Research in Economics 70 (2016) 518–535

Contents lists available at ScienceDirect

Research in Economics

journal homepage: www.elsevier.com/locate/rie

Bagels and donuts for sale: A case study in profit maximization[☆]

Steven D. Levitt

University of Chicago, USA

ARTICLE INFO

Article history:
Received 23 October 2015
Accepted 1 November 2015
Available online 19 January 2016

Keywords:
Profit maximization
Price setting
Freakonomics

ABSTRACT

Profit maximizing behavior on the part of firms is a fundamental, but rarely tested, assumption of economics. In this paper, I analyze the decisions made by an MIT trained economist running a company that delivers bagels and donuts. The simplicity and transparency of the business (e.g. marginal cost is easily observed) allow for relatively direct tests of profit maximization in the quantities delivered each day and the prices that are charged. Using twelve years of data representing more than 80,000 deliveries, I find that the company is extremely adept and determining how many bagels and donuts to deliver to a particular customer on a given day. The company appears to price on the inelastic portion of the short-run demand curve for the entire period. These pricing choices are inconsistent with short-run profit maximization, although they can potentially be reconciled with a dynamic optimization model.

1. Introduction

The assumption of profit maximizing behavior by firms is one of the oldest, most fundamental, and widely applied in all of economics. Virtually all models of production start with profit maximization. The assumption of profit maximization is particularly central to modern empirical industrial organization techniques which rely on indirect identification approaches to overcome the fact that critical components of the inputs to a firm's decision (e.g. marginal cost) are not observable to the econometrician.[1]

In principle, testing for violations of profit maximization is straightforward. At least in the long run, no profit maximizing firm should set price below marginal cost, choose a price at which residual demand is inelastic, or stop production when marginal revenue exceeds marginal cost. Yet, direct attempts to empirically test the assumption of profit maximization are quite rare. One reason is that real-world firms are typically quite complex, producing multiple goods and using large numbers of inputs. As a consequence, detailed information on marginal cost is rarely available. Estimating the price elasticity of demand is also not a trivial task, although great progress has been made in this area in recent years. A few prior analyses have attempted to overcome these difficulties. Genesove and Mullin (1998) analyzes the sugar industry in the

☆ I would like to thank Gary Becker, Stephen Dubner, Paul Feldman, Jeremy Fox, Ali Hortacsu, Kevin Murphy, Ariel Pakes, and Chad Syverson for helpful conversations and comments. Paul Feldman generously provided the data used in the paper. Adam Castor, Andrew Francis, Marina Niessner, and Iolanda Palmieri provided outstanding research assistance. This research was supported by research grants from Sherman Shapiro and the National Science Foundation (Grant no. 452738).

[1] See, for instance, Rosse (1970), Pakes (1986), Bresnahan (1987), Bresnahan and Reiss (1991), Berry et al. (1995), and the dozens of papers that these seminal contributions have spawned.

http://dx.doi.org/10.1016/j.rie.2015.11.001
1090-9443/

period 1890–1914, arguing that the production technology used by various refiners was fixed and known, which allows them to measure marginal costs accurately. Wolfram (1999) and Hortacsu and Puller (2005) both study the electricity industry, in which the short run marginal costs are almost exclusively driven by fuel costs.

In this paper, I analyze data from a company that delivers donuts and bagels to Washington, DC area businesses. Workers at these firms buy these goods on the honor system, depositing their payment in a lock-box that is picked up later the same day.

The data generated by this business provide a unique window into the question of profit maximization.[2] There are a number of reasons why this firm would a priori be a leading candidate to maximize profits. First, the service the firm provides is very simple. The firm has only one line of business. The marginal cost of delivering one more bagel or donut to an existing customer is simply the wholesale price of the good, a number which is easily observed by the company's owner. Second, the firm gets frequent and detailed signals of demand. Each day, for each customer, the owner chooses the quantity of bagels and donuts to deliver. Later that day he observes how many of the goods go uneaten, as well as the revenue collected.[3] Third, the owner is well-trained in the principles of profit maximization. He is ABD in Economics from MIT (studying under Paul Samuelson in the mid-1960s), has published research in *Journal of Political Economy* (Feldman, 1971), and spent more than twenty years working as a professional economist before starting this business. Fourth, the owner both makes the decisions and is the residual claimant on the profit flows; there is no principal-agent problem at work.

Analyzing 13 years of data representing more than 80,000 deliveries of bagels and donuts to clients, the findings concerning profit maximization are mixed. The company is extremely adept at choosing the optimal quantities to deliver on a daily basis at a given price. The fact that the firm gets quantities almost exactly right, combined with the fact that the profit function is flat near the optimum, implies that there is little scope for improving profits in this dimension. Although I show that a regression model can improve on the firm's choices, it improves the bottom line by just a fraction of a percent.

In stark contrast, however, the prices set by the firm deviate substantially from those predicted by static pricing models. The firm prices on the inelastic part of the short-run demand curve over the entire period under study, i.e. price increases lead to an increase in both revenues and profits. If the static pricing model is the correct one, then back-of-the-envelope calculations suggest that mispricing lowers firm profits by perhaps 30%. To the extent that dynamic considerations (such as entry deterrence or longer run behavioral adjustments of consumers) are at work, this crude estimate will be an upper bound on the losses due to mispricing.

In light of the information available to the firm, the patterns observed are not particularly surprising. The firm receives daily feedback regarding each customer's demand for the goods at a given price, and given this information, incorporates it extremely efficiently. The daily activities of the firm, in contrast, yield little information that is useful in determining the optimal price. The firm perceives price changes as costly to make, thus they occur infrequently. Thus, one would expect that the firm would do a better job of choosing quantity conditional on price then identifying the optimal price.

The remainder of the paper is structured as follows. Section 2 describes the data set in greater detail. Section 3 derives the profit maximizing choices in a model that captures key aspects of the firm's decision, which is a two-good version of the "newsboy" problem (Mills, 1959). Section 4 presents the results testing for profit maximization. Section 5 concludes.

2. Background and data

The company began operations in 1984. The nature of the business is straightforward. It purchases bagels, cream cheese, and donuts wholesale, which it delivers to local businesses in the morning. The food is left in a central area, along with a sign that states the prices, and a wooden lock-box in which customers leave their payments on the honor system. Later that day, a company employee returns to collect the money and any uneaten pastries are taken away and discarded.[4] Other than securing approval from the office manager for permission to provide the bagel and donuts, the company has little direct interaction with the customers it supplies. The company does not charge the offices to which it delivers for the service provided. Revenues accrue solely from payments for the bagels and donuts.[5] The typical client firm receives a delivery once per week. In an average week in the sample, the company delivers more than 3000 bagels and 1500 donuts to roughly 125 different clients.

The founder of the company generously provided detailed records of the firm's operations. These data consist of the number of bagels and donuts and amount of cream cheese delivered, the amount of money collected per day, and the number of bagels and donuts that go uneaten. All of this information is reported separately for each client, each day. In addition, the data include both the posted prices that customers are charged and the wholesale costs of the bagels, cream cheese, and donuts.[6] These data cover the period January 1993 to December 2005. As noted earlier, the company began

[2] The analysis I undertake encompasses only a part of the firm's overall profit maximization. I do not examine whether the firm is cost minimizing in its choice of inputs, how it attracts new customers, or its decision concerning which products and services to supply.

[3] Because customers pay on the honor system, revenues are not simply the quantity eaten multiplied by the posted price.

[4] For more details of the daily operations, see Levitt and Dubner (2005).

[5] A small number of companies pay full price in advance for the company's services as a perk for company employees. These companies are excluded from the analysis of this paper.

[6] I also know something about the wages and other operating costs paid by the firm, but these are not used in the analysis.

Table 1
Summary statistics.

Variable	Mean	Standard Deviation	Within-customer standard deviation
Number of bagels delivered	30.032	25.125	14.326
Number of bagels eaten	27.450	24.400	14.003
Number of donuts delivered	13.639	13.490	6.446
Number of donuts eaten	12.814	13.144	6.324
Posted price of bagel (nominal $)	0.828	0.107	0.069
Posted price of donuts (nominal $)	0.506	0.023	0.020
Payment rate	0.894	0.118	0.108
Marginal cost of bagel (nominal $)	0.273	0.060	0.035
Marginal cost of donut (nominal $)	0.219	0.017	0.015
Year	1998.941	3.718	2.154

Notes: The unit of observation is a delivery. Data reflect 80,044 deliveries over the period 1993–2005.
The payment rate is the fraction of the posted price that the company actually receives on the honor system. The marginal cost is the wholesale price of a bagel or donut. The final column presents the standard deviation of the variables within a particular customer over time. See the data appendix for a more detailed description of the data set and its construction.

deliveries almost a decade earlier, but the early years of data are not in a machine readable format.[7] The data appear to be of very high quality, but required some cleaning, the details of which are described in the Data Appendix.

In spite of the richness of the data, there are dimensions along which the information available is limited. First, because the payments are deposited into a single lock box for each customer, it is not possible to observe how much individual customers pay for their goods. Payment rates may, for instance, vary across bagels and donuts, or vary with the number of goods remaining uneaten. Second, the data do not include any information on the variety of bagel (e.g. plain, poppy seed, etc.) or donuts delivered or left uneaten. Consequently, I am unable to investigate questions related to the choice of product variety.

Table 1 presents summary statistics for the data set. The unit of observation is a delivery to a particular office on a specific day.[8] The table presents means, overall standard deviations, and within-office standard deviations over time. Because virtually all of the analysis performed are concerned with decisions at a single point in time, prices used in the paper are in nominal dollars, except where otherwise noted. There are a total of 80,044 deliveries, with an average of approximately 30 bagels and 14 donuts per delivery. More than 90% of the bagels and donuts delivered are actually eaten.

The nominal price customers are charged for bagels changes three times over the course of the sample. Initially, the nominal price for a bagel (which comes with cream cheese) was 60 cents. That price jumped to 75 cents in August 1993, 85 cents in August 1988, and $1.00 in May 2003. The nominal price of donuts is 50 cents over almost the entire sample, increasing to 60 cents in March 2005. In real terms, the price of bagels in 2005 dollars varies between a low of roughly 80 cents at the beginning of the sample before the first price increase and a high of slightly more than a dollar after the last price change. The real price of bagels slowly falls from about 70 cents at the beginning of the sample to roughly 50 cents over the sample period due to inflation, before rising to 60 cents after the price change. The nominal price of donuts is 50 cents over almost the entire sample, increasing to 60 cents in March 2005. In real terms, the price of bagels in 2005 dollars varies between a low of roughly 80 cents at the beginning of the sample before the first price increase and a high of slightly more than a dollar after the last price change. The real price of bagels slowly falls from about 70 cents at the beginning of the sample to roughly 50 cents over the sample period due to inflation, before rising to 60 cents after the price change.

Because payments are made on the honor system, the firm's revenue is less than the posted price. On average, the payment rate (defined as actual revenue divided by expected revenue if the posted prices were paid for each good consumed) is slightly below 90% in the data.

The marginal cost of delivering one more bagel or donut to an existing customer is simply the wholesale cost. The wholesale cost of a bagel (including cream cheese) rises steadily over the period from 20 cents to 37 cents. The wholesale cost of a donut is near 20 cents for most of the sample, but is slightly higher in the late 1990s, peaking at 26 cents in 2000. The wholesale cost per unit is constant over relevant quantity ranges at a given point in time.

Fig. 1 presents monthly time-series data on the total number of bagels and donuts delivered. Donuts represent an increasing share of the product mix over time. At the beginning of the sample, bagels represented more than 80% of the sales in both units and revenues. By the end of the data, the number of bagels and donuts sold are nearly equal, although due to changes in the relative prices charged, bagels continued to account for 70% of revenues. A clear seasonal pattern emerges in the time-series, December deliveries are always low.

[7] The data for 1993–1996 were also available only in hard copy, but were scanned or hand entered.

[8] In cases where deliveries are made to multiple floors of an office building housing employees of the same company, the firm treats these as separate deliveries and I follow that practice as well.

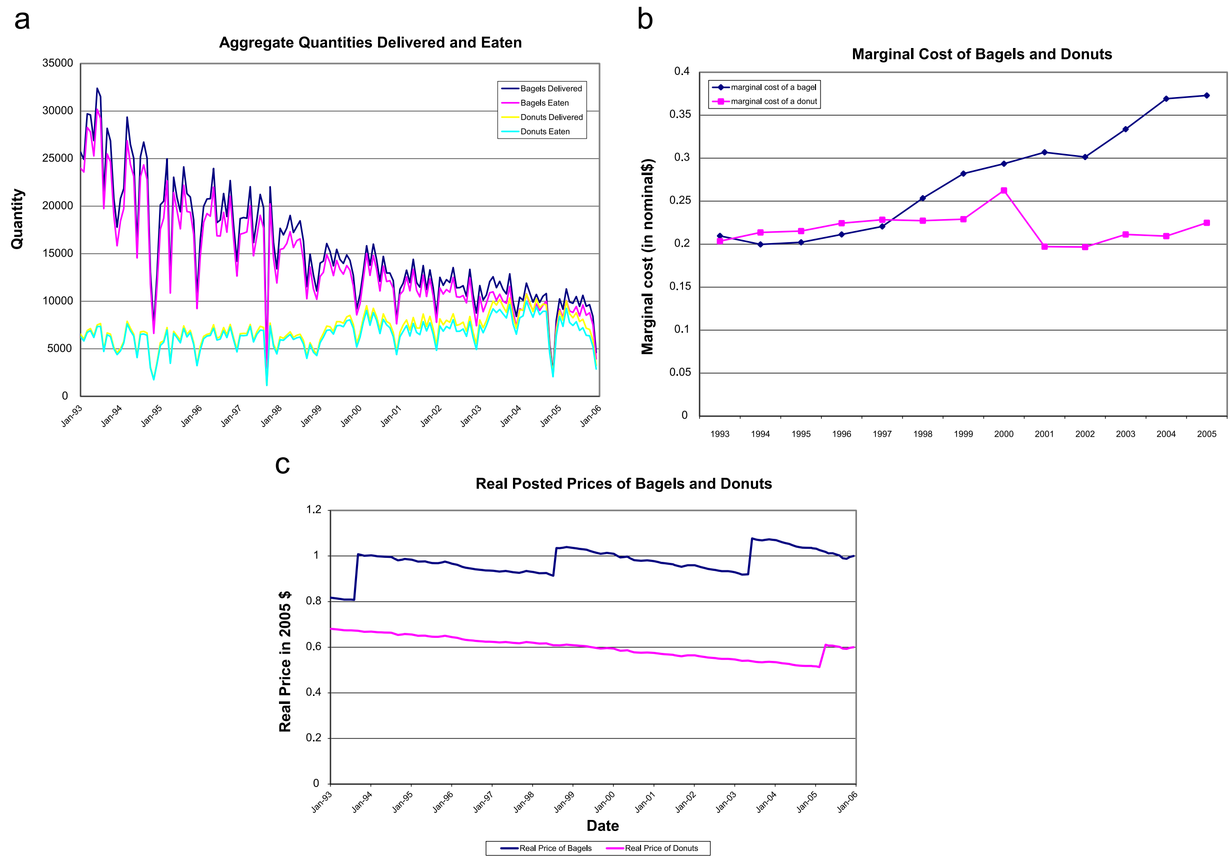

Fig. 1. Aggregate quantities delivered and eaten.

3. Modeling the firm's decision problem

Conditional on the set of products offered and the offices to which the company delivers, the two basic choices the firm faces are the prices to charge and the quantity of bagels and donuts to deliver to a particular office on a given day. In modeling these choices, I will assume that price is taken as fixed in the short run, with the quantity delivered (which is akin to capacity) adjusted on a daily basis in response to anticipated fluctuations in demand. Given that there are only four price changes in the 13 years of my sample, this appears to be an accurate reflection of how the company operates. The firm's decision is an application of the "newsboy" problem, in which a company must commit to choosing an inventory before uncertainty in demand is resolved (Mills, 1959; Carlton, 1978; Dana, 2001; Deneckere and Peck, 1995).

Defining notation, let X_B and X_D represent the quantity of bagels and donuts delivered respectively. Bagels are assumed to be homogeneous, as are donuts. This ignores the fact that consumer preferences may vary across different varieties of bagels (e.g. plain versus pumpernickel) or donuts. P_B and P_D are the posted prices of the two goods to the consumers. Because the goods are sold on an honor system, the price a consumer pays need not equal the posted price. I define the parameter θ as the fraction of the posted price that the company actually receives from a sale taking into account shirking on the honor system. Throughout the analysis, I will assume that the marginal payment rate θ is identical for bagels and donuts, and that the marginal payment rate is equal to the average payment rate (i.e. the marginal consumer is as honest as the average consumer).[9] Consistent with the data, I will also assume that at the posted price and marginal payment rate, the revenue from the sale of the good exceeds the marginal cost (denoted C_B and C_D respectively). The firm's marginal cost of increasing the quantity delivered to an existing customer is constant and equal to the wholesale price of the good delivered. F is the fixed cost of the delivery. The firm maximizes expected profits.

Demand for the goods is characterized as follows. Consider first the case where the firm delivers bagels and donuts in sufficient quantities such that every consumer's demand for each of the products at a given price is satisfied. There will be a demand for each product, denoted N_b and N_d respectively for bagels and donuts. If instead the firm only delivered one of the

[9] The results in Levitt (2006) suggest that the marginal payment rate may be a bit below the average payment rate and that buyers of donuts pay a slightly smaller fraction of the posted price. If this is true, the empirical results that follow will slightly overstate the profitability of the marginal unit of product delivered, particularly for donuts.

products, the implied price of the product that is not delivered is infinite. Thus, the demand for the product that is delivered will be greater when the other product is not available as long as bagels and donuts are substitutes. The quantity of bagels demanded at the fixed price when no donuts are available is denoted as $N_b^{\sim d}$ and the quantity of donuts consumed when no bagels are available is denoted as $N_d^{\sim b}$.

If demand for one of the goods at a particular office on a given day exceeds the quantity delivered of that good, I assume that the good in short supply is rationed randomly among the consumers demanding the good.[10] This assumption implies that the firm is not able to price discriminate between those consumers who will substitute towards the other good in case of a shortfall of their preferred product and those who will not. Thus, the residual demand for bagels is given by

$$\begin{aligned} &N_b+(N_d-X_d)\left(\frac{N_b^{\sim d}-N_b}{N_d}\right) \quad \text{for } N_d \geq X_d \text{ and} \\ &N_b \qquad \text{for } N_d < X_d \end{aligned} \tag{1}$$

A parallel expression holds for donuts. The term inside the first set of parentheses in the first line of Eq. (1) is the excess demand for donuts given prices and the number of donuts delivered. The expression in the second set of parentheses is the number of extra bagels sold when the number of donuts delivered falls by one. (The numerator of that expression is the increase in bagel demand going from the case where no donuts are supplied to the situation where all donut demand is satisfied; the denominator scales it in terms of donuts.)

The degree of cannibalization between bagels and donuts, $\frac{N_b^{\sim d}-N_b}{N_d}$, appears repeatedly throughout the analysis. To simplify notation going forward, I will define $\lambda_{bd} \equiv \frac{N_b^{\sim d}-N_b}{N_d}$, and likewise, let $\lambda_{db} \equiv \frac{N_d^{\sim b}-N_d}{N_b}$ represents the number of extra donuts demanded when one less bagel is delivered.

3.1. The decision problem when demand is known with certainty

I begin with the simplest case in which the firm knows with certainty N_b, N_d,$N_b^{\sim d}$, and $N_d^{\sim b}$ for a particular customer.[11] The firm's optimal decision boils down to a choice of one of three possible strategies: (1) bring both bagels and donuts and exactly meet the demand for each good by supplying N_b and N_d respectively, (2) only bring bagels and deliver $N_b^{\sim d}$ of them, or (3) only bring $N_d^{\sim b}$ donuts. Because the firm is assumed unable to price discriminate between consumers who do substitute and those who do not, there is no middle ground in which the firm elects to deliver a small quantity of one good – the same tradeoff exists between bagels and donuts when going from supplying zero to one donuts or from N_d-1 to N_d donuts. Thus, either it is optimal to fully satisfy the first choices of the consumers, or to bring only one of the goods. Formally, the firm chooses to bring both bagels and donuts unless $\lambda_{bd} > \frac{\theta P_D-C_D}{\theta P_B-C_B}$, in which case bringing only bagels is optimal, or $\lambda_{db} > \frac{\theta P_B-C_B}{\theta P_D-C_D}$, in which case delivering only donuts is optimal. The terms on the left-hand side of these expressions reflect the degree of substitutability between the two goods. The terms on the right-hand side are the ratios of the markups. When one good has a much higher markup and consumption between the two is highly substitutable, delivering none of the low markup good will is optimal. Each unit of the low markup good sold generates less profit than is lost because fewer of the high markup good are sold. Empirically, for the parameters I estimate from the available data, the degree of substitutability between bagels and donuts is well below the threshold that makes it more profitable to deliver only the high markup good (bagels). Delivering both bagels and donuts, as the firm does to most customers, appears to be the more profitable strategy.

3.2. The decision problem when demand is uncertain

When demand is not known with certainty, the choice of optimal quantities to deliver of each good becomes more difficult. Because the goods are substitutes the residual demand for each good is a function not only of demand shocks for that good, but also is indirectly affected by demand shocks for the other good and the quantity delivered of the other good, due to stock outs. Khouja et al. (1996) demonstrate the difficulty of solving this problem explicitly, resorting to Monte Carlo characterizations because an analytical solution proves elusive.

In light of these difficulties, I follow a different path, which is to limit my focus to the first order conditions which hold at the optimum, rather than trying to solve explicitly for quantities. These first order conditions are expressed in terms of the probability that the last bagel or donut delivered is eaten, rather than in terms of absolute numbers of the product eaten and delivered. This approach provides an indirect test of optimizing behavior that is much less demanding of the data and does not require the imposition of arbitrary functional form assumptions.

[10] In practice, consumers may alter their behavior in ways that make this assumption unrealistic. If consumers who preferences run bagels, nothing, donuts gain more from consuming a bagel than those whose preferences are bagels, donuts, nothing (because the latter group may still get a donut even if no bagels remain), then the former set of consumers may be more likely to show up early to ensure they obtain a bagel, for instance.

[11] Because the firm faces constant marginal costs and because demand spillovers across offices are unlikely, the optimal quantity decision for each customer can be made without reference to the circumstances of other customers.

3.3. Uncertain demand with only one product

To begin, take the simple case in which there is only one product, say bagels. At the profit maximizing point, the marginal cost of providing one bagel (C_B) must exactly offset the expected marginal revenue it generates, which is given by $\theta P_B * \Pr(N_b^{\sim d} - X_b \geq 0)$. The expected marginal revenue is the incremental revenue when the bagel is sold, θP_B (remember that the firm does not receive the full posted price, but rather, that price scaled by the average amount of shirking), multiplied by the probability that the last bagel delivered is eaten. Expressed in terms of the firm's choice variable, which is the number of bagels to deliver (X_B), the first order condition for profit maximization is

$$\Pr(\text{Last bagel eaten } | X_B, N_b^{\sim d}) = \frac{C_B}{\theta P_B}. \quad (2)$$

The greater the markup over the marginal cost, the lower is the equilibrium probability that the last unit of the good supplied will be purchased.[12] Note that the earlier assumption that price is fixed is still being maintained.

3.4. Uncertain demand with two products

In the one product case, the consumer faces a choice between buying the product or not. With the introduction of a second good, one must take into account substitution across products. When attempting to derive an explicit solution for the optimal quantities, the addition of a second good dramatically increases the complexity of the problem, as demonstrated in Khouja et al. (1996). Analyzed in terms of first-order conditions, however, introducing a second good is relatively straightforward. Take the case of determining the first-order condition for bagels, holding fixed the number of donuts delivered. There are four different regimes to consider, each corresponding to whether, at the current choice of delivery quantities, there is excess supply or excess demand for each of the goods.[13] If there is excess demand for both bagels and donuts (i.e. both stock out), then cross-good substitution is not a concern when adding an additional bagel. The extra bagel may lead a consumer who prefers bagels over donuts to switch her purchase from a donut to a bagel, but that donut will be taken by another customer because there is unsatisfied demand for donuts. Similarly, if there is excess demand for donuts and an excess supply of bagels, bringing an extra bagel will not affect the consumption of donuts. There are already excess bagels, so bringing one more bagel should have no impact on anyone's consumption; the last bagel will simply go uneaten with certainty and nothing else will change. Similarly, if there is excess supply of both bagels *and* donuts, bringing an extra bagel will have no impact on either bagel or donut consumption. The extra bagel will once again go uneaten. Thus, in three of the four regimes, substitution away from donuts towards bagels is not an issue. In those regimes, the same first order condition that was present in the one-good case continues to hold in the two-good case.

The only case in which the marginal bagel influences the quantity of donuts consumed is when there is excess supply of donuts and excess demand for bagels. In that case, if the consumer who chooses the last bagel switches away from a donut to purchase that bagel, the number of uneaten donuts will increase by one. Using the notation and assumptions introduced above, the probability that the consumer who consumes the last bagel switches off of a donut is equal to λ_{db}. When such a switch occurs, the firm earns the revenue associated with selling the marginal bagel, but sacrifices the revenue from the lost donut sale. The only difference between the first-order condition in the two-good case and the one-good case is that this substitution across goods leads to an additional term in the two-good case, but only when the firm is in the fourth regime where there is excess demand for bagels and excess supply of donuts:

$$\Pr(\text{last bagel eaten}) = \frac{C_B}{\theta P_B} + \lambda_{db}\frac{P_D}{P_B} * \Pr(\text{last bagel eaten, last donut not eaten}) \quad (3)$$

where λ_{db} is the degree of substitution to donuts from bagels. The term after the addition sign on the right-hand side of the equation is the foregone revenue from lost donut sales due to bringing one extra bagel.[14]

[12] See, Hadley and Whitin (1963) for a more general characterization of these first order conditions allowing for other considerations such as a non-zero salvage value of unsold products or good will losses associated with stock-outs. Note that (except in extreme cases such as positive feedback effects) the conditional probability of the last bagel being eaten is (weakly) declining with the number of bagels delivered; if there are two delivery quantities that satisfy Eq. (2), the greater of these maximizes profits.

Another point worth noting is that I have assumed that all bagels are identical, whereas in practice there is heterogeneity (e.g., poppy vs. plain bagel). Even if a poppy bagel goes uneaten, if the firm had brought one extra plain bagel, that plain bagel may have been consumed. Product heterogeneity will thus bias me towards concluding that the firm is delivering too many bagels and donuts when indeed the firm might be choosing optimally (and similarly, to possibly argue that the firm is choosing quantity optimally, when in fact they are delivering too few).

Ideally, we would like to observe not just the presence of a stock out, but also the time at which the stock-out occurred. Unfortunately, the data do not contain such information.

[13] In the discussion that follows, I ignore the fact that donuts and bagels come in discrete units, as opposed to varying continuously. Near the margins of the four regimes, of course, the discreteness can push the firm from one regime to another, but these changes are of second order importance in solving the problem.

[14] I limit my derivation of Eq. (3) to an informal discussion in the text. Readers interested in a more formal proof are directed to Eq. (1) of Khouja et al. (1996), which presents the profit function for a two-item newsboy problem with substitution. Assigning a salvage value of zero for leftover bagels and donuts, taking the first-order condition of Eq. (1) (Khouja et al., 1996), setting that first-order condition equal to zero, and rearranging terms, yields my Eq. (3).

Table 2
Estimated profitability of the last bagel delivered when bagels are the only product delivered.

Year	(1) Average posted price of bagel	(2) Marginal cost of bagel	(3) Payment rate	(4) Probability all but one bagel eaten	(5) Probability all bagels eaten	(6) Expected profit from *next to last* bagel delivered	(7) Expected profit from *last* bagel delivered	(8) Expected profit if *one extra* bagel had been delivered
1993	0.646	0.210	0.926	0.489	0.365	0.083 (0.006)	0.009 (0.006)	−0.047 (0.006)
1994	0.750	0.200	0.914	0.438	0.315	0.101 (0.007)	0.016 (0.006)	−0.044 (0.007)
1995	0.750	0.202	0.909	0.466	0.347	0.116 (0.007)	0.035 (0.007)	−0.026 (0.007)
1996	0.750	0.211	0.908	0.448	0.307	0.094 (0.007)	−0.002 (0.006)	−0.068 (0.007)
1997	0.750	0.221	0.901	0.438	0.320	0.075 (0.008)	−0.005 (0.007)	−0.063 (0.008)
1998	0.785	0.253	0.896	0.448	0.339	0.062 (0.008)	−0.015 (0.007)	−0.073 (0.008)
1999	0.850	0.282	0.904	0.578	0.436	0.162 (0.009)	0.053 (0.009)	−0.029 (0.011)
2000	0.850	0.294	0.919	0.625	0.458	0.195 (0.010)	0.064 (0.011)	−0.032 (0.012)
2001	0.850	0.307	0.901	0.633	0.475	0.178 (0.012)	0.057 (0.013)	−0.033 (0.015)
2002	0.850	0.301	0.903	0.583	0.421	0.146 (0.013)	0.022 (0.013)	−0.068 (0.014)
2003	0.950	0.334	0.893	0.494	0.324	0.085 (0.015)	−0.059 (0.013)	−0.154 (0.014)
2004	1.000	0.369	0.883	0.558	0.379	0.123 (0.021)	−0.034 (0.020)	−0.141 (0.023)
2005	1.000	0.373	0.863	0.578	0.375	0.126 (0.017)	−0.049 (0.017)	−0.163 (0.018)
Average over all years	0.784	0.246	0.906	0.497	0.361	0.107 (0.002)	0.011 (0.002)	−0.059 (0.003)

Note: Each row corresponds to average values for all deliveries in the listed calendar year for customers receiving bagels only.
Columns (1) through (5) are observed in the data. Columns (6) through (8) are estimates using Eq. (1) of the paper to compute the expected profit from the next to last bagel delivered, the last bagel delivered, and if one extra bagel was delivered.
Standard errors (in parentheses below the estimates) are calculated using a nonparametric bootstrap with one thousand replications.
Column (8) is calculated under the assumption that the likelihood that the $N+1$ bagel is eaten conditional on the Nth being eaten is equal to the probability the Nth was eaten conditional on the $N-1$ bagel being eaten.

Comparing a situation in which only bagels are delivered (Eq. (2)) to one in which both products are offered (Eq. (3)), the equilibrium probability that the last bagel will be eaten is (weakly) higher when both goods are delivered. This implies a *reduction* in the number of bagels delivered when both goods are offered. This is true not only because demand for bagels falls when donuts are offered, but also because increases in bagel sales lead to foregone revenue associated with cannibalizing donut sales on the margin. The only cases in which the probability the last bagel eaten is the same across the two situations is when there is no substitution between the products, or when there is never simultaneously a shortage of bagels and an excess supply of donuts.[15]

The corresponding first order condition for donuts is identical to that of bagels, except that the roles of bagels and donuts have been reversed:

$$\Pr(\text{last donut eaten}) = \frac{C_D}{\theta P_D} + \lambda_{bd}\frac{P_B}{P_D} * \Pr(\text{last donut eaten, last bagel not eaten}). \tag{4}$$

Empirically, bagels are the more expensive good and carry a larger markup. Comparing Eqs. (3) and (4), the term before the addition sign pushes towards bagels stocking out less frequently than donuts. In the term after the addition sign, both the higher price of bagels, and the probability term reinforce the tendency for donuts to stock out. Therefore, for reasonable values of the rate of substitution from bagels to donuts and vice-versa, therefore, the likelihood that donuts stock out will exceed that of bagels stocking out.

[15] In the case of perfectly correlated demand shocks for bagels and donuts, for instance, it can be shown that donuts will always stock out before bagels (which have a higher markup).

3.5. Setting Optimal Prices

The firm's price decision is standard: it wants to price such that marginal cost equals marginal revenue, following the usual inverse elasticity rule of choosing markups. To the extent that the two goods are substitutes or complements in demand, the pricing decision should also take into account spillovers across the two goods.

4. Testing for profit maximization in the choice of quantities

Detailed data, combined with the straightforward nature of this company's business, provide an unusually direct opportunity to test for profit maximizing behavior on the part of the firm along two important dimensions of decision making: the choice of quantity and the choice of price. In this section I explore the choice of quantity to deliver.

4.1. The quantity decision when only one product is delivered

A subset of the customers to which the company delivers receive no donuts, only bagels. When only a single product is delivered, Eq. (1) captures the first-order conditions governing the optimal quantity to deliver conditional on a particular price. The determinants of the optimal quantity choice are the posted price of a bagel, the degree of underpayment on the honor system, the marginal cost, and the observed probability that all the bagels delivered are eaten. All of these factors are either known, or can be readily computed from the available data.

Table 2 analyzes the degree to which the firm's actions, on average over the course of a year, correspond to the prediction in Eq. (1).[16] Each row of the table corresponds to a different year's data. The first five columns report the means of the five factors that enter into the first-order condition. Based on those data, columns 6–8 present the estimated profit generated from the penultimate bagel delivered, the last bagel delivered, and the predicted profit had one extra bagel been delivered.[17] Standard errors, computed using a nonparametric bootstrap, are shown in parentheses. If the firm is optimizing with respect to quantity delivered, the profit on the last bagel should be close to zero. By extension of that same logic, there should be positive profit associated with the next-to-last bagel, and losses associated with bringing one extra bagel. The results reported in Table 2 demonstrate that the firm's quantity choices correspond closely with the predictions of the model. Averaged over the whole sample, the expected profit on the last bagel delivered is 1.1 cents.[18] This estimate has small standard errors, as do the others in the table. Given marginal costs, prices, and the payment rate, the probability that all bagels are eaten that sets the profit of the last bagel to zero is .345. The actual value observed in the data is .36. Across years, the estimated profit on the last bagel delivered ranges from 5.9 cents to 6.4 cents.

Comparing the expected profit on the last bagel delivered to that of the penultimate and $N+1$ bagel provides a further measure of the firm's choice of quantity.

Columns 6 and 8 of Table 2 present estimates of the expected profit associated with the $N-1$ bagel and the $N+1$ bagel. In column 6, the next-to-last bagel yields positive profits in all years, averaging between 10 and 11 cents over the entire sample. Thus, although there are a few years (especially in recent times) in which the last bagel delivered has a negative impact on profits, there are no years in which the company would want to reduce average delivery size by more than one bagel. Because there are relatively few customers who receive only bagels and the loss on the last bagel is small, these deviations from the optimal quantity have a minimal impact on the bottom line. Had the company lowered bagel deliveries by one to all clients in the years of the sample in which the last bagel delivered returned negative profits, the total increase in profit would be less than $200 combined.

The hypothetical profit associated with delivering one extra bagel, shown in column 8, is consistently negative, producing a loss of roughly 6 cents on average. There are no years in which the firm could have increased profits through an

[16] In this first initial analysis of the data I aggregate across customers because all of the components going into the test of first order conditions are constant across firms at a given point in time. Later in the paper I analyze whether including observable characteristics of firms and their consumption histories can improve the quantity decisions of the bagel company.

[17] For a hypothetical increase in the number of bagels delivered from N to $N+1$, additional assumptions are necessary since in the data I only observe whether N bagels were eaten, but what is needed for this estimate is the probability that the $N+1$ bagel will be eaten. One way to approximate that probability is to compute the likelihood that the Nth bagel will be eaten, conditional on bagel $N-1$ being eaten, and to multiply that number by the probability that the Nth bagel is eaten. This value is observed in the data, and among customers receiving only bagels, the likelihood that the Nth bagel is consumed given that the $N-1$ bagel is eaten is roughly 73% over the course of the sample, with a high of 75.8% in 1998 and a low of 65.6% in 2003. Those numbers are likely to be upper bounds on the likelihood that the $N+1$ donut is eaten given that donut N is eaten. By using this overly optimistic number in my hypothetical calculations, I overstate the true profitability of increasing bagel quantities, making it easier to reject that the firm actually took the correct action on average in a year. It is also possible, however, that the likelihood the $N+1$ donut is eaten with a higher probability than is given by my assumption. For instance, if at a particular office demand comes in discrete clumps (i.e. either 10 donuts eaten or 20, but never in between), then the probability that the 11th donut gets eaten is identical to the probability that the 12th donut gets eaten, not declining as I assume. It is also possible (but probably unlikely), that the probability that the $N+1$ donut is eaten conditional on N donuts being eaten rises with the number of donuts delivered if there are positive feedback effects (e.g. each extra person standing around the table eating a donut attracts more than one additional consumers to buy a donut).

[18] The average revenue from a bagel in the sample, taking into account a payment rate less than 1, is 64.1 cents. The marginal cost of a bagel, averaged over the sample, is 24.4 cents. So if the last bagel delivered were eaten with certainty, the marginal profit on that bagel would be 39.7 cents. If the last bagel delivered was never eaten the profit would be 24.4 cents.

Table 3
Estimates of the cannibalization rate across products.

	(1)	(2)	(3)	(4)	(5)	(6)
	Number of bagels delivered is unchanged and...			Number of donuts delivered is unchanged and...		
	Donuts delivered unchanged	Donuts delivered increases	Donuts delivered decreases	Bagels delivered unchanged	Bagels delivered increases	Bagels delivered decreases
Change in:						
Donuts delivered	0.000	4.369	−4.250	0.000	0.000	0.000
Bagels delivered	0.000	0.000	0.000	0.000	2.922	−3.315
Donuts eaten	−0.088	3.272	−1.997	−0.088	−0.328	0.060
Bagels eaten	−0.169	−0.678	0.150	−0.169	0.865	−0.522
Implied cannibalization rate: (diffs-in-diffs using column 1 as control)						
From delivering one more donut	–	0.151	0.167	–	–	–
From delivering one more bagel	–	–	–	–	0.232	0.419

Notes: Estimates in the table are mean changes in donuts and bagels delivered and eaten from one delivery to a client to the next when the quantity delivered of at least one of the products is unchanged across the two deliveries.
In columns (1) and (4), both quantities delivered remain constant. In columns (2) and (3) only the number of donuts delivered changes. In columns (5) and (6) only the number of bagels delivered changes. The bottom two rows of the table present the implied cannibalization rates, using the results in column (1) as the baseline.

across the board increase of one bagel per delivery. The firm appears to do an extremely good job of choosing the quantity of bagels to deliver when that is the only product delivered.

4.2. The quantity decision when both bagels and donuts are delivered

The optimal quantity decision becomes more difficult when two products are delivered rather than one. Now, in addition to the other factors that enter the first-order conditions in the one-good case, the probability that each of the goods stocks out, and the degree of cannibalization across the two goods become relevant. The stock-out rates are readily observed in the data. The degree of cannibalization across the two goods is not directly observable and must be estimated.

The ideal setting for estimating the substitution from donuts to bagels would be a randomized experiment in which the researcher systematically varies the number of donuts delivered, and then measures the change in consumption of both donuts and bagels. Likewise, one would want to induce random variation in the number of bagels delivered to identify the extent of cannibalization of donuts by bagels.

Unfortunately, this sort of randomization is not present in the data. The observed fluctuations in quantities delivered are clearly not exogenous. The company increases the number of bagels delivered when anticipated demand for bagels is high. To the extent that demand shocks for the two goods are positively correlated, demand for donuts will generally be higher when there is a large delivery of bagels than when the delivery is small. In practice, estimates of the cannibalization rates for these two goods are biased toward zero due to observed variation in the data.

As a crude proxy for the degree of cannibalization, I compare changes in the number of bagels and donuts consumed from one delivery to the next at a client for those cases where the quantity delivered of one good remains constant, but the quantity of the other good does not.[19] I measure cannibalization of bagels by donuts as follows:

$$\lambda_{bd} = \frac{\Delta(Bagels\_eaten)}{\Delta(Donuts\_eaten)}\Big|\ \Delta X_b = 0,\ \Delta X_d \neq 0$$

where Δ reflects the change from one week to the next at a given customer. If, in response to six extra donuts being delivered, four more donuts and one fewer bagels are eaten, I would estimate a cannibalization effect on bagels by donuts of −.25. The formula for cannibalization of donuts by bagels is identical, except with the values for bagels and donuts reversed in the equation.

Focusing on cases where the quantity delivered of one good remains constant has two benefits. First, in the case of stock outs of that good, the truncation point remains the same in both weeks, facilitating comparisons of the quantity demanded.

[19] One institutional factor that aids in this estimation is that the firm almost always delivers donuts in multiples of six. Because of the lumpiness of the delivery quantity, small changes in perceived demand may lead to discrete jumps in donuts delivered, as implied by the standard sS model, somewhat mitigating the issues raised by endogeneity of the delivery size. For bagels, however, this lumpiness is not present.

Table 4
Estimated profitability of the last bagel eaten (when both products delivered).

Year	(1) Average posted price of bagel	(2) Marginal cost of bagel	(3) Payment rate	(4) Probability all but one bagel eaten	(5) Probability all bagels eaten	(6) Probability all bagels eaten and excess of donuts	(7) Expected profit from *next to last* bagel delivered	(8) Expected profit from *last* bagel delivered	(9) Expected profit if *one extra* bagel had been delivered
1993	0.647	0.210	0.908	0.456	0.365	0.052	0.047 (0.004)	−0.003 (0.004)	−0.043 (0.005)
1994	0.750	0.200	0.894	0.426	0.328	0.052	0.074 (0.005)	0.012 (0.005)	−0.036 (0.005)
1995	0.750	0.202	0.888	0.466	0.354	0.049	0.097 (0.005)	0.027 (0.005)	−0.027 (0.005)
1996	0.750	0.211	0.884	0.472	0.354	0.056	0.088 (0.005)	0.016 (0.004)	−0.040 (0.005)
1997	0.750	0.221	0.882	0.463	0.358	0.058	0.074 (0.005)	0.008 (0.004)	−0.043 (0.005)
1998	0.787	0.253	0.874	0.469	0.360	0.069	0.055 (0.005)	−0.015 (0.004)	−0.070 (0.005)
1999	0.850	0.282	0.883	0.563	0.430	0.114	0.117 (0.005)	0.024 (0.005)	−0.047 (0.006)
2000	0.850	0.294	0.882	0.634	0.487	0.127	0.154 (0.005)	0.053 (0.005)	−0.025 (0.006)
2001	0.850	0.307	0.880	0.605	0.447	0.151	0.112 (0.005)	0.006 (0.005)	−0.073 (0.006)
2002	0.850	0.301	0.900	0.560	0.411	0.132	0.097 (0.005)	−0.006 (0.005)	−0.083 (0.006)
2003	0.954	0.334	0.894	0.522	0.373	0.109	0.086 (0.005)	−0.031 (0.005)	−0.116 (0.006)
2004	1.000	0.369	0.893	0.585	0.412	0.131	0.122 (0.006)	−0.020 (0.006)	−0.122 (0.006)
2005	1.000	0.373	0.902	0.575	0.405	0.160	0.103 (0.005)	−0.035 (0.005)	−0.134 (0.006)
Average over all years	0.843	0.281	0.890	0.530	0.395	0.101	0.093 (0.001)	0.000 (0.001)	−0.070 (0.002)

Note: Each row corresponds to average values for all deliveries in the listed calendar year for customers receiving both bagels and donuts. Columns (1) through (6) are observed in the data. Columns (7) through (9) are estimates using Eq. (2) of the paper to compute the expected profit from the next to last bagel delivered, the last bagel delivered, and if one extra bagel were delivered. These estimates use the cannibalization estimates obtained in Table 3. Standard errors (in parentheses below the estimates) are calculated using a nonparametric bootstrap with one thousand replications. Column (9) is calculated under the assumption that the likelihood that the $N+1$ bagel is eaten conditional on the Nth being eaten is equal to the probability the Nth was eaten conditional on the $N-1$ bagel being eaten.

Second, the fact that the quantity delivered of one good did not change suggests that the firm did not anticipate a large fluctuation in demand relative to the previous week (although if they deliver more of the other good it does suggest that they are relatively optimistic).

The results on cannibalization across products are presented in Table 3. The first column shows the changes in bagels and donuts consumed in cases in which there is no change from the prior week in the number of bagels and donuts delivered: the number of bagels eaten falls by .169 and the number of donuts eaten is reduced by .088. These changes will serve as the counterfactual as to what would be expected had the quantities delivered not changed. Columns 2 and 3 present results when the number of donuts delivered changes and bagels delivered remains constant. Column 2 reflects cases where donuts delivered increase; column 3 has decreases in delivered donuts. In those instances when donut deliveries rise, the increase in the number of donuts delivered is 4.37, with 3.27 of those extra donuts eaten. Bagel consumption falls by .68. Relative to the counterfactual in column 1, an extra 3.36 donuts are eaten, and bagels consumed fall by .509. This implies that each extra donut eaten cannibalizes .151 bagel sales. In column 3, the number of donuts delivered declines, leading to approximately two fewer donuts eaten and an extra .15 bagels consumed. Compared to the counterfactual in column 1, donuts consumed fall by 1.91 and bagels eaten rise by .319, for an implied cannibalization rate of .167, nearly identical to the parallel figure for Column 2.

Columns 5 and 6 show the results when the number of bagels varies, with donut quantities held constant. When more bagels are delivered, an extra .865 bagels are eaten, and .328 fewer donuts. Relative to the column 4 counterfactual, bagels consumed rise by 1.034 and a .24 reduction in donuts, implying a cannibalization rate of .232. When fewer bagels are delivered, bagels eaten fall by .522, with an increase of .060 donuts. Compared to the counterfactual, this is a decline of .353 bagels and a .148 rise in donuts, for a cannibalization rate of .419 – substantially larger than implied by the results in column 5.

Table 5
Estimated profitability of the last donut eaten (when both products delivered).

Year	(1) Average posted price of donut	(2) Marginal cost of donut	(3) Payment rate	(4) Probability all but one donut eaten	(5) Probability all donuts eaten	(6) Probability all donuts eaten and excess of bagels	(7) Expected profit from *next to last* donut delivered	(8) Expected profit from *last* donut delivered	(9) Expected profit if *one extra* donut had been delivered
1993	0.500	0.204	0.908	0.853	0.723	0.410	0.136 (0.002)	0.086 (0.003)	0.044 (0.004)
1994	0.500	0.214	0.894	0.840	0.700	0.425	0.105 (0.002)	0.054 (0.003)	0.011 (0.004)
1995	0.500	0.215	0.888	0.841	0.695	0.389	0.105 (0.002)	0.052 (0.003)	0.008 (0.004)
1996	0.500	0.224	0.884	0.822	0.670	0.372	0.087 (0.002)	0.033 (0.003)	−0.013 (0.004)
1997	0.500	0.228	0.882	0.853	0.703	0.402	0.093 (0.002)	0.039 (0.003)	−0.006 (0.004)
1998	0.500	0.227	0.874	0.832	0.670	0.379	0.082 (0.002)	0.024 (0.003)	−0.023 (0.003)
1999	0.500	0.229	0.883	0.777	0.615	0.299	0.066 (0.002)	0.007 (0.003)	−0.041 (0.003)
2000	0.500	0.263	0.882	0.756	0.605	0.245	0.030 (0.002)	−0.025 (0.003)	−0.070 (0.003)
2001	0.500	0.197	0.880	0.653	0.512	0.216	0.054 (0.003)	0.003 (0.003)	−0.039 (0.003)
2002	0.500	0.197	0.900	0.660	0.520	0.241	0.059 (0.003)	0.008 (0.003)	−0.033 (0.003)
2003	0.500	0.211	0.894	0.676	0.545	0.281	0.040 (0.002)	−0.006 (0.002)	−0.043 (0.003)
2004	0.500	0.209	0.893	0.692	0.553	0.272	0.048 (0.002)	−0.001 (0.003)	−0.041 (0.003)
2005	0.575	0.225	0.902	0.630	0.471	0.227	0.055 (0.003)	−0.013 (0.003)	−0.065 (0.003)
Average over all years	0.507	0.219	0.890	0.751	0.605	0.311	0.070 (0.001)	0.017 (0.001)	−0.027 (0.001)

Note: Each row corresponds to average values for all deliveries in the listed calendar year for customers receiving both bagels and donuts. Columns (1) through (6) are observed in the data. Columns (7) through (9) are estimates using Eq. (3) of the paper to compute the expected profit from the next to last donut delivered, the last donut delivered, and if one extra donut were delivered. These estimates use the cannibalization estimates obtained in Table 3. Standard errors (in parentheses below the estimates) are calculated using a nonparametric bootstrap with one thousand replications. Column (9) is calculated under the assumption that the likelihood that the $N+1$ donut is eaten conditional on the Nth being eaten is equal to the probability the Nth was eaten conditional on the $N-1$ donut being eaten.

Based on these results, I proceed by using a simple average of the two estimates obtained for cannibalization in each direction. Thus, I assume that each extra donut eaten lowers bagels eaten by .159, and that the cannibalization rate in the other direction is .325.[20]

Table 4 reports the results on the estimated profitability of the last bagel delivered for customers who receive both bagels and donuts using Eq. (2). For purposes of comparison, I also report the implied profitability of the $N-1$ and $N+1$ bagel delivered, based on the same assumptions used in Table 2. As was true for clients who received only bagels, the firm does an outstanding job of choosing the quantity of bagels for customers who receive both products. On average across the entire sample, the firm earns a zero profit from the last bagel delivered. In six of the thirteen years, including the last four years, the last bagel delivered had a negative profit expectation, although in even the worst year the loss was less than four cents. There are no years in which the company could have improved profits by lowering bagel deliveries by more than one per customer (i.e. the profits of the next to last bagel is always positive) or by increasing bagel deliveries across the board (i.e. the expected profit of the $N+1$ bagel is always negative).

Because the markup on bagels is greater than that of donuts, bagel stock-outs occur less frequently than donut stock-outs. As a consequence, it is relatively rare that all bagels are consumed, but donuts remain uneaten; this occurs between 5 and 16% of the time depending on the year. The infrequency of this circumstance coupled with the relatively low profit margin on donuts means that cannibalization considerations are of minor importance in the choice of bagel quantity.

[20] The results do not prove to be particularly sensitive to these cannibalization parameters. Assuming no cannibalization, or cannibalization that is twice as large as assumed here shifts the marginal return on the last bagel delivered up by an average of 1.5 cents (if no cannibalization) or down by that same amount (if twice the amount of cannibalization). For donuts, the effect is larger: 3.7 cents up or down for the marginal donut.

Table 6
Predicting stock-outs from observable characteristics.

Can regression improve on the company's choices?

Bagels	Distribution of uneaten bagels, by regression prediction of degree of shortfall or excess of bagels delivered			
Observed number of bagels left	(1) 10% of observations with greatest predicted bagel *shortfall*	(2) 50% of observations with greatest predicted bagel *shortfall*	(3) 50% of observations with greatest predicted bagel *surplus*	(4) 10% of observations with greatest predicted bagel *surplus*
Zero	0.520	0.460	0.330	0.261
One	0.125	0.139	0.143	0.126
Two	0.096	0.108	0.124	0.115
Three	0.071	0.085	0.102	0.104
Four	0.049	0.060	0.081	0.083
Five or more	0.139	0.148	0.220	0.311
Expected profit on:				
next to last bagel	0.185	0.151	0.059	−0.002
last bagel	0.099	0.055	−0.039	−0.089
extra bagel	0.028	−0.020	−0.108	−0.148

Donuts	Distribution of uneaten donuts, by regression prediction of degree of shortfall or excess of donuts delivered			
Observed number of donuts left	10% of observations with greatest predicted donut *shortfall*	50% of observations with greatest predicted donut *shortfall*	50% of observations with greatest predicted donut *surplus*	10% of observations with greatest predicted donut *surplus*
Zero	0.733	0.685	0.496	0.382
One	0.115	0.133	0.165	0.156
Two	0.062	0.073	0.115	0.125
Three	0.039	0.044	0.077	0.094
Four	0.019	0.026	0.051	0.071
Five or more	0.032	0.041	0.096	0.172
Expected profit on:				
next to last donut	0.116	0.100	0.036	−0.013
last donut	0.073	0.051	−0.025	−0.070
extra donut	0.036	0.010	−0.072	−0.112

Predicting stock-outs from observable characteristics **(INCLUDES COMPANY FIXED EFFECTS).**
Can regression improve on the company's choices?

Bagels	Distribution of uneaten bagels, by regression prediction of degree of shortfall or excess of bagels delivered			
Observed number of bagels left	(1) 10% of observations with greatest predicted bagel *shortfall*	(2) 50% of observations with greatest predicted bagel *shortfall*	(3) 50% of observations with greatest predicted bagel *surplus*	(4) 10% of observations with greatest predicted bagel *surplus*
Zero	0.460	0.441	0.347	0.309
One	0.134	0.140	0.142	0.137
Two	0.102	0.107	0.125	0.122
Three	0.077	0.085	0.103	0.103
Four	0.060	0.064	0.078	0.079
Five or more	0.168	0.163	0.206	0.251
Expected profit on:				
next to last bagel	0.144	0.137	0.072	0.040
last bagel	0.053	0.040	−0.026	−0.054
extra bagel	−0.019	−0.034	−0.096	−0.120

Donuts	Distribution of uneaten donuts, by regression prediction of degree of shortfall or excess of donuts delivered			
Observed number of donuts left	10% of observations with greatest predicted donut *shortfall*	50% of observations with greatest predicted donut *shortfall*	50% of observations with greatest predicted donut *surplus*	10% of observations with greatest predicted donut *surplus*
Zero	0.709	0.664	0.514	0.453

Table 6 (*continued*)

Predicting stock-outs from observable characteristics **(INCLUDES COMPANY FIXED EFFECTS).**
Can regression improve on the company's choices?

Bagels	Distribution of uneaten bagels, by regression prediction of degree of shortfall or excess of bagels delivered			
	(1)	(2)	(3)	(4)
One	0.129	0.140	0.158	0.143
Two	0.063	0.077	0.111	0.117
Three	0.039	0.047	0.074	0.075
Four	0.023	0.028	0.049	0.063
Five or more	0.037	0.044	0.093	0.150
Expected profit on:				
next to last donut	0.108	0.094	0.041	0.009
last donut	0.060	0.042	−0.017	−0.043
extra donut	0.020	−0.001	−0.062	−0.083

Notes: The results in the table report the number of bagels (top panel) and donuts (bottom panel) that go uneaten depending on the predictions of a probit model designed to improve on the quantity choices of the company. The probit model uses only information observable to the company ex-ante and is estimated on the prior two calendar years of data and extrapolated out of sample.
Deliveries are rank ordered within a year based on the predicted likelihood of a stockout. Column (1) reports actual bagels uneaten in the 10% of deliveries predicted to be most likely to stock out.
Columns (2) and (3) are outcomes for deliveries above and below the median predicted stock out.
Column (4) corresponds to the 10% of deliveries predicted to be least likely to stock out.
For each type of good in each column, the estimated profitability of the next to last, last, and one extra donut are reported.

Referring back to Eq. (2), the presence of donuts reduces the optimal probability that the last bagel should be eaten by less than one percentage point on average.

The firm exhibits substantial skill in altering its behavior over time to achieve the profit maximizing delivery quantities. For instance, comparing 1996 and 2001, in both years the firm earned a small profit on the last bagel delivered. In 1996, however, the probability the last bagel is eaten is 35.4%, versus almost 45% in 2001. Because the ratio of price to marginal cost is higher in 1996, the optimal probability that the last bagel eaten is lower, partially offset by the fact that in 2001 it is more frequently the case that there uneaten donuts available to absorb unsatisfied demand when bagels stock out.

Table 5 is identical to Table 4, except that the values reported represent the profit associated with the last donut delivered instead of the last bagel. The firm does slightly worse in optimizing donut quantities than it does for bagels early in the sample, but otherwise chooses the number of donuts extremely well. In the first three years of the sample, the firm delivers too few donuts, i.e. the expected profit from adding one donut to each delivery is positive. There are four years (including the last three years) in which the last bagel delivered yields negative returns, implying that the firm would increase profit by reducing the quantity delivered. The foregone profits of these slight deviations from the optimal choice of quantity, however, are trivial: less than $300 total across the whole sample. It is interesting to note that at the end of the sample, the quantity of both bagels and donuts deliver appear to be too high.

Cannibalization is a much more important concern for donuts than bagels for two reasons. First, it is much more likely that there are excess bagels when all donuts are eaten then vice versa. Second, there is a larger markup on bagels than on donuts, so a customer who substitutes a donut for a bagel is costly to the firm. Substituting the values observed in the data, cannibalization concerns raise the optimal probability of having the last donut eaten by 3 percentage points on average in the sample off of a baseline of a little more than 60%.

The estimates in Tables 4 and 5 are averages over all deliveries in a year. It is possible that the firm does well on average, but systematically delivers too much product to some clients and too little to others. To investigate this possibility, I estimate probit regressions in which the dependent variable is equal to one if all bagels are eaten on a particular delivery and zero otherwise. Explanatory variables include the number of bagels delivered this time as well as in the last two deliveries to this customer, the number of uneaten bagels in each of the last two deliveries, indicator variables for stock-outs of bagels in the last two deliveries. Also included are the parallel variables for donut deliveries, as well as year, month, and day of the week indicators.[21] All of the variables in the regression are observable to the company at the time the quantity decision is made. Using data from the preceding two calendar years, I generate out-of-sample predictions regarding the likelihood of a stock-out of bagels on each delivery.[22] By basing my predictions only on the outcome of prior years of data, I ensure that only information available to the firm when making their quantity choice underlies my estimates. I then rank–order observations within a year with respect to the predicted likelihood of a stock-out. I also carry out a parallel exercise for donut deliveries including the exact same set of explanatory variables.

[21] I have experimented with using continuous measures of time as opposed to month dummies, as well as including customer-level fixed effects in the regressions with little impact on my estimates.

[22] Because I require two years of earlier data to carry out this exercise, I do not make predictions for the first two years of the sample.

Table 7
Changes in profits and revenue in response to price changes.

	(1) August 31, 1993	(2)	(3)	(4) August 4, 1998	(5)	(6)	(7) May 5, 2003	(8)	(9)	(10) March 28, 2005	(11)	(12)
Variable	Before	After	*P*-value	Before	After	*P*-value	Before	After	*P*-value	Before	After	*P*-value
Price of bagel	0.600	0.750		0.750	0.850		0.850	1.000		1.000	1.000	
Price of donut	0.500	0.500		0.500	0.500		0.500	0.500		0.500	0.600	
Marginal cost of bagel	0.210	0.210		0.253	0.253		0.334	0.334		0.373	0.373	
Marginal cost of donut	0.204	0.204		0.227	0.227		0.211	0.211		0.225	0.225	
Bagels delivered	52.084	49.077	0.000	30.170	28.693	0.000	22.810	22.101	0.000	18.610	18.247	0.059
Bagels eaten	48.884	44.516	0.000	27.091	25.120	0.000	20.540	19.638	0.000	16.857	16.554	0.187
Donuts delivered	12.207	11.758	0.000	10.789	10.405	0.000	17.328	17.514	0.180	16.234	15.792	0.103
Donuts eaten	11.705	11.333	0.008	10.319	9.715	0.000	16.098	16.181	0.623	15.247	13.996	0.000
Payment rate	0.916	0.896	0.004	0.890	0.887	0.625	0.914	0.887	0.003	0.910	0.895	0.105
Total variable cost	13.431	12.758	0.000	9.791	9.825	0.652	11.504	11.480	0.724	10.542	10.322	0.042
Total revenue	32.010	34.626	0.000	22.436	22.952	0.067	23.121	24.293	0.000	22.281	22.274	0.985
Total profit	18.579	21.867	0.000	12.646	13.127	0.053	11.616	12.813	0.000	11.744	11.958	0.468

Note: Results in table are means of the three deliveries immediately preceding and following a price change for all customers who have three such deliveries in the month before and after the price change. All values are in nominal dollars. The *P*-values shown are the *p*-values of a before-period indicator variable in a regression where the indicator variable and company dummy variables are regressed on the variable of interest.

A summary of the results is presented in Table 6. The columns of the table correspond to my estimate of the likelihood of a stock-out. Column 1 represents the 10% of cases with my highest predicted likelihood of a stock-out. Column 2 reflects the half of the observations most likely to stock-out. Columns 3 and 4 are the bottom half and bottom 10% of predicted stock-outs respectively. The rows of the table report what actually happened in terms of leftovers for deliveries in each of these categories. The top panel presents results for bagels, the second panel corresponds to donuts.

The results in Table 6 demonstrate that the regression model has some power to predict delivery outcomes out of sample. Bagel stock-outs occur 52.0% of the time versus only 26.1% of the time where the model predicts bagel stock-outs to be most likely, top panel column (1), and least likely, top panel column (2), respectively. The fact that the probit model is able to generate systematic out-of-sample predictions means that the firm is not optimally incorporating all of the available information into its delivery decisions and could improve its performance using this information.[23] The table also reports the expected profits from the delivery of the $N-1$, N, and $N+1$ bagel or donut within each of the subsets identified by the regressions. For the 10% of observations where the probit model predicts the greatest likelihood of a bagel stock-out, it would have been profitable to deliver one more bagel. For the 50% of deliveries with the lowest stock-out probabilities, lowering the quantity delivered by one would have increased profits, and for 10% of the sample a reduction of two bagels would have been improved profits. Overall, the regression model adds only trivially to what the firm already does. The increment to profits that would have been associated with incorporating the probit predictions into what the firm already was doing is less than $2000 combined over the entire sample.

The probit model identifies slightly greater deviations for donuts, as shown in the bottom panel of Table 6.[24] Profits could have been increased by delivering two more donuts to half of the customers and one or two less to the other half. The contribution to profits of implementing the model's predictions with respect to donut quantities is approximately $1500 total over the sample.

5. Testing for profit maximization in the choice of prices

In contrast to quantities supplied, which vary for each customer on a delivery-by-delivery basis, there is little price variation. Over the course of 13 years, nominal prices changed on only four occasions, as noted earlier. Bagel prices increased from 60 to 75 cents on August 31, 1993, from 75 to 85 cents on August 4, 1998, and from 85 cents to one dollar on March 5, of 2003. The price of donuts is 50 cents until March 28, 2005, after which it rises to 60 cents.[25] While there is no reason to believe that the timing of these price jumps were completely exogenous, discussions with the firm suggest that the price changes were not precipitated by any perceived discrete shifts in consumer demand.

[23] Note that this model takes as an input the firm's choice of quantity to deliver, so the way that the firm would utilize the model would be to come up with quantity choices using their current approach, and then adjust quantities on the margin using the model. As a testimony to the skill of the firm in choosing quantities, I have been unable to build a model using only previous delivery outcomes that outperforms the firm's choices.

[24] In defense of the firm's quantity choices, it is worth remembering that the firm delivers in increments of six donuts, whereas this analysis assumes that the firm can costlessly vary the quantity delivered by any discrete number of units.

[25] According to the founder of the firm, the price increases were related to changes in costs and the perception that the prices they were charging were lower than competitors, such as Dunkin Donuts. His choice of prices was not constrained by agreements with office managers at his clients to charge particular prices or an implicit or explicit promise to keep prices low.

Table 8
The pattern of profits around price changes.

	Total Profits		
	(1)	(2)	(3)
Months in advance of price change:			
6 months prior	–	–	−15.045
	–	–	(18.305)
5 months prior	–	–	−17.110
	–	–	(18.695)
4 months prior	–	–	−7.406
	–	–	(17.175)
3 months prior	–	12.882	17.954
	–	(16.068)	(17.125)
2 months prior	–	10.860	16.349
	–	(16.158)	(17.197)
1 month prior	−16.242	−5.893	−0.256
	(16.101)	(16.479)	(17.590)
Months following a price change:			
1 month after	39.002	49.868	55.610
	(15.426)	(15.755)	(16.878)
2 months after	–	49.112	54.922
	–	(15.907)	(16.902)
3 months after	–	46.338	53.033
	–	(16.068)	(17.064)
4 months after	–	–	32.411
	–	–	(16.895)
5 months after	–	–	30.358
	–	–	(17.323)
6 months after	–	–	7.361
	–	–	(16.593)
Year dummies included?	Yes	Yes	Yes
Month dummies included?	Yes	Yes	Yes
Day of the week dummies included?	Yes	Yes	Yes
R^2	0.651	0.653	0.654
F-test:Months prior jointly equal to zero?	0.313	0.740	0.640
F-test:Months after jointly equal to zero?	0.012	0.000	0.001
F-test:Months after jointly different from months prior	0.009	0.001	0.000

Notes: The results in the table are regression coefficients from a time-series regression with one observation per day. The dependent variable is total profit for the day. Year, month, and day of the week dummies are included in all specifications. Standard errors in parentheses. The bottom three rows of the table report *p*-values from tests that the coefficients on the month(s) immediately preceding or following the price changes are jointly statistically significant different than zero. The number of observations included in the regression is 2262.

Table 7 presents a simple comparison of the raw data (on a per delivery basis) between the three deliveries that immediately precede and the three that follow each of the price changes for the set of customers for whom there are at least three deliveries in the month before and the month after the price change. Columns 1 and 2 correspond to numbers before and after the 1993 bagel price change; column 3 presents the *p*-value of the difference between the values in those two columns, controlling for fixed effects by customer. For the great majority of entries in the table, the differences before and after are highly statistically significant. Before the price increase, an average of 52.1 bagels are delivered, of which 48.9 are eaten. After the price change, the number of bagels delivered falls to 49.1, with 44.5 consumed. The number of donuts delivered and eaten, surprisingly, also falls slightly. Given the earlier cannibalization results, one would expect consumers to substitute from bagels to donuts in response to the price increase, but that is not observed. The data instead suggest that the price increase leads to ill-will among customers. Also consistent with this hypothesis is that the payment rate falls from .916 to .896 with the price increase – a statistically significant pattern that is repeated for all four price increases.

After the price rises in August 1993, the firm's revenues increase, implying that (at least with respect to demand over a one month horizon) the firm was operating on the inelastic portion of the demand curve. This cannot be short-run profit maximizing, and indeed profits are roughly 18% higher in the month after the price change than they were in the prior month. Under the (admittedly crude) assumption of a linear demand curve with a slope implied by the change in quantity associated with this price increase, I estimate the bagel price that maximizes short run profit to be $1.26 at that time – far above the actual price, even after the price increase.[26] If that linear demand curve is indeed accurate – which

[26] This estimate ignores any substitution from bagels to donuts when the price of bagels rises. If consumers switch to donuts, then the optimal price of bagels would be even higher, because some fraction of the lost bagel sales will be recouped by increased donut sales. Empirically, however, there is no evidence from any of the four price changes that raising the price of one good materially increases the quantity consumed of the other good.

is highly uncertain given how far out of sample the extrapolation to the optimal price is – then the firm's profit at the optimal price would be roughly \$11 higher per delivery, or nearly 60% greater than the profits achieved before the price change.[27]

The remaining columns of Table 7 report the results from the other three price changes observed in the data. In every case, after the price increase, the quantity eaten of that good falls.[28] These declines, however, are not sufficient to offset the increased revenue due to the higher price. Thus, at the time of the first three price changes, it appears that the company was pricing on the inelastic part of the demand curve. Profit rises an average of 5–6% over a one-month horizon in response to these latter three price changes. The optimal price of a bagel, based on the assumption of a linear demand curve, is \$1.21 in 1998 and \$2.33 in 2003. It is encouraging that the estimates generated from the first two bagel price changes produce similar values for the optimal price, although the implied value corresponding to the 2003 price change appears too high. On a recent visit to Dunkin Donuts, a bagel and cream cheese sold for \$1.79. The implied optimal price of a donut based on the 2005 price change is 99 cents, compared to 69 cents at Dunkin Donuts.

One important caveat regarding the apparent mispricing of the goods is that, although price changes are associated with immediate increases in profits, demand may be more elastic in the long run than the short run. In that case, raising prices may have adverse consequences for profits down the road. To explore this possibility, I examine the pattern of profits in the two months preceding the price change and up to six months after the price change. In particular, I estimate time series regressions of the form

$$Profit_t = \alpha + \sum_{m=1}^{2} \beta_m PriceChange_{t+m} + \sum_{m=1}^{6} \theta_m PriceChange_{t-m} + \Gamma X_t + \epsilon_t$$

where t indexes time and m corresponds to the number of months preceding or following a price change. The data are aggregated across all deliveries in a given day to generate one total profit number for each day. Dummy variables for the calendar year, the month, and the day of the week are included in the specification.

Table 8 reports the results of the estimation for specifications that consider various time windows before and after the price change. Column 1 presents results for a one-month window on either side of the price change. Relative to other days in that calendar year, profits are roughly \$16 (standard error of 16) a day lower in the month before the price change and \$39 higher in the month after (standard error of 15), for an increase of profits of \$55 a day, which is highly statistically significant. In column 2, estimates are presented for the three months before and after a price change. The months leading up to the price change are all insignificantly different from zero. The coefficients on the months after the price change are positive and significant with values between 46 and 49. There is no evidence that the increment to profits associated with the price change decays over this time period. When looking at the six months before and after the price change, the point estimates on the period four to six months after the price change, while still positive, are smaller in magnitude than those in the first three months, and lose statistical significance. Thus, consistent with economic theory, it does appear that demand is more elastic at longer time horizons. Using the estimates from the last column, daily profits are roughly \$40 per day higher after the price change, implying an increase in profits over the six month period of more than \$ 4000. [29] Subject to the limitations of the data, there is no strong evidence that the price changes provide benefits beyond that six month window. A further caveat on the issue of pricing is that the observed price increases did not lead to entry by competitors. A larger price increase would have been more likely to induce a direct competitor to enter, negatively impacting this firm's profitability.

6. Conclusion

Profit maximization is one of the most fundamental assumptions in economics, yet is rarely directly testable because of data limitations and the complexity of most firms. This paper provides a case study of the decision making of a firm whose activities and administrative data records lend themselves quite naturally to an analysis of the choices it makes regarding choice of prices and choice of quantities to deliver conditional on these prices. In addition, the primary decision maker in the firm is an MIT trained economist. An analysis of the data suggests that the firm does an exceptionally good job of making the daily decision regarding the quantities to deliver to customers. In stark contrast, the evidence seems to suggest that the firm has charged prices below the optimal level for more than a decade, although the conclusions about pricing are subject to important caveats since short-run responses to price changes are more readily observable than long-run responses.

The findings of this paper apply directly only to the firm in question, but there may be reasons to believe that the results obtained here are more broadly generalizable. It is not by chance that the firm does well in choosing what quantity to deliver.

[27] The estimated number of bagels consumed at a price of \$1.26 is roughly 30. In my optimal profit calculation, I assume that the firm will have the same number of uneaten bagels at the optimal price, that the payment rate for bagels will fall to .85, and that the profit generated by donuts will be unchanged.

[28] The quantities of the other good change only slightly in response to the price increase, declining in three of the four instances.

[29] There are no statistically significant differences between the rate at which customers appear or disappear from the sample before and after price increases.

Rather, by observing daily feedback about sales relative to product delivered, the firm is able to draw on an enormous volume of data in making the quantity decision. Any mistakes on the quantity dimension are quickly revealed and corrected. In contrast, the firm rarely changes prices – only four times in the entire 13 year period – and thus gets little feedback regarding the right price. Without feedback, the firm has no direct mechanism for learning whether it is pricing correctly.

The fundamental question that remains unanswered, however, is why this firm did not actively seek to create mechanisms to provide feedback on pricing. [30] For instance, since the firm serves customers at many locations who presumably have little or no contact with one another, it would have been possible for the firm to charge different prices across locations without the customers being aware. By varying prices exogenously, the firm could have learned about the elasticity of customer demand, as well as potentially price discriminating across customers.

This firm is not alone in missing out on opportunities to increase feedback to drive learning. The most straightforward and powerful tool for generating feedback is a randomized experiment. Whereas randomization is the norm in clinical studies, it is rarely used in business settings (Levitt and List, 2009). Understanding why this is the case stands as an important direction for future research.

7. Data appendix

The data used in this paper were provided by the founder of the company. Recent years of data were in Excel spreadsheets; earlier years were in hard copy that were scanned or hand entered.

The primary use of the data was for accounting purposes. A number of steps were taken to clean the data for the analysis in this paper. I exclude all clients in which the company prepays, as opposed to the goods being purchased on the honor system. In a handful of cases, there were small variations in company names over time, although it is clear from the data that these represent the same customers and were treated accordingly. When multiple deliveries were made to different floors of the same company, or to the same floor on different days of the week, these were treated as different customers. I excluded all observations in which the quantity delivered of bagels and donuts are both reported as zero, and all observations in which no quantities eaten were reported in the data. On the advice of the company, I excluded cases in which the revenues received were either less than 25% of the number that would have been expected based on the posted prices and the number of products consumed, or more than 150% of the expected number; the company says that occasionally they lump together the results of multiple deliveries to a client in their bookkeeping. In a small number of cases, the accounting identity regarding costs, revenues, and profits for a delivery did not hold. These observations were dropped, as were duplicate entries corresponding to the same delivery. I also dropped a small number of cases in which the quantities delivered were far higher than in the preceding and following weeks.

The data I was given has information on the total cost associated with the products delivered to a given client on a particular day, but is not specifically broken out between bagels, cream cheese, and donuts. To estimate marginal cost of a bagel (including cream cheese), I limit the sample to customers who only receive bagels (no donuts) and run regressions by calendar year with the total product cost on the left-hand side and the number of bagels delivered on the right-hand side, with no constant. Because the wholesale cost of the product fluctuates within a calendar year and the ratio of bagels to cream cheese varies by customer, the fit of the regression is less than one. In all but two years of the sample, the R-squared of this regression is above .99. I then estimate the marginal cost of a donut by subtracting off the fitted value of the costs of the bagels delivered and running a parallel regression on donuts. The R-squared of this regression is above .97 in all but two years. The resulting data set includes just over 80,000 valid customer deliveries.

References

Berry, Steven, Levinsohn, James, Pakes, Ariel, 1995. Automobile prices in market equilibrium. Econometrica 63 (4), 841–890.
Bresnahan, Timothy, 1987. Competition and collusion in the American automobile industry: the 1955 price war. J. Ind. Econ. 35 (4), 457–482.
Bresnahan, Tim, Reiss, Peter, 1991. Entry and competition in concentrated markets. J. Political Econ. 99 (5), 977–1009.
Carlton, Dennis, 1978. Market behavior with demand uncertainty and price inflexibility. Am. Econ. Rev. 68, 571–587.
Dana, James, 2001. Competition in price and availability when availability is unobservable. RAND J. Econ. 32 (3), 497–513.
Deneckere, Raymond, Peck, James, 1995. Competition over price and service rate when demand is stochastic: a strategic analysis. RAND J. Econ. 26 (1), 148–162.
Feldman, Paul, 1971. Efficiency, distribution, and the role of government in a market economy. J. Political Econ. 79 (3), 508–526.
Genesove, David, Mullin, Wallace P., 1998. Testing static oligopoly models: conduct and cost in the sugar industry, 1890–1914. RAND J. Econ. 29 (2), 355–377.
Hadley, G., Whitin, T.M., 1963. Analysis of Inventory Systems. Prentice-Hall, Inc, Englewood Cliffs, NJ.
Hortacsu, Ali, Puller, Steven, 2005. Testing strategic models of bidding in deregulated electricity markets: a case study of ERCOT. University of Chicago Working Paper.
Khouja, Moutaz, Mehrez, Abraham, Rabinowitz, Gad, 1996. A two-item newsboy problem with substitutability. Int. J. Prod. Econ. 44, 267–275.
Levitt, Steven, 2006. White-collar crime writ small: a case study of bagels, donuts, and the honor system. Am. Econ. Rev. 96 (2), 290–294.
Levitt, Steven, Dubner, Stephen, 2005. Freakonomics: A Rogue Economist Explores the Hidden Side of Everything. William Morrow, New York.

[30] Discussions with the firm's owner suggest that one reason for infrequent price changes was menu costs – the prices were attached to the boxes in which the money was collected in a way that made changing the listed prices time consuming. He also feared a negative reaction to price changes on the part of customers, and thus favored infrequent changes. Experimenting with price changes is something he never seriously considered.

Levitt, Steven, List, John, 2009. Field experiments in economics: the past, the present, and the future. Eur. Econ. Rev. 53 (1), 1–18.
Mills, Edwin, 1959. Uncertainty and price theory. Q. J. Econ. 73 (1), 116–130.
Pakes, Ariel, 1986. Patents as options: some estimates of the value of holding European patent stocks. Econometrica 54 (4), 755–784.
Rosse, James N., 1970. Estimating cost function parameters without using cost data: illustrated methodology. Econometrica 38 (2), 256–275.
Wolfram, Catherine, 1999. Measuring duopoly power in the British electricity spot market. Am. Econ. Rev. 89 (4), 805–826.