

Production targets

Guillermo Caruana*

and

Liran Einav**

We analyze a dynamic model of quantity competition, where firms continuously adjust their quantity targets, but incur convex adjustment costs when they do so. Quantity targets serve as a partial commitment device and, in equilibrium, follow a hump-shaped pattern. The final equilibrium is more competitive than in the static analog. We then use data on monthly production targets of the Big Three U.S. auto manufacturers and show a similar empirical hump-shaped dynamic pattern. Taken together, this suggests that strategic considerations may play a role in setting auto production schedules, and that static models may misestimate the industry's competitiveness.

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, and so forth. 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

*CEMFI; caruana@cemfi.es.

**Stanford University and NBER; leinav@stanford.edu.

We thank dozens of auto industry experts for insightful conversations. We thank two anonymous referees and Phil Haile (the editor), as well as Susan Athey, Steve Berry, Tim Bresnahan, Uli Doraszelski, Shane Greenstein, Ken Judd, Aviv Nevo, Ignacio Palacios-Huerta, Chris Snyder, Ralph Winter, and especially Esteban Rossi-Hansberg for helpful conversations, and seminar participants at IO Fest 2004, SICS 2005, AEA 2006 annual meeting, SED 2006 annual meeting, UAB, Bilbao, UC Berkeley, UCLA, Carlos III, CEMFI, ECARES, UIUC, LSE, Minneapolis Fed, Northwestern, NYU Stern, Olin School of Business, Oviedo, Tel-Aviv, and Yale for comments. We are extremely thankful to Maura Doyle and Chris Snyder for so kindly sharing their data with us. An earlier version of this article had the title “Quantity Competition with Production Commitment: Theory and Evidence from the Auto Industry.” Einav acknowledges financial support from the National Science Foundation.

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 targets, and use this to their advantage, adjusting their own intentions strategically.

The main goal of the article 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 article 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 article makes three separate contributions. First, we present new theoretical predictions for quantity setting games regarding the nonmonotone evolution of production targeting. Because 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 nonmonotonic pattern of production targets in the U.S. auto industry. Because this is one of the largest industries in the United States, 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 misestimate the competitiveness of the industry.

Section 2 contains the theoretical part of the article. We first present a simple two-period example which provides the key intuition. We then present the baseline 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 onward, 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 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 "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, whereas 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 intentions 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 baseline 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). The final and most important extension nests the baseline model as the stage game of an infinitely repeated game. We solve for a Markov Perfect Equilibrium of this game, and show that its equilibrium path exhibits the same nonmonotonic pattern. Moreover, the repeated game provides a natural way to endogenize the initial production targets, which are taken as given in the baseline model. This extension also takes the model one step closer to the reality of the empirical application we study later in the article.

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 the commitment advantage monotonically increases with targeted 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, upward and downward, 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 (2008), 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, 1991; Driskill and McCafferty, 1989; Lapham and Ware, 1994; Jun and Vives, 2004). These papers focus on the stable equilibrium of an infinite-horizon model. They typically find that (when actions are strategic substitutes) the equilibrium is more competitive than its static analog, as players engage in Stackelberg warfare.² Our model shares this feature but, unlike this literature, our main focus is on the planning phase. One advantage in studying the dynamics of the planning phase is its strong nonstationarity; it provides a clear prediction with respect to an observed and exogenous state variable, namely time. Predictions of stationary dynamic models are more difficult to verify in the data, as the static benchmark is typically not available (for example, marginal costs are typically not observed).

Section 3 empirically explores the predictions of the model using data on monthly production targets by the Big Three auto manufacturers in the United States 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 nonmonotonic pattern, which is consistent with the theoretical prediction. Early targets, about six months prior to production, overstate eventual production by about 5%. Then they start to slowly increase, until they peak at 10% about two to three months before production. At this point, they start to gradually decline toward the eventual production levels. This result is robust to alternative measurements and across different subsamples. Establishing the relationship between the empirical pattern and the theoretical predictions suggests that adjustment costs and strategic considerations may play an important role in the planning phase of production and that static models may therefore underestimate the competitiveness of the industry.³

A significant part of Section 3 is devoted to a discussion of the relationship between the data analyzed and the theory previously developed. First, we argue, on both theoretical and empirical grounds, why production decisions in the industry act as strategic substitutes. Second, we explain why it is reasonable to view the reported production targets as a good proxy for actual production scheduling decisions made by the Big Three. Finally, we describe the multiple sources of adjustment costs that exist in the empirical setting.

At some general level, this work can be classified within the recent empirical studies of dynamic oligopolies (e.g., Benkard, 2004; Ryan, 2006). 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 qualitative predictions. Therefore, the primary objective here is to analyze whether the data are consistent with these predictions.

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

² As Jun and Vives (2004) point out, this can be viewed as a dynamic extension of a “top dog” strategy within the Fudenberg and Tirole (1984) taxonomy of strategic behavior.

³ This is the case if marginal costs are observed. If the equilibrium model is used to back out marginal costs, as in much of the recent literature in Industrial Organization, it is the marginal costs that would be underestimated when static competition models are assumed.

The data we use in this work are also used in Doyle and Snyder (1999). They investigate the role of reported 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 article 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 targets 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. Therefore, our framework is more relevant for the average pattern rather than for the period-by-period pattern, whereas information-sharing motives are more likely to be important and observed *within* production periods.

2. Theory

■ **A two-period example.** We start by illustrating the key qualitative predictions of the model using a simple two-period example. Consider a game where two firms start with exogenously given (symmetric) initial production targets $y_1 = y_2 = y$. At $t = 1$, each firm $i = 1, 2$ can revise its target from y_i to z_i , but pays a quadratic adjustment cost of $\frac{\theta}{2}(z_i - y_i)^2$ to do so. Then, at $t = 2$, each firm has a second (and final) opportunity to revise its production target to q_i , paying adjustment costs of $\frac{\theta}{2}(q_i - z_i)^2$. Given these final 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 this game using backward induction. At $t = 2$, each firm i chooses q_i to solve

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

Best-response functions are

$$q_i = \frac{1 - q_j + \theta z_i}{2 + \theta}, \tag{2}$$

and the second-period equilibrium strategies are

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

One can check that if firms target the Cournot quantities, $z_1 = z_2 = \frac{1}{3}$, then setting $q_1 = q_2 = \frac{1}{3}$ is an equilibrium. In general, as Figure 1 illustrates, the first-order conditions define a best-response function which is a rotation of the static best response at the previously targeted production level. Each firm’s response to a change in its opponent’s quantity is not as strong as it would have been in the absence of adjustment costs. Thus, if the symmetric targets, $z_1 = z_2$, are greater (less) than $\frac{1}{3}$, the firms end up adjusting in the direction of their static best responses, but not fully, thereby ending up in a more (less) competitive equilibrium.

At $t = 1$ firms choose z_1 and z_2 , accounting for the equilibrium strategies at $t = 2$. Thus, each player i chooses z_i to solve

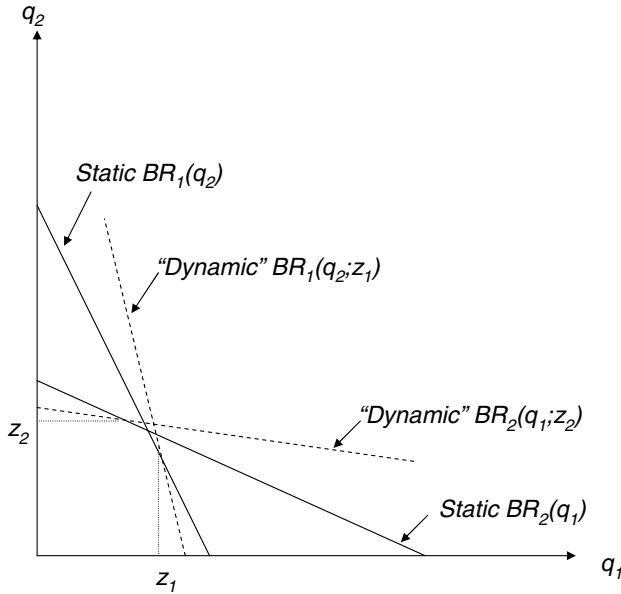
$$\max_{z_i} (1 - q_i(z_i, z_j) - q_j(z_i, z_j))q_i(z_i, z_j) - \frac{\theta}{2}(q_i(z_i, z_j) - z_i)^2 - \frac{\theta}{2}(z_i - y)^2 \tag{4}$$

implying the following first-order condition for each player:

$$\frac{\partial q_i}{\partial z_i} (1 - q_i - q_j) - q_i \left(\frac{\partial q_i}{\partial z_i} + \frac{\partial q_j}{\partial z_i} \right) - \theta (q_i - z_i) \left(\frac{\partial q_i}{\partial z_i} - 1 \right) - \theta (z_i - y) = 0. \tag{5}$$

FIGURE 1

BEST-RESPONSE FUNCTIONS IN THE TWO-PERIOD MODEL



This figure sketches the dynamic effect of adjustment costs in the context of the two-period example. 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.

This yields a (symmetric) solution $z(y, \theta)$ and $q(y, \theta)$.⁴ If $y = \frac{1}{3}$, that is, firms’ initial targets are at the Cournot level, their final productions would be

$$q\left(\frac{1}{3}, \theta\right) = \frac{1}{3} + \frac{\theta}{3\theta^3 + 30\theta^2 + 78\theta + 54}, \tag{6}$$

which are always above $\frac{1}{3}$ for any $\theta > 0$. For example, when $\theta = 1$, equilibrium targets at $t = 1$ are $z \approx 0.357$ and final production levels are $q \approx 0.339$. In summary, production targets in this example increase first and decrease later, and the overall result is more competitive than in a static game (when $\theta = 0$). In the first period, firms have an incentive to exaggerate their production intentions as a way to achieve commitment. In the second period, absent the commitment motives, the incentive is to best respond to the opponent’s high production target, but only partially, given the adjustment costs. As we will see below, these qualitative conclusions hold under more general assumptions.

□ **The baseline model.** We now introduce the baseline model, which is in continuous time. There are two players. At time $t = 0$, they start with exogenously inherited initial production targets of $(q_1(0), q_2(0))$. At all points $t \in [0, T]$ each player i chooses $x'_i \in \mathbb{R}$, which controls the rate at which she changes her production target, that is, $q'_i(t) = x'_i$. Note that x'_i can be either positive or negative. If a player changes her target 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 targets, $q_1(T)$ 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; we assume that inverse demand is linear, given by $p = a - b(q_1 + q_2)$, and marginal costs are constant and

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

given by c . Thus, we have that

$$\begin{aligned} \pi_i(q_i(T), q_j(T)) &= (a - bq_i(T) - bq_j(T))q_i(T) - cq_i(T) \\ &= (a - c)q_i(T) - bq_i^2(T) - bq_i(T)q_j(T). \end{aligned} \tag{7}$$

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

$$c_i(x_i, t) = \frac{\theta}{2}x_i^2. \tag{8}$$

Note that adjustment costs are constant over time,⁵ symmetric across players, and symmetric for 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 $V'_i(q_i, q_j)$ be the value function for player i at time t , with state variables q_i and q_j . If $V'_i(q_i, q_j)$ exists and is continuous and continuously differentiable in its arguments, then it satisfies the following Bellman equation:

$$\max_{x'_i} \left(-\frac{\theta}{2}(x'_i)^2 + \frac{\partial V'_i}{\partial q_i}x'_i + \frac{\partial V'_i}{\partial q_j}x'_j + \frac{\partial V'_i}{\partial t} \right) = 0. \tag{9}$$

The first-order condition for x'_i implies that

$$x'_i = \frac{1}{\theta} \frac{\partial V'_i}{\partial q_i}. \tag{10}$$

We can now substitute this back into equation (9), and obtain the following differential equation:

$$\frac{1}{2\theta} \left(\frac{\partial V'_i}{\partial q_i} \right)^2 + \frac{1}{\theta} \left(\frac{\partial V'_i}{\partial q_j} \right) \left(\frac{\partial V'_j}{\partial q_j} \right) + \frac{\partial V'_i}{\partial t} = 0. \tag{11}$$

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 value function is a quadratic function of the state variables.⁶ Note that due to the inherent nonstationarity of the model, the parameters of this quadratic equation will depend on t in an unspecified way. We can express the value function as

$$V'_i(q_i, q_j) = A_t + B_tq_i + C_tq_j + D_tq_i^2 + E_tq_j^2 + F_tq_iq_j \tag{12}$$

which, using equation (10), implies that

$$x'_i(q_i, q_j) = \frac{1}{\theta}(B_t + 2D_tq_i + F_tq_j). \tag{13}$$

Given that players are symmetric, we use equations (11) and (12) to obtain

$$\begin{aligned} 0 &= \frac{1}{2\theta}(B_t + 2D_tq_i + F_tq_j)^2 + \frac{1}{\theta}(C_t + 2E_tq_j + F_tq_i)(B_t + 2D_tq_j + F_tq_i) \\ &\quad + (A'_t + B'_tq_i + C'_tq_j + D'_tq_i^2 + E'_tq_j^2 + F'_tq_iq_j). \end{aligned} \tag{14}$$

This is a polynomial in q_i and q_j . Because it has to be satisfied for all values of q_i and q_j , all of 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 backward.

⁵ For 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.

⁶ See 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

$$\begin{pmatrix} A' \\ B' \\ C' \\ D' \\ E' \\ F' \end{pmatrix} = \frac{1}{\theta} \begin{pmatrix} \frac{1}{2}B^2 + BC \\ 2BD + BF + CF \\ BF + 2BE + 2CD \\ 2D^2 + F^2 \\ \frac{1}{2}F^2 + 4DE \\ 4DF + 2EF \end{pmatrix} \tag{15}$$

with boundary condition (for $t = T$)

$$\begin{pmatrix} A_T \\ B_T \\ C_T \\ D_T \\ E_T \\ F_T \end{pmatrix} = \begin{pmatrix} 0 \\ a - c \\ 0 \\ -b \\ 0 \\ -b \end{pmatrix} \tag{16}$$

which is provided by the profit function in equation (7).

□ **Equilibrium properties.** The system of ordinary differential equations given by equation (15), 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 (13). The system cannot be solved analytically, but the equilibrium can be approximated through the solution of the discrete-time analog of the game for very small time intervals.

Throughout this section, we illustrate the properties of the equilibrium for a specific choice of parameters. This choice is inconsequential and the properties we report are generic. Specifically, 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 $c_i(x_i, t) = \frac{1}{2}x_i^2$.⁷ 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 $q = \frac{1}{3}$, whereas the Stackelberg leader and follower production levels are $q = \frac{1}{2}$ and $q = \frac{1}{4}$, respectively.


The top panel of Figure 2 presents the symmetric equilibrium path for the game in which both players inherit an initial production target 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 toward it. Given that this is costly, the parties do not adjust all the way to the static Nash equilibrium.⁸ 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 target at $t = T - 4$ and has her target 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 target even further. Meanwhile, player i 's best response is to rebuild her commitment by increasing her target. Nevertheless, the advantageous position acquired by player j never fully diminishes and is kept until the production date.

⁷ One should note that some of these restrictions 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.

⁸ With 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, whereas the marginal benefit is strictly positive for small adjustments but zero for full adjustments.

FIGURE 2

EQUILIBRIUM IN THE BASELINE MODEL



This figure plots the equilibrium path in the baseline model, when $a = b = 1$, $c = 0$, and $\theta = 1$, and initial production targets are symmetric. In the top panel, initial targets are at the Cournot level of $1/3$, illustrating the nonmonotone equilibrium path: targets peak at about 0.4 and then decline toward 0.37, which is the equilibrium production level. The dashed lines illustrate off-equilibrium-path strategies. They simulate a change, occurring at date $t = -4$, which exogenously and unexpectedly drops one of the players' production target to $1/3$. The subsequent dashed lines present the equilibrium path after the change (different for each player). It shows that the leadership position persists, illustrating why players cannot credibly coordinate on sticking to the Cournot levels. The bottom panel presents how the equilibrium path changes with different (symmetric) initial targets, ranging from $1/3$ (the Cournot level) to 0.5 (the Stackelberg level). We note that although final equilibrium production levels are much closer (around 0.37 in all cases) than the initial targets, they are not the same. In particular, final production levels are monotone in the initial targets. The nonmonotone pattern persists as long as initial targets are sufficiently low (less than about 0.44 in this case). When initial actions are higher, the equilibrium path is monotone but concave (due to the commitment effect).

The bottom panel of Figure 2 presents the symmetric equilibrium path for different initial production targets. If these are not too high, one observes the same pattern as in the previous figure. If initial targets 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 targets over time is not constant, however, due to the commitment effect. They first decrease targets slowly, so they remain committed to high quantities, and only later they speed up adjustments in the direction of their static best-response levels.⁹

Figure 3 presents 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 (15) 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. Thus, changes in the adjustment cost parameter are similar to a rescaling of time.¹⁰ The top panel of Figure 3 shows how the length of the horizon affects the equilibrium path. As the horizon gets longer, there is more time to build up commitment. However, because the buildup is more gradual, it is not as costly. Indeed, in the limit (as $T \rightarrow \infty$) the equilibrium profits converge to a constant of 0.0925,¹¹ which is approximately 17% lower than the static Cournot profits of $\frac{1}{9}$.¹² This illustrates how the dynamic interaction leads to a reduction in profits. If they could, the two parties would avoid the “preparation race” and commit to the static Cournot outcome throughout. Similarly, the bottom panel of Figure 3 shows that as the adjustment costs decrease, building commitment becomes cheaper. In both cases, this leads to higher targets and an ultimately faster decline. This is an important observation that we emphasize later: smaller (but positive) adjustment costs lead to a more pronounced hump shape.¹³

□ **Extensions to the baseline model.** Here we present some of the most natural extensions to the baseline 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 baseline 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 (e.g., firms may vary in their marginal costs) or through the adjustment costs (e.g., labor may be more unionized in one firm than in the other). In the Appendix we treat them jointly, but we report comparative statics on each dimension separately.

The top panel of Figure 4 illustrates the case of asymmetric marginal costs. In particular, it uses the same parameter values as in Figure 2, but introduces a (constant) marginal cost of 0.2 for player 2. The figure presents the equilibrium paths for different initial targets. The general pattern is similar to the baseline 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 by the thin solid line. The reason for this is that asymmetric marginal costs introduce asymmetries in commitment

⁹ Note that if the initial targets were very low and the adjustment parameters were high, one could also see a fully increasing equilibrium path.


¹⁰ It 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, stretching of time allows for more opportunities to adjust behavior.

¹¹ This limit is invariant to initial targets: the parameter A of the value function converges to approximately 0.0925, while all other parameters approach zero.

¹² Fourteen percent is due to higher production and lower equilibrium prices, whereas 3% is due to adjustment costs.

¹³ Although not the focus of our analysis, the bottom panel of Figure 3 also shows that final production levels decrease with θ , suggesting a likely discontinuity at $\theta = 0$, as is the case in similar models of convex adjustment costs (Fershtman and Kamien, 1987; Driskill and McCafferty, 1989).

FIGURE 3
COMPARATIVE STATICS IN THE BASELINE MODEL




This figure plots the equilibrium path in the baseline model, when parameters are set to $a = b = 1, c = 0$, and $\theta = 1$, and initial production targets are $1/3$ (the Cournot level). The top panel does so for different horizons (100, 50, 10, and 1), and shows that as the horizon gets longer, players have more time to smooth out their initial increase in targets, therefore peaking at higher levels. Once the deadline gets closer, however, this higher buildup declines faster. Final production levels do not change by much, unless the horizon is very short (as is the case when $T = 1$). The bottom panel plots the equilibrium path for different values of the adjustment cost parameters, θ (0.1, 1, and 10). As adjustment costs are lower, production targets peak higher, as it is both cheaper to achieve these levels and lower targets do not provide sufficient 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 at 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.

FIGURE 4

ASYMMETRIC PLAYERS



This figure plots the equilibrium path in a two-player model with asymmetric players. Parameters are set to $a = b = 1$. In the top panel, the asymmetry is in the marginal cost of production, c . That is, $\theta = 1$ for both players, one player has zero marginal costs ($c_1 = 0$), and the other has positive marginal costs ($c_2 = 0.2$). The figure plots three different cases, for different initial production targets: with thin solid lines they start at the Cournot level ($q_1 = 0.4, q_2 = 0.2$), with dashed lines we reverse these initial targets ($q_1 = 0.2, q_2 = 0.4$), and with thick solid lines they start with identical initial targets ($q_1 = q_2 = 0.3$). As the horizon is reasonably long, in all cases the lower marginal cost player eventually gains higher market share. Her market share is monotone in her initial production target. In the bottom panel, the asymmetry is in adjustment costs. That is, $c = 0$ for both players, but one player has a higher adjustment cost parameter than her opponent. As discussed in the text, when initial conditions are sufficiently low (as in the figure), the more flexible player is able to obtain higher market shares. It is not clear whether 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 achieve any commitment advantage.

The bottom panel of Figure 4 presents the case of asymmetric adjustment costs for different values of the θ coefficients. The shape of the equilibrium 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

targets so rapidly (recall that initial targets 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 targets further as a way to commit to high output.¹⁴

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 natural to think of adjustments becoming more expensive as the production date gets closer. As an example, hiring temporary labor three months before production may be cheap, whereas labor availability one day before production is scarce, and will require higher wages or higher search costs on the employer's part.

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

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

where $\theta(t)$ is continuous in t . The derivation of the system of ordinary differential equations is the same as in equation (15), 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 slowly. The qualitative predictions of the model remain unchanged.

Uncertainty. In the presence of uncertainty, there is a general tradeoff 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 demand can be either high or low, and that it (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 of demand is not particularly informative about demand at time T . Given that firms only care about the eventual realization of demand, on equilibrium they start by having a similar behavior independently of the current state. As the production date draws near, however, firms become more responsive to changes in the state of demand. This typically results in upward (downward) adjustments of 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 baseline model, and therefore build up commitment more slowly. Whereas the equilibrium path is random as it depends on the realization of uncertainty, the expected equilibrium path (computed numerically) exhibits a nonmonotonic pattern as in the baseline model.

□ **Repeated interaction and endogenous initial targets.** 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 baseline 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. This extension also allows us to endogenize the initial targets, which were taken as given in the baseline model.

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 targets $q_1(0)$ and $q_2(0)$ for next period, but pay a cost of $\frac{\alpha}{2}(q_i(0) - y_i)^2$ when they do so. For the next T units of time, they play the

¹⁴ Note that if the initial inherited targets were higher, say $q(0) = 0.4$, and the adjustment costs high as well, the previous result could be reversed. In this case, the nonflexible player would be at a credible position not to change her target far away from 0.4, which would force the flexible player to adjust downward.

baseline model with inherited initial targets of $(q_1(0), q_2(0))$ 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 targets 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 baseline model and thus satisfies equation (15). The boundary condition is different in this case, as 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, player i sets her initial production target $q_i(0)$ to satisfy

$$\max_{q_i} (A_0 + B_0q_i + C_0q_j(0) + D_0q_i^2 + E_0q_j(0)^2 + F_0q_iq_j(0)) - \frac{\varphi}{2} (q_i - q_i(T))^2 \quad (18)$$

which leads to the following first-order condition:

$$B_0 + 2D_0q_i(0) + F_0q_j(0) - \varphi (q_i(0) - q_i(T)) = 0. \quad (19)$$

Equation (19), together with its analog for player j , provides a closed-form relationship between $(q_1(0), q_2(0))$ and $(q_1(T), q_2(T))$. Because, 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 (16) of the baseline model. The solution to equation (15) and this new boundary condition constitute an MPE of the repeated game. Finally, we focus on the steady state of the equilibrium, in which the production decisions (but not production targets) are constant.


The equilibrium is computed by numerically searching for a solution. We start with a guess for A_T, \dots, F_T , and then iterate the law of motion in equation (15) to obtain A_0, \dots, F_0 . Then, using the boundary condition, new values for A_T, \dots, F_T are obtained. 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. We have checked for a wide range of parameters that the procedure converges (rapidly) and that the fixed point does not depend on the starting values chosen. Thus, on numerical grounds, we believe that the repeated interaction game has a unique symmetric MPE, or at least a unique symmetric MPE with linear-quadratic value functions.¹⁵

The top panel of Figure 5 presents the equilibrium path for the baseline 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 baseline model. The production levels are now higher than what would be produced in the baseline model if the inherited targets were the ones from the steady-state equilibrium. This is because, in addition to the commitment

¹⁵ To our knowledge, there does not exist a general existence or uniqueness result for linear-quadratic games with infinite horizon. Thus, one needs to establish these results on a "case-by-case" basis (as in Reynolds, 1987; Driskill and McCafferty, 1989; Jun and Vives, 2004). However, our model is not totally standard, as we nest a continuous-time adjustment stage within a discrete-time infinite-horizon model. For this reason, one cannot have closed-form equations for the law of motion, which are likely needed to derive certain sufficient stability conditions.

FIGURE 5

REPEATED INTERACTION MODEL



This figure plots the equilibrium path of the (symmetric) infinite-horizon repeated interaction game. In the top panel, parameters are set to $a = b = 1, c = 0, \theta = 1, T = 10, \varphi = 0.1$, and $\beta = 0.9$. Here, initial targets are determined endogenously as part of the (stationary) equilibrium path. In the bottom panel, we concentrate on a single period (which repeats itself forever), and present comparative statics with respect to the two adjustment cost parameters, φ and θ (the rest of the parameter values remain as before). As one can observe, φ mainly affects the initial targets whereas θ mainly affects the dynamic pattern of targets. As noted in the text, the case of $\varphi = 0$ is a special case in which the equilibrium of the repeated game is identical to the baseline model with free initial targets. Note, however, that even small values for φ are sufficient to generate a nonmonotonic pattern. This is because continuation values do not change much with initial targets.

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, 1987; Reynolds, 1987, 1991; Driskill and McCafferty, 1989; 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 . The bottom panel of Figure 5 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 targets for the subsequent production period, with high values of φ implying small jumps. In contrast, θ primarily affects the shape of the production target 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 baseline model with free initially chosen targets. That is, at $t = 0$, players decide simultaneously and costlessly on their initial targets ($q_1(0)$, $q_2(0)$) and then continue playing as in the baseline 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 targets give rise to an equilibrium path in which production targets are flat at $t = 0$ and gradually decline thereafter (this is also one of the cases presented in the bottom panel of Figure 5).¹⁶ For any $\varphi > 0$, however, the equilibrium path presents the hump-shaped pattern emphasized throughout.

3. Evidence

■ **Data.** We use data on domestic production targets of the major auto manufacturers in the United States.¹⁷ 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.¹⁸ 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 1965–1995, for a total of 372 production months.¹⁹ Every time a production target is published, it includes production targets for *all* three manufacturers. Thus, manufacturers do not decide when to post their targets, as this is requested by *Ward's*. Overall, we observe 1621 production targets for each manufacturer.²⁰ This amounts to an average of 4.42 production targets per production

¹⁶ This path is initially flat because, in equilibrium, initial production targets (q_i^* , q_j^*) satisfy $\frac{\partial V_i^0(q_i^*, q_j^*)}{\partial q_i} = 0$. From equation (10), 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 $x_i^0(q_i^*, q_j^*) = 0$.

¹⁷ For additional details, see Doyle and Snyder (1999), who use the same data.


¹⁸ These targets are being described by various synonyms: “assembly targets,” “assembly schedules,” “production plans,” “production forecasts,” and so forth.

¹⁹ Some of the observations in the data include postproduction revisions. We discard these observations. We only focus on targets posted before production. Five production months have no preproduction targets, and are therefore omitted from the analysis.

²⁰ The 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 had a small market share (2.3% on average) and exhibits a similar pattern to the Big Three, with the exception of its last three years of operation, during which its market share, production, and production targets rapidly declined. The qualitative results of the article remain unchanged if we use pre-1984 AMC data.

FIGURE 6

TIMING OF PRODUCTION TARGET REPORTS



This is a histogram of the number of days before production of each report in our sample. There are 1621 reports in our sample, which correspond to 372 distinct production months. The histogram shows 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 (more than four months in advance) are not as regular.

month, ranging from some cases with a single production target to others with up to 12 associated targets.²¹

Figure 6 presents the total number of published targets made at each ten day interval prior to actual production.²² It shows that production targets are published approximately once a month, often during the last week of the month. One can also observe that the number of observations is quite stable over the three to four months before production. There are significantly fewer earlier observations.

□ **Linking the data to the theory**

Are production decisions strategic substitutes? Our theoretical model assumes a Cournot structure for the product market competition. Of course, this is a simplification. Competition in the auto industry is more complex, and involves production, inventory management, distribution, pricing to consumers, and dealer incentives. In this section, we argue that on both theoretical and empirical grounds, the assumptions of the theory provide a reasonable approximation to the production planning decisions made by the Big Three. We start by noting that the qualitative prediction of the theory does not rely on a model of quantity competition *per se*. Rather, it only requires that production decisions are strategic substitutes. That is, one should interpret our theoretical profit function as a reduced form summarizing the subsequent stages of competition. The key for the theoretical results to hold is that, within this reduced form, production decisions, the variables whose targeting we observe, are strategic substitutes.

²¹ The 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. The frequency is also higher for months in the end of quarters, as important strategic scheduling decisions are often made on a quarterly basis.

²² Because 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.

But should one expect production decisions to be strategic substitutes? From the perspective of economic theory, the simplest way to address this question is to use the well-known static framework of Kreps and Scheinkman (1983). They show that even if the product market clears through prices, first-stage capacity decisions are, as in Cournot, strategic substitutes. More generally, Athey and Schmutzler (2001) show that in a variety of market clearing models, irrespective of whether product market choices are strategic substitutes or complements, first-stage investment decisions act as strategic substitutes. It is natural to view production decisions as such investments. Of course, dynamic considerations, through inventories, play an important role in the auto industry.²³ Judd (1996) considers a simple dynamic oligopoly with firms making production decision but then competing on prices. Inventories are the dynamic state variable linking periods. Judd shows that, as long as there are adjustment costs in production, production decisions act as strategic substitutes.²⁴ All these results support our modelling assumption that production decisions are strategic substitutes.

In addition to the theoretical arguments, one can directly test for strategic substitutability empirically. Berndt, Friedlander, and Wang Chiang (1990), for example, cannot reject Cournot as a model for competition among the Big Three. We reach a similar conclusion performing an additional test with our data. For this purpose, we use auxiliary data on labor strikes also reported in *Ward's*.²⁵ We regress the actual monthly production of one firm on a dummy variable that is equal to one if there was a labor strike during that month. We find that a strike in one firm reduces its own production, which is to be expected, but increases production by its rivals. If production decisions were strategic complements, rivals would decrease their production instead. Similarly, we test for the sign of the reaction function by regressing one firm's production on the production of its rivals. To solve an obvious endogeneity problem, we use the strike dummy variables as instruments. Again, we find downward-sloping reaction functions, consistent with strategic substitutes, although significance levels are lower.

Another important aspect of the industry that we have abstracted from is the presence of product differentiation and multiproduct firms. We think that explicitly modelling these features is unlikely to change the nature of the theoretical predictions. Firm-level vehicle production is the sum of the firm's production in all market segments. As long as our theory provides a good approximation for each segment separately, its qualitative predictions should also hold at the aggregate.

Actual versus reported production targets. In principle, there is a distinction between actual production targets, the object of the theoretical model, and reported production targets, the data used in the empirical analysis. The maintained assumption in our empirical analysis is that reported targets are a good proxy for real decisions. This is an important assumption: if these reported targets were not anchored to any real decision, they would constitute pure cheap talk. Therefore, we discuss this issue in more detail below.

First, we do not think that *Ward's* plays an essential role in facilitating the strategic interaction. According to our conversations with many industry experts, including *Ward's* analyst, it is difficult to hide real actions from competitors. Hiring temporary labor, scheduling extra shifts, temporary shutdowns of plants, or orders of large numbers of windshields can be easily monitored. In this sense, we view the production targets reported by *Ward's* as providing monthly snapshots of the underlying, more continuous decisions made by manufacturers regarding their production targets.

Second, production target figures reported by *Ward's* are considered to be one of the most reliable sources of information regarding actual production targets. Industry analysts and consultants use the reported targets in *Ward's* as a key input in their production forecasting

²³ See Kahn (1992) and Bresnahan and Ramey (1994) for theory and evidence about the relationship between sales and production.

²⁴ See also Jun and Vives (2004), who report similar results in a more general framework.

²⁵ See Doyle and Snyder (1999) for more details on these strike data.

models, the press often quotes these reported targets, and part suppliers use these targets as a way to verify the corresponding orders they receive from manufacturers directly.²⁶ It is natural to wonder why firms truthfully and publicly report such internal decisions. There are several related reasons. First, as we mentioned above, plans cannot be easily hidden, neither from competitors, nor from *Ward's*, the press, or other external analysts. The main role of *Ward's* is to report these actions to third parties (suppliers, dealers, etc.), who cannot perform the monitoring so easily.²⁷ As we describe in the next section, third parties' plans crucially depend on this information. Second, manufacturers are aware of this monitoring, making it costly for them to consistently report targets that misrepresent actual decisions. These costs can take the form of bad press, as well as reduced ability to control (future) information flows.²⁸ Third, strategic considerations may also work toward providing incentives for truthful reporting. If commitment is achieved by credible production targets, then credibility can only be achieved by a reputation for truthful reporting.

Finally, we note that our focus on the qualitative predictions of the theory makes the empirical exercise meaningful even if reported targets do not perfectly describe actual decisions. As long as reported and actual decisions are positively correlated and this correlation structure is time invariant, the qualitative prediction of the theory regarding actual production targets should hold for the reported ones as well.

Sources of adjustment costs. One of our key assumptions is that the adjustment of production targets is costly, as this allows production targets to serve as a commitment device. That is, the model assumes that production targets are associated with real actions that cannot be costlessly reversed. These costs do not have to be large. In fact, as mentioned earlier, lower costs lead to a more pronounced hump shape in production targets, as long as they are positive. Here we discuss the possible sources of such adjustment costs in the context of the empirical application. This discussion draws on interviews with many individuals in the auto industry, as well as on information from newspaper and trade journal articles.²⁹

Production scheduling is an important task for auto manufacturers. Considerable amounts of labor and resources are devoted to this function. Manufacturers are continuously taking actions that affect their future production capabilities: they produce certain parts internally, order others from suppliers, hire and fire temporary labor, cancel vacations, schedule plant shutdowns, and so forth. Each adjustment in plans generates responses along all these margins, which are often costly to manufacturers.

The relationship with suppliers is an important dimension of these costs. When orders change, suppliers adjust.³⁰ They incur adjustment costs in production. These include costs associated with changes in their own labor and plant operation schedule, hiring and firing temporary labor, excess inventory holding, the requirement to buy materials in the spot market at higher prices, or the opportunity cost of idle capacity. Some of these costs are only borne

²⁶ An online appendix (available on our web pages) provides both quotes from the press that mention *Ward's* reported targets as well as quotes from our conversations with industry players who refer to these targets.

²⁷ One 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* as follows: "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 companies' futures" (wardsauto.com/war/index.htm).

²⁸ As an example, for a certain period in the late 1990s, Ford was known to provide targets that were "too high." Industry players then ignored Ford's numbers and formed their own forecasts. It took a big effort by Ford to gain its credibility back. The online appendix provides some related quotes.

²⁹ Selected quotes from both interviews and newspaper articles are provided in the online appendix.

³⁰ See the online appendix for evidence on this issue. If one were to take our model literally, suppliers should foresee the final equilibrium production levels and not follow manufacturers all along the adjustment process. In practice, however, targets are also affected by various sources of uncertainty, such as demand conditions. If some new information about such sources is only revealed to manufacturers, suppliers' optimal response to a change in production targets is to at least partially adjust to it.

by suppliers, but others are passed on to manufacturers through different channels. First, order changes with too short notice result in less timely deliveries, and are occasionally not even fulfilled (due to capacity constraints along the vertical supply chain). Second, 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 seldom happens, as contracts often stipulate clauses that deal with these instances. Typical part contracts in the industry specify minimum and maximum 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. Because parts are complements in production,³¹ 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 ones, carries an implicit adjustment cost, as it would have been cheaper if previous contracts had been set differently. Third, there are also dynamic consequences. Order changes lead suppliers to increase their projected cost for subsequent models, making manufacturers face less competitive bidding in the future. Finally, there is the increased risk of pushing suppliers into bankruptcy if orders are cut substantially.³²

Suppliers of parts are not the only channel through which manufacturers incur adjustment costs. Another important channel is the production of parts produced by the manufacturers themselves. There are substantial costs associated with scheduling adjustments to labor on the assembly lines.³³ Moreover, even though a large fraction of parts are currently produced by separate part suppliers, the industry was almost fully vertically integrated until the early 1980s (Scherrer, 1991). The adjustment costs incurred by part suppliers when production targets change are also incurred by manufacturers when they produce these parts in house.

Financial markets provide a final channel through which changes in production targets may be costly to manufacturers. As we document in the online appendix, changes in production targets are closely followed by the media and financial analysts. They are often interpreted as bad news, regardless of whether targets are adjusted upward or downward. This gives manufacturers another reason to be reluctant to make such changes.

□ **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.^{34, 35} Our key

³¹ Consider, for example, an O-ring production function (Kremer, 1993).

³² See a related discussion on this issue in Ben-Shahar and White (2006).

³³ With a large fraction of workers in the industry unionized, workers on temporary layoffs normally collect company-funded layoff protection plans.

³⁴ This 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.

³⁵ There are six instances of extreme outliers. Five of them are due to unexpectedly low Q_{it} 's, which generate highly a_{it}^d 's, more than three times eventual production ($a_{it}^d > 2$). The sixth instance is of zero announcements by Chrysler. Although 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

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 the production date gets closer.

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 though they are made in the same context. We then use quartic (biweight) kernel regressions of a_{it}^d on d to nonparametrically 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 separately, for the Big Three average $a_{Big3,t}^d = \frac{1}{3}(a_{GM,t}^d + a_{Ford,t}^d + a_{Chrysler,t}^d)$, 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 7, which pools all production months in the data. The qualitative picture is of a nonmonotonic pattern. On average, production targets start about 5% above eventual production levels and gradually increase. They peak two to three months before production at about 10%, and then gradually decline toward actual production levels. This pattern is not uniform across manufacturers. Whereas Ford and Chrysler, the two smaller firms, follow a similar nonmonotonic pattern of production targets, GM exhibits a different behavior. GM's average initial production target is about 15% 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 targets 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.³⁶

The dashed lines in Figure 7 report 95% 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 toward the production deadline is quite precisely estimated. This pattern is extremely consistent across manufacturers and for different subsamples. Figure 7 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 6: the number of observed production targets three to six months before production is significantly smaller than the number of observations zero to three months before production.

Our theoretical prediction concerns a nonmonotonic pattern of production targets with respect to the same production month. A potential concern may be that whereas 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.³⁷ To address this concern, we repeat the same exercise for different subsamples of the data. Figure 8 divides the sample into

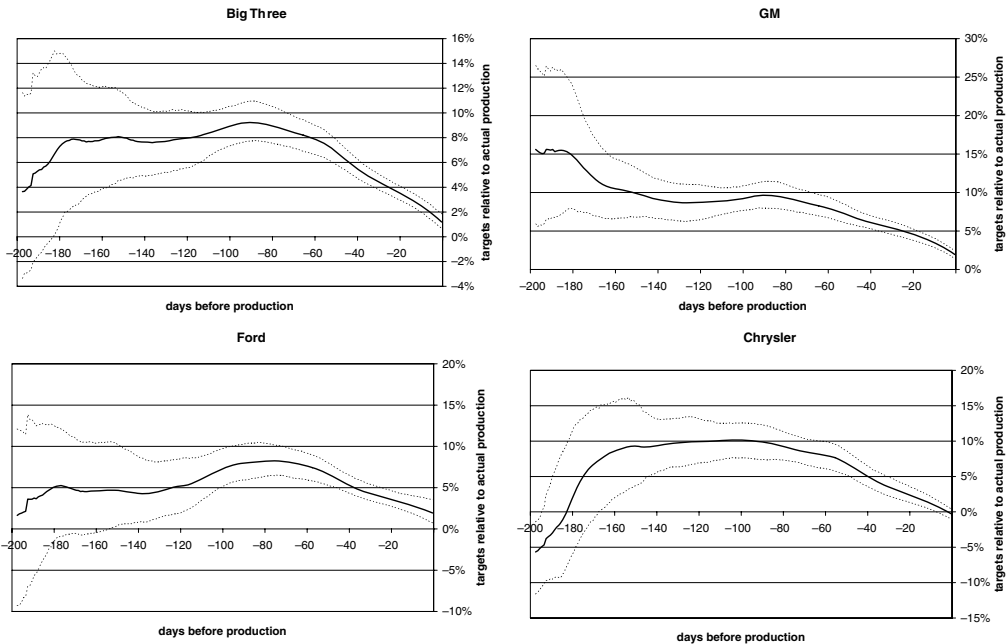
with the same production month. This leaves us with 361 production months and 1598 targets by each manufacturer for the empirical analysis.

³⁶ We note that the nonmonotonic pattern is primarily driven by months with positive production growth relative to the previous month (which happens in slightly less than half the data). This is consistent with the idea—formalized by the positive φ parameter in the repeated interaction game—of linkages between a month's actual production level and the production targets for subsequent months. When production declines, it is less costly to report sufficiently high initial targets.

³⁷ For 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 nonmonotonicity even though none of the individual patterns is such.

FIGURE 7

PRODUCTION TARGETS OVER TIME, POOLING ALL DATA



The 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 Big Three separately, as well as for the (unweighted) average. Each series is based on 1598 observations. All estimates use a bandwidth of 30 days. The dashed lines present 95% confidence intervals computed by bootstrapping the data, running the same kernel regression on each bootstrapped sample, and taking a 95% confidence interval point by point.

three decades.³⁸ Figure 9 performs the analysis for each calendar month separately to account for potential seasonal variation.³⁹ All these exercises show similar qualitative patterns. First, the pattern of declining production targets during the last two to three months before production is 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.

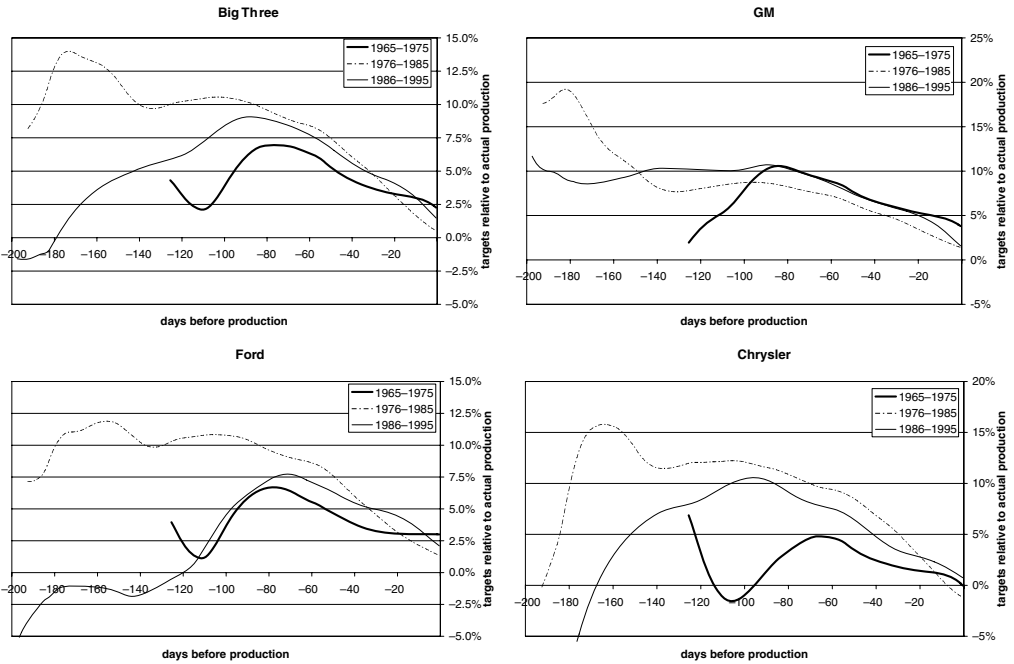
As already discussed, the nonmonotonic pattern predicts a positive slope of a_{it}^d with respect to d early on, and a negative slope toward production. In order to focus on this precise prediction of the theory, we analyze directly the directional change in production targets, rather than their levels. To do so, 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

³⁸ Over the period we analyze, there have been many changes to the industry which greatly affected production planning. These include a dramatic loss in market share by the Big Three to foreign manufacturers, an important increase in the use of external part suppliers, and an increase in model differentiation that led to less flexibility in production.

³⁹ Seasonal variation may occur due to model-year product-life-cycle effects (Copeland, Dunn, and Hall, 2005). As new model-years are introduced (typically during July and August), excess inventory of old models becomes significantly more costly to manufacturers.

FIGURE 8

PRODUCTION TARGETS OVER TIME, BY DECADE



This figure repeats the exercise performed in Figure 7 (see its notes for details), but here the kernel regressions are run for each ten year period separately. The estimates for the first decade (1965–1975) only start about 120 days before production, as during this period there were no earlier production target reports. We do not present confidence intervals for presentation reasons, but they are noticeably greater than those in Figure 7, as each line uses only one third of the data.

12 frequencies except one are greater than 0.5. None of them is significantly lower than 0.5 and the majority of them are significantly higher. This is all consistent with the theory, and supports the existence of a nonmonotonic pattern. Second, we define the percentage change, per day, in production targets by

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

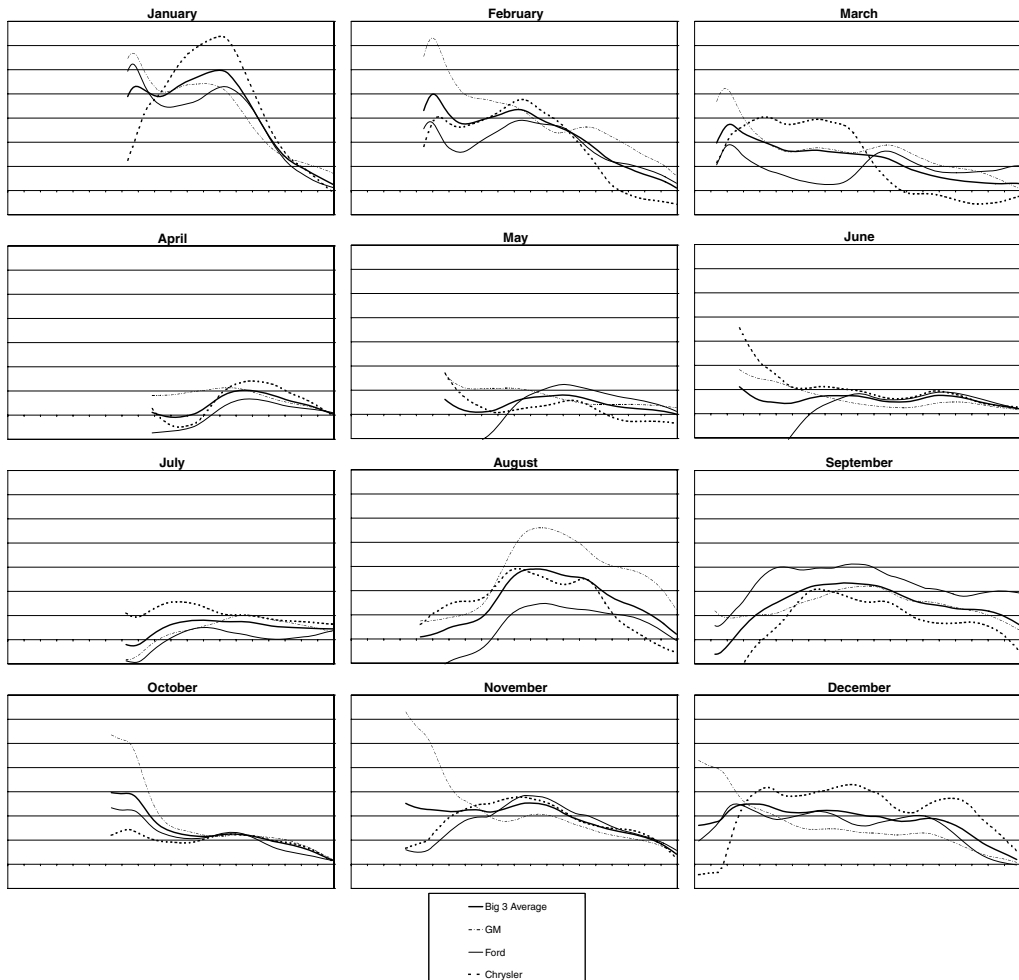
where A_{it}^d and $A_{it}^{d'}$ are two consecutive production targets associated with the same production month. We then run similar kernel regressions of s_{it}^d with respect to d . Figure 10 reports these regressions. One can observe that in all cases the slope of production targets is positive 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 nonmonotonic pattern.

□ **Discussion of alternative explanations.** Production targets follow a hump shape over time: they start low, then increase, and eventually decrease to actual production levels. This is consistent with the main prediction of the theoretical section of the article.

One possible concern is that the theory can, in principle, give rise to either a decreasing or a hump-shaped pattern. Thus, only an empirical pattern that is increasing close to the production deadline would be inconsistent with the theory. Although we empirically do not find such a pattern, we do not consider this sufficiently interesting in the context of the theory. One can think of a

FIGURE 9

PRODUCTION TARGETS OVER TIME, BY CALENDAR MONTH



This figure repeats the exercise performed in Figure 7 (see its notes for details), but here the kernel regressions are run for each calendar month separately to account for potential seasonality arising from model-year product-life-cycle effects. All figures use the same scale, with the horizontal axis running from 200 days before production to 0, and the vertical axis running from -5% to 35% . We do not present confidence intervals for presentation reasons, but they are much greater than those in Figure 7 or 8 as each line uses only about $1/12$ of the data.

variety of theoretical models delivering production targets that are higher than eventual production levels. One example involves asymmetric adjustment costs in the presence of uncertainty. If adjusting targets downward is cheaper than adjusting them upward, overtargeting would have an option value. If uncertainty is resolved gradually, then production targets would slowly decrease toward actual production levels. A second explanation builds on a notion of optimism or overconfidence. If, as often argued in the press, Big Three executives are overoptimistic and reality only sinks in gradually, a similar monotone decreasing pattern would arise. Neither story, however, would give rise to a hump.

We therefore feel that it is important to emphasize the hump-shaped pattern of production targets observed in the data. It is significantly more difficult to come up with alternative explanations for this nonmonotonic pattern. One could roughly think of two classes of theories

TABLE 1 Frequency Estimates of Revision Signs

	$Pr(A_{it}^{Middle} > A_{it}^{Early})^a$	$Pr(A_{it}^{Late} < A_{it}^{Middle})^a$	$Pr(Q_{it} < A_{it}^{Late})^a$
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
Observations	135	286	359

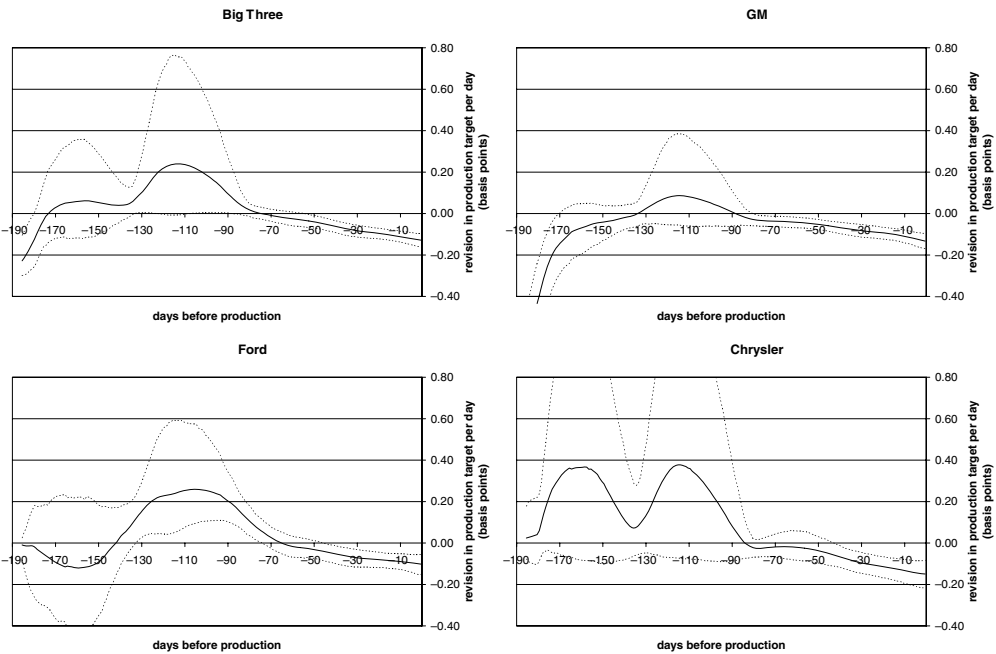
*Significantly different from 0.5 at 95% confidence level.

^aFor 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 does not affect the results.

The table reports frequency tests for whether the average early, middle, and late targets conform with a hump-shaped pattern. The inequalities (top row) 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. Indeed, all numbers but one are greater than 0.5, the majority of them are significantly greater than 0.5, and none is significantly less than 0.5.

FIGURE 10

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 pattern is presented for each of the Big Three separately, as well as for the (unweighted) average. Each series is based on 1239 observations (taking first differences, we lose the earliest report for each production month). All estimates use a bandwidth of 30 days. The dashed lines present 95% confidence intervals computed by bootstrapping the data, running the same kernel regression on each bootstrapped sample, and taking a 95% confidence interval point by point.

delivering it. One includes theories that have a nonmonotone variable as one of their foundations. For example, suppose that Big Three executives are realistic far in advance, then gradually become overoptimistic, and then close to the production date become realistic again. Of course, this would yield a hump shape. A second class of theories are those that involve two offsetting forces, with each dominating at a different time. For example, suppose there is uncertainty and

asymmetric adjustment costs as described above, coupled with inherent pessimism by executives. If pessimism dominates early but adjustment costs dominate late, this would lead to a hump shape. Although our empirical findings cannot rule out these alternative stories, one could argue that Occam’s razor principle should favor our theory. This encourages us to view our findings as empirical support for the existence of adjustment costs, and for the relevance of the strategic role of preproduction preparations in determining final production decisions.

4. Concluding remarks

■ This article studies the dynamics of preproduction 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 look for the main predictions of the theory in data on production targets in the U.S. auto industry. The observed production targets exhibit a nonmonotonic pattern similar to the one predicted by the theory.

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 more structural estimation approach. The hump-shaped pattern we find in this article would help in identifying strategic effects from uncertainty. Estimating structural parameters would be interesting for policy purposes, as one could quantify the intensity of competition (estimating how far the equilibrium is from the Cournot levels) or perform welfare analysis.⁴⁰ We leave this exercise for future work.

On a methodological level, we think that our exercise illustrates the empirical potential of nonstationary 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.

Appendix

■ This appendix derives the equations that describe the solutions to three extensions of the baseline model, as discussed in Section 2.

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

$$\max_{x'_i} \left(-\frac{\theta}{2} (x'_i)^2 + \frac{\partial V'_i}{\partial q_i} x'_i + \sum_{j \neq i} \frac{\partial V'_i}{\partial q_j} x'_j + \frac{\partial V'_i}{\partial t} \right) = 0. \tag{A1}$$

The first-order condition for x'_i implies

$$x'_i = \frac{1}{\theta} \frac{\partial V'_i}{\partial q_i}. \tag{A2}$$

We can now substitute this back into equation (A1), as well as the symmetric solution for all other x'_j ’s, rearrange, and obtain the following differential equation:

$$0 = \frac{1}{2\theta} \left(\frac{\partial V'_i}{\partial q_i} \right)^2 + \frac{1}{\theta} \sum_{j \neq i} \left(\frac{\partial V'_i}{\partial q_j} \right) \left(\frac{\partial V'_i}{\partial q_j} \right) + \frac{\partial V'_i}{\partial t}. \tag{A3}$$

⁴⁰ The 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 “wasted” adjustment costs during the planning phase.

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

$$\begin{aligned} V_i^t(q_i, q_j) &= A_t + B_t q_i + \sum_{j \neq i} C_t q_j + D_t q_i^2 + \sum_{j \neq i} E_t q_j^2 + \sum_{j \neq i} F_t q_i q_j + \sum_{j \neq i} \sum_{k \neq i, j} G_t q_j q_k \\ &= A_t + B_t q_i + C_t Q_{-i} + D_t q_i^2 + E_t R_{-i} + F_t q_i Q_{-i} + G_t S_{-i}, \end{aligned} \tag{A4}$$

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

$$x_i^t(q_i, q_j) = \frac{1}{\theta} (B_t + 2D_t q_i + F_t Q_{-i}). \tag{A5}$$

Thus, we can rewrite equation (A3) as

$$\begin{aligned} 0 &= \frac{1}{2\theta} \left(B_t + 2D_t q_i + F_t Q_{-i} \right)^2 + \frac{1}{\theta} \sum_{j \neq i} (C_t + 2E_t q_j + F_t q_i + 2G_t(Q_{-j} - q_i)) \\ &\quad \times (B_t + 2D_t q_j + F_t Q_{-j}) \\ &\quad + (A_t' + B_t' q_i + C_t' Q_{-i} + D_t' q_i^2 + E_t' R_{-i} + F_t' q_i Q_{-i} + G_t' S_{-i}). \end{aligned} \tag{A6}$$

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

$$\begin{pmatrix} A' \\ B' \\ C' \\ D' \\ E' \\ F' \\ G' \end{pmatrix} = \frac{1}{\theta} \begin{pmatrix} \frac{1}{2} B^2 + (N-1)BC \\ 2BD + BF(N-1) + CF(N-1) \\ BF + 2BE + 2CD + CF(N-2) + 2BG(N-2) \\ 2D^2 + F^2(N-1) \\ \frac{1}{2} F^2 + 4DE + 2FG(N-2) \\ 4DF + 2EF + F^2(N-2) + 2FG(N-2) \\ \frac{1}{2} F^2 + 2EF + 4GD + 4FG(N-3) \end{pmatrix} \tag{A7}$$

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

$$\begin{pmatrix} A_T \\ B_T \\ C_T \\ D_T \\ E_T \\ F_T \\ G_T \end{pmatrix} = \begin{pmatrix} 0 \\ a - c \\ 0 \\ -b \\ 0 \\ -b \\ 0 \end{pmatrix}. \tag{A8}$$

□ **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_i^t to F_i^t denote i 's value function coefficients.

One can start by following the same steps as in Section 2. The first difference appears in equation (11), which now reads

$$\frac{1}{2\theta^i} \left(\frac{\partial V_i^t}{\partial q_i} \right)^2 + \frac{1}{\theta^j} \left(\frac{\partial V_i^t}{\partial q_j} \right) \left(\frac{\partial V_j^t}{\partial q_j} \right) + \frac{\partial V_i^t}{\partial t} = 0. \tag{A9}$$

The value function for each player is

$$V_i^t(q_i, q_j) = A_i^t + B_i^t q_i + C_i^t q_j + D_i^t q_i^2 + E_i^t q_j^2 + F_i^t q_i q_j. \tag{A10}$$

Substituting it into equation (A9) gives

$$\begin{aligned} 0 &= \frac{1}{2\theta^i} (B_i^t + 2D_i^t q_i + F_i^t q_j)^2 + \frac{1}{\theta^j} (C_i^t + 2E_i^t q_j + F_i^t q_i) (B_j^t + 2D_j^t q_j + F_j^t q_i) \\ &\quad + (A_i^{t'} + B_i^{t'} q_i + C_i^{t'} q_j + D_i^{t'} q_i^2 + E_i^{t'} q_j^2 + F_i^{t'} q_i q_j). \end{aligned} \tag{A11}$$

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

$$\begin{pmatrix} A^{i'} \\ B^{i'} \\ C^{i'} \\ D^{i'} \\ E^{i'} \\ F^{i'} \end{pmatrix} = \begin{pmatrix} \frac{1}{2\theta^i} B^{i2} + \frac{1}{\theta^i} B^j C^i \\ \frac{2}{\theta^i} B^i D^i + \frac{1}{\theta^j} B^j F^i + \frac{1}{\theta^i} C^i F^j \\ \frac{1}{\theta^i} B^i F^i + \frac{2}{\theta^j} B^j E^i + \frac{2}{\theta^i} C^i D^j \\ \frac{2}{\theta^i} D^{i2} + \frac{1}{\theta^j} F^i F^j \\ \frac{1}{2\theta^i} F^{i2} + \frac{4}{\theta^j} D^j E^i \\ \frac{2}{\theta^i} D^i F^i + \frac{2}{\theta^j} D^j F^i + \frac{2}{\theta^i} E^i F^j \end{pmatrix} \tag{A12}$$

with the boundary condition given by

$$\begin{pmatrix} A_T^i \\ B_T^i \\ C_T^i \\ D_T^i \\ E_T^i \\ F_T^i \end{pmatrix} = \begin{pmatrix} 0 \\ a - c^i \\ 0 \\ -b \\ 0 \\ -b \end{pmatrix}. \tag{A13}$$

□ **Uncertainty** We follow the same steps as in Section 2 with a few modifications. Now there are two value functions, depending on the state of demand. Let these two value functions be V_L and V_H . Thus, the Bellman equation if one is in state L is

$$\max_{x_i^t} \left(-\frac{\theta}{2} (x_i^t)^2 + \frac{\partial V_{L,i}^t}{\partial q_i} x_i^t + \frac{\partial V_{L,i}^t}{\partial q_j} x_j^t + \frac{\partial V_{L,i}^t}{\partial t} + \lambda (V_{H,i}^t(q_t, q_j) - V_{L,i}^t(q_t, q_j)) \right) = 0 \tag{A14}$$

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

$$x_i = \frac{1}{\theta} \frac{\partial V_{L,i}^t}{\partial q_i} \tag{A15}$$

and symmetrically for H . Now one can obtain the corresponding differential equations as in equations (11) and (14), resulting in a system of 12 ordinary differential equations (ODEs). The law of motion for the coefficients associated with the L state is

$$\begin{pmatrix} A'_L \\ B'_L \\ C'_L \\ D'_L \\ E'_L \\ F'_L \end{pmatrix} = \lambda \begin{pmatrix} A_H - A_L \\ B_H - B_L \\ C_H - C_L \\ D_H - D_L \\ E_H - E_L \\ F_H - F_L \end{pmatrix} + \frac{1}{\theta} \begin{pmatrix} \frac{1}{2} B_L^2 + B_L C_L \\ 2B_L D_L + B_L F_L + C_L F_L \\ B_L F_L + 2B_L E_L + 2C_L D_L \\ 2D_L^2 + F_L^2 \\ \frac{1}{2} F_L^2 + 4D_L E_L \\ 4D_L F_L + 2E_L F_L \end{pmatrix}. \tag{A16}$$

Additional six analogous equations are associated with the H state. This structure is identical to the baseline 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.

References

ATHEY, S. AND SCHMUTZLER, A. "Investment and Market Dominance." *RAND Journal of Economics*, Vol. 32 (2001), pp. 1–26.

BENKARD, C.L. "A Dynamic Analysis of the Market for Wide-Bodied Commercial Aircraft." *Review of Economic Studies*, Vol. 71 (2004), pp. 581–611.

BEN-SHAHAR, O. AND WHITE, J.J. "Boilerplate and Economics Power in Auto Manufacturing Contracts." *Michigan Law Review*, Vol. 104 (2006), pp. 953–982.

BERNDT, E.R., FRIEDLANDER, A.F., AND WANG CHIANG, J.S. "Interdependent Pricing and Markup Behavior: An Empirical Analysis of GM, Ford, and Chrysler." Working Paper no. 3396, NBER, 1990.

BRESNAHAN, T.F. AND RAMEY, V.A. "Output Fluctuations at the Plant Level." *Quarterly Journal of Economics*, Vol. 109 (1994), pp. 593–624.

CARUANA, G. AND EINAV, L. "A Theory of Endogenous Commitment." *Review of Economic Studies*, Vol. 75 (2008), pp. 99–116.

COPELAND, A., DUNN, W., AND HALL, G. "Prices, Production, and Inventories over the Automotive Model Year." Working Paper no. 11257, NBER, 2005.

CYERT, R.M. AND DEGROOT, M.H. "Multiperiod Decision Models with Alternating Choice as the Solution to the Duopoly Problem." *Quarterly Journal of Economics*, Vol. 84 (1970), pp. 410–429.

DOYLE, M.P. AND SNYDER, C.M. "Information Sharing and Competition in the Motor Vehicle Industry." *Journal of Political Economy*, Vol. 107 (1999), pp. 1326–1364.

DRISKILL, R.A. AND MCCAFFERTY, S. "Dynamic Duopoly with Adjustment Costs: A Differential Game Approach." *Journal of Economic Theory*, Vol. 49 (1989), pp. 324–338.

- FERSHTMAN, C. AND KAMIEN, M.I. "Dynamic Duopolistic Competition with Sticky Prices." *Econometrica*, Vol. 55 (1987), pp. 1151–1164.
- FUDENBERG, D. AND TIROLE, J. "The Fat-Cat Effect, the Puppy-Dog Ploy, and the Lean and Hungry Look." *American Economic Review*, Vol. 74 (1984), pp. 361–366.
- HANIG, M. "Differential Gaming Models of Oligopoly." Ph.D. Dissertation, Massachusetts Institute of Technology, 1986.
- JUDD, K.L. "Cournot versus Bertrand: A Dynamic Resolution." Mimeo, Hoover Institution, 1996. Available from bucky.stanford.edu/papers/invold.pdf.
- JUN, B. AND VIVES, X. "Strategic Incentives in Dynamic Duopoly." *Journal of Economic Theory*, Vol. 116 (2004), pp. 249–281.
- KAHN, J.A. "Why Is Production More Volatile Than Sales? Theory and Evidence on the Stockout-Avoidance Motive for Inventory-Holding." *Quarterly Journal of Economics*, Vol. 107 (1992), pp. 481–510.
- KREMER, M. "The O-Ring Theory of Economic Development." *Quarterly Journal of Economics*, Vol. 108 (1993), pp. 551–575.
- KREPS, D.M. AND SCHEINKMAN, J.A. "Quantity Precommitment and Bertrand Competition Yield Cournot Outcomes." *Bell Journal of Economics*, Vol. 14 (1983), pp. 326–337.
- KYDLAND, F. "Noncooperative and Dominant Player Solutions in Discrete Dynamic Games." *International Economic Review*, Vol. 16 (1975), pp. 329–341.
- LAPHAM, B. AND WARE, R. "Markov Puppy Dogs and Related Animals." *International Journal of Industrial Organization*, Vol. 12 (1994), pp. 569–593.
- LUKES, D.L. "Equilibrium Feedback Control in Linear Games with Quadratic Costs." *SIAM Journal on Control*, Vol. 9 (1971), pp. 234–252.
- MASKIN, E. AND TIROLE, J. "A Theory of Dynamic Duopoly III: Cournot Competition." *European Economic Review*, Vol. 31 (1987), pp. 947–968.
- PAPAVASSILOPOULOS, G.P. AND CRUZ, J.B. "On the Uniqueness of Nash Strategies for a Class of Analytic Differential Games." *Journal of Optimization Theory and Applications*, Vol. 27 (1979), pp. 309–314.
- AND OLSDER, G.J. "On the Linear-Quadratic, Closed-Loop, No-Memory Nash Game." *Journal of Optimization Theory and Applications*, Vol. 42 (1984), pp. 551–560.
- REYNOLDS, S.S. "Capacity Investment, Preemption and Commitment in an Infinite Horizon Model." *International Economic Review*, Vol. 28 (1987), pp. 69–88.
- . "Dynamic Oligopoly with Capacity Adjustment Costs." *Journal of Economic Dynamics and Control*, Vol. 15 (1991), pp. 491–514.
- ROMANO, R. AND YILDIRIM, H. "On the Endogeneity of Cournot-Nash and Stackelberg Equilibria: Games of Accumulation." *Journal of Economic Theory*, Vol. 120 (2005), pp. 73–107.
- RYAN, S.P. "The Costs of Environmental Regulation in a Concentrated Industry." Mimeo, Massachusetts Institute of Technology, 2006. Available from econ-www.mit.edu/files/1166.
- SALONER, G. "Cournot Duopoly with Two Production Periods." *Journal of Economic Theory*, Vol. 42 (1987), pp. 183–187.
- SCHERRER, C. "Governance of the Steel Industry: The Transformation of Labor and Supplier Relations." In J.L. Campbell, J.R. Hollingsworth, and L.N. Lindberg, eds., *Governance of the American Economy*. Cambridge, UK: Cambridge University Press, 1991.