= \frac{1}{2\theta} (B_t + 2D_t q_i + F_t Q_{-i})^2 + \frac{1}{\theta} \sum_{j \neq i} (C_t + 2E_t q_j + F_t q_i + 2G_t(Q_{-j} - q_i)) (B_t + 2D_t q_j + F_t Q_{-j}) + \\ &+ (A'_t + B'_t q_i + C'_t Q_{-i} + D'_t q_i^2 + E'_t R_{-i} + F'_t q_i Q_{-i} + G'_t S_{-i}) \end{aligned} \quad (29)$$

After collecting terms (and reversing signs for A' , B' , etc. as in the benchmark model) and equating coefficients, we obtain the following law of motion:

$$\begin{pmatrix} A' \\ B' \\ C' \\ D' \\ E' \\ F' \\ G' \end{pmatrix} = \frac{1}{\theta} \begin{pmatrix} \frac{1}{2} B^2 + (N-1)BC \\ 2BD + BF(N-1) + CF(N-1) \\ BF + 2BE + 2CD + CF(N-2) + 2BG(N-2) \\ 2D^2 + F^2(N-1) \\ \frac{1}{2} F^2 + 4DE + 2FG(N-2) \\ 4DF + 2EF + F^2(N-2) + 2FG(N-2) \\ \frac{1}{2} F^2 + 2EF + 4GD + 4FG(N-3) \end{pmatrix} \quad (30)$$

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

$$\begin{pmatrix} A_T \\ B_T \\ C_T \\ D_T \\ E_T \\ F_T \\ G_T \end{pmatrix} = \begin{pmatrix} 0 \\ a - c \\ 0 \\ -b \\ 0 \\ -b \\ 0 \end{pmatrix} \quad (31)$$

Asymmetric Players 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 $c_i(x_i, t) = \frac{\theta^i}{2}x_i^2$, her (constant) marginal cost is c^i , and A_t^i to F_t^i denote i 's value function coefficients.

One can start by following the same steps as in Section 2.1. The first difference appears in equation (5), which now reads

$$\frac{1}{2\theta^i} \left(\frac{\partial V_i^t}{\partial q_i} \right)^2 + \frac{1}{\theta^j} \left(\frac{\partial V_i^t}{\partial q_j} \right) \left(\frac{\partial V_j^t}{\partial q_j} \right) + \frac{\partial V_i^t}{\partial t} = 0 \quad (32)$$

The value function for each player is

$$V_i^t(q_i, q_j) = A_t^i + B_t^i q_i + C_t^i q_j + D_t^i q_i^2 + E_t^i q_j^2 + F_t^i q_i q_j \quad (33)$$

Substituting it into equation (32) gives

$$0 = \frac{1}{2\theta^i} (B_t^i + 2D_t^i q_i + F_t^i q_j)^2 + \frac{1}{\theta^j} (C_t^i + 2E_t^i q_j + F_t^i q_i) (B_t^j + 2D_t^j q_j + F_t^j q_i) + (A_t^{i'} + B_t^{i'} q_i + C_t^{i'} q_j + D_t^{i'} q_i^2 + E_t^{i'} q_j^2 + F_t^{i'} q_i q_j) \quad (34)$$

By collecting terms one obtains the law of motion for the coefficients in player i 's value function (symmetrically for player j):

$$\begin{pmatrix} A^{i'} \\ B^{i'} \\ C^{i'} \\ D^{i'} \\ E^{i'} \\ F^{i'} \end{pmatrix} = \begin{pmatrix} \frac{1}{2\theta^i} B^{i2} + \frac{1}{\theta^j} B^j C^i \\ \frac{2}{\theta^i} B^i D^i + \frac{1}{\theta^j} B^j F^i + \frac{1}{\theta^j} C^i F^j \\ \frac{1}{\theta^i} B^i F^i + \frac{2}{\theta^j} B^j E^i + \frac{2}{\theta^j} C^i D^j \\ \frac{2}{\theta^i} D^{i2} + \frac{1}{\theta^j} F^i F^j \\ \frac{1}{2\theta^i} F^{i2} + \frac{4}{\theta^j} D^j E^i \\ \frac{2}{\theta^i} D^i F^i + \frac{2}{\theta^j} D^j F^i + \frac{2}{\theta^j} E^i F^j \end{pmatrix} \quad (35)$$

with the boundary condition given by

$$\begin{pmatrix} A_T^i \\ B_T^i \\ C_T^i \\ D_T^i \\ E_T^i \\ F_T^i \end{pmatrix} = \begin{pmatrix} 0 \\ a - c^i \\ 0 \\ -b \\ 0 \\ -b \end{pmatrix} \quad (36)$$

Uncertainty We follow the same steps as in Section 2.1 with few modifications. Now there are two value functions, depending on the state of the economy. Let these two value functions be V_L and V_H . Thus, the Bellman equation if one is in state L is

$$\max_{x_i^t} \left(-\frac{\theta}{2} (x_i^t)^2 + \frac{\partial V_{L,i}^t}{\partial q_i} x_i^t + \frac{\partial V_{L,i}^t}{\partial q_j} x_j^t + \frac{\partial V_{L,i}^t}{\partial t} + \lambda (V_{H,i}^t(q_i, q_j) - V_{L,i}^t(q_i, q_j)) \right) = 0 \quad (37)$$

and symmetrically for V_H . The optimal adjustment rate is given by

$$x_i = \frac{1}{\theta} \frac{\partial V_{L,i}^t}{\partial q_i} \quad (38)$$

and symmetrically for H . Now one can obtain the corresponding differential equations as in equations (5) and (8), resulting in a system of twelve ODE's. The law of motion for the coefficients associated with the L state is

$$\begin{pmatrix} A'_L \\ B'_L \\ C'_L \\ D'_L \\ E'_L \\ F'_L \end{pmatrix} = \lambda \begin{pmatrix} A_H - A_L \\ B_H - B_L \\ C_H - C_L \\ D_H - D_L \\ E_H - E_L \\ F_H - F_L \end{pmatrix} + \frac{1}{\theta} \begin{pmatrix} \frac{1}{2} B_L^2 + B_L C_L \\ 2B_L D_L + B_L F_L + C_L F_L \\ B_L F_L + 2B_L E_L + 2C_L D_L \\ 2D_L^2 + F_L^2 \\ \frac{1}{2} F_L^2 + 4D_L E_L \\ 4D_L F_L + 2E_L F_L \end{pmatrix} \quad (39)$$

Additional six symmetric equations are associated with the H state. This structure is identical to the benchmark model except for the fact that, in each equation, with probability λ we switch to the other value function. Finally, the boundary conditions are given by the different profit functions at each state.

Figure 1: Parameters of the value function

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

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

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 (and we verify with longer horizons) that as the horizon becomes longer all parameters, except for the constant A_0 , approach zero. A_0 converges to approximately 0.0925. Thus, for games with long horizons the values converge to 0.0925, which are almost 20% lower than the static Cournot profits of $\frac{1}{9}$.




Figure 1: Parameters of the value function

Figure 2: Equilibrium path and off-equilibrium strategies

This figure plots the equilibrium path in the benchmark model, when parameters are set to $a = b = 1$, $c = 0$, and $\theta = 1$, and initial production plans are at the Cournot level of $\frac{1}{3}$. It shows that equilibrium path is non-monotonic: it peaks at about 0.4 and then declines towards 0.37, which is the equilibrium production level. This level is higher than the Cournot level of $\frac{1}{3}$. The dashed lines illustrates off-equilibrium path strategies. It simulates a shock, occurring at date $t = -4$, which exogenously and unexpectedly drops one of the player's production target to $\frac{1}{3}$. The subsequent dashed lines present the equilibrium path in the continuation game, after the shock. It shows that the leadership position persists, illustrating why players cannot credibly coordinate on sticking to the Cournot levels.




Figure 2: Equilibrium path and off-equilibrium strategies

Figure 3: Equilibrium path with different (symmetric) initial actions

This figure plots the equilibrium path in the benchmark 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 cases are of full symmetry, in parameters and in initial actions, so the equilibrium path is identical for both players in each case. Clearly, equilibrium paths of the different cases do not cross each other. Note, however, that final production levels are much closer (around 0.37 in all cases, but not the same, keeping the same ordering as that of initial levels) 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 3: Equilibrium path with different (symmetric) initial actions

Figure 4: Equilibrium path with different horizons

This figure plots the equilibrium path in the benchmark 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 an 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 4: Equilibrium path for different horizons

Figure 5: Equilibrium path with different adjustment cost parameters

This figure plots the equilibrium path in the benchmark 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 (0.1, 1, and 10). One can clearly observe that as adjustment costs are lower, production targets peak higher, as it is both cheaper to achieve these levels, and lower targets do not provide enough commitment. The picture also suggests that final production levels decrease with θ . This may be misleading. Since production level is $\frac{1}{3}$ for $\theta = 0$, we suspect that they may increase with θ close to zero. We could not, however, verify this conjecture numerically.




Figure 5: Equilibrium path with different adjustment cost parameters

Figure 6: 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 than the Cournot level.




Figure 6: Illustration of how the best response functions change as a result of adjustment costs

Figure 7: 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 gains 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). Total production is higher than the Cournot level in all cases.




Figure 7: Asymmetric players with different marginal costs

Figure 8: Asymmetric players with different adjustment costs

This figure plots the equilibrium path in a two-player model with asymmetric players. Parameters are set to $a = b = 1$, and $c = 0$. One player (player 1), however, has higher adjustment cost parameter than her opponent. As discussed in the text, it is somewhat interesting that once initial conditions are sufficiently low (as in the figure), it is the more flexible player who is able to obtain higher market shares. It is not clear if this commitment advantage persists for any value of $\theta > 0$. For $\theta = 0$ the flexible player always best responds at time T and therefore cannot enjoy a commitment advantage.




Figure 8: Asymmetric players with different adjustment costs

Figure 9: Equilibrium path in the repeated game

This figure plots the equilibrium path of the repeated game described in Section 2.5. Parameters are set to $a = b = 1$, $c = 0$, $\theta = 1$, $T = 10$, $\varphi = 0.1$, and $\beta = 0.9$. Production plans follow the same pattern before every production period, and production levels (approximately 0.37) are constant. Initial plans in this case are not exogenous, but are set endogenously as part of the equilibrium.




Figure 9: Equilibrium path in the repeated game

Figure 10: Comparative statics in the repeated game

This figure plots the equilibrium path of the repeated game for different levels of adjustment costs. The rest of the parameters are set to $a = b = 1$, $c = 0$, $T = 10$, and $\beta = 0.9$. The figure presents a snapshot of one stage of the game. Since we solve for a steady state, this snapshot repeats itself forever, as in Figure 9. As one can observe, φ mainly affects the initial plans, while θ affects the dynamic pattern of plans. As discussed in the text, the case of $\varphi = 0$ is a special case in which the equilibrium of the repeated game is identical to the benchmark model with free initial decisions. Note, however, that even small values for φ are sufficient to generate non-monotonic pattern. This is because continuation values do not change much with initial plans.




Figure 10: Comparative statics in the repeated game

Figure 11: Frequency and timing of production target observations

This figure provides information about the timing of the observations available. Each observation includes separate production target by each of the Big Three associated with a particular production month. 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 (more than four months in advance) observations about production targets are not as regular.




Figure 11: Frequency and timing of production target observations

Figure 12: Production targets over time, pooling all data

This figure presents quartic (biweight) kernel regressions of production targets, measured by a_{it}^d (see equation (22)), as a function of the number of days before production, d . It does so for each of the major three manufacturers (GM, Ford, and Chrysler), as well as for the (unweighted) average (Big Three). Each series is based on 1,598 observations. All estimates use bandwidth of 30 days. The dashed lines present 95 percent confidence intervals. 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.




Figure 12: Production targets over time, pooling all data

Figure 13: Production targets over time, by period

This figure presents quartic (biweight) kernel regressions of production targets, measured by a_{it}^d (see equation (22)), as a function of the number of days before production, d . It does so for each of the major three manufacturers (GM, Ford, and Chrysler), as well as for the (unweighted) average (Big Three). Each figure reports the kernel regression estimates, estimated separately for each decade of the data: 1965-1975 (thick solid line), 1976-1985 (dashed line), and 1986-1995 (thin solid line). The estimates for the first decade (1965-1975) only start about 120 days before production, as during this period they were no early production target observations.




Figure 13: Production targets over time, by period

Figure 14: Production targets over time, by calendar month

This figure presents quartic (biweight) kernel regressions of production targets, measured by a_{it}^d (see equation (22)), as a function of the number of days before production, d . It does so for each calendar month separately, to account for potential seasonality of model-year product-life-cycle effects. Each figure presents the kernel regression estimates for each of the Big Three, as well as for their average.




Figure 14: Production targets over time, by calendar month

Figure 15: Production targets over time, for positive and negative production growth

This figure presents quartic (biweight) kernel regressions of production targets, measured by a_{it}^d (see equation (22)), as a function of the number of days before production, d . It does so for each of the major three manufacturers (GM, Ford, and Chrysler), as well as for the (unweighted) average (Big Three). Each figure reports the kernel regression estimates, estimated separately for production months with positive production growth (thick line) and negative production growth (thin line). Much of the variation in production growth is due to seasonality. As discussed in the text, one can notice that much of the non-monotonic pattern in the data is driven by months with positive production growth.




Figure 15: Production targets over time, for positive and negative production growth

Figure 16: Revisions in production targets

This figure presents quartic (biweight) kernel regressions of the *revisions* in production targets, measured by s_{it}^d (see equation (23)), as a function of the number of days before production, d . The units of s_{it}^d are in basis point change, per day. The figures present the pattern for each of the major three manufacturers (GM, Ford, and Chrysler), as well as for the (unweighted) average (“Big Three”). Each series is based on 1,239 observations (taking first differences, we lose the earliest observed announcement for each production month). All estimates use bandwidth of 30 days. The dashed lines present 95 percent confidence intervals. 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.

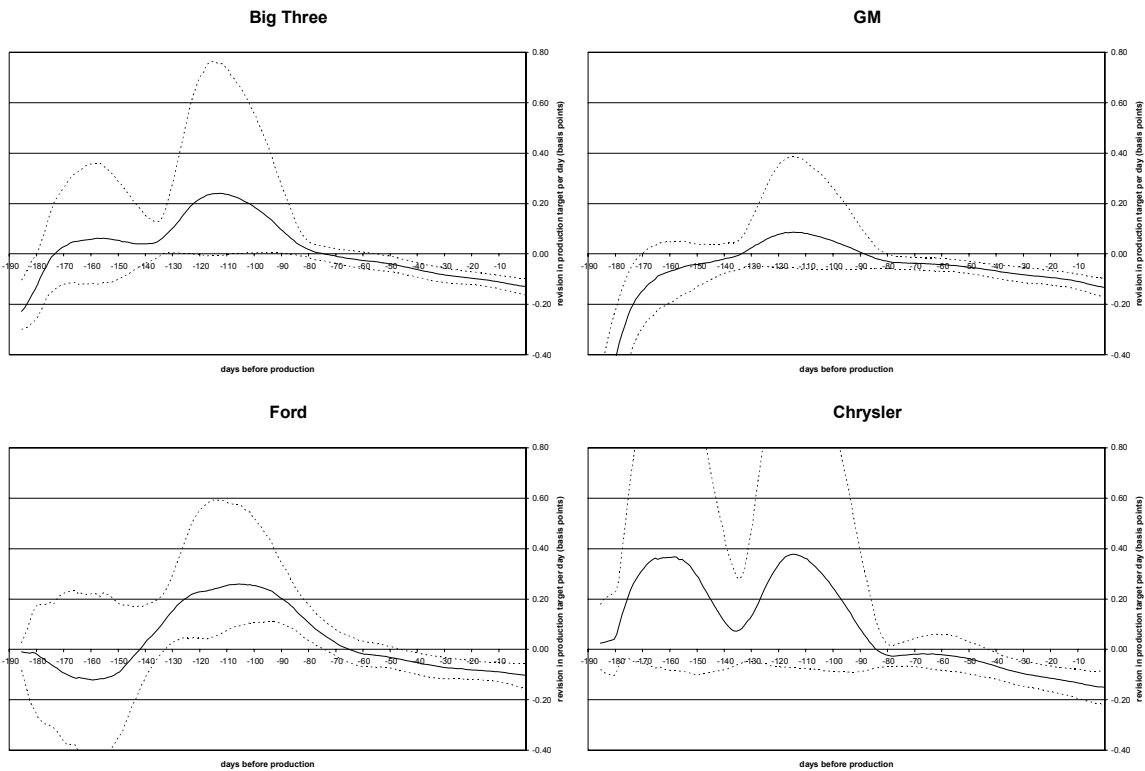


Figure 16: Revisions in production targets

Table 1: Frequency estimates of revision signs

This table provides frequency estimates of the direction of production target revisions. The inequalities are constructed in such a way that estimates of 0.5 imply random revisions and estimates greater than 0.5 are consistent with the theoretical predictions. As one can observe, all numbers but one are greater than 0.5, none of them is significantly less than 0.5, and the majority of them are significantly greater than 0.5.

	$Pr(A_{it}^{Middle} > A_{it}^{Early})^a$	$Pr(A_{it}^{Late} < A_{it}^{Middle})^a$	$Pr(Q_{it} < A_{it}^{Late})^a$
Big 3 Average	0.614**	0.628**	0.740**
GM	0.474	0.584**	0.763**
Ford	0.667**	0.509	0.676**
Chrysler	0.511	0.528	0.543
Obs.	135	286	359

** Significantly different from 0.5 at 95% confidence level.

^a For each i and t we construct A_{it}^{Early} as the average of A_{it}^d such that $d < -110$. Respectively, for A_{it}^{Middle} we use $d \in [-110, -50]$ and for A_{it}^{Late} we use $d > -50$. Changing the cutoff levels for these variables has no effect on the results.