Dynamic Demand Elasticities in Promotion Cycles: Implications for Experimental Design and Endogeneity Correction*

Pedro M. Gardete, Tânia Cotrim

June 2025

Abstract

This paper develops a continuous-time model in which cyclic price promotions endogenously generate time-varying demand elasticities, even when consumer preferences are static. In the model, price promotions sort consumers with heterogeneous valuations into waiting or purchasing states, leading price sensitivity to evolve over time as promotions unfold. We characterize two empirical challenges that arise in such settings. First, we show that experiments introducing promotions out of sync with the market's natural promotion cycles can produce biased elasticity estimates, with the bias depending on both the timing and duration of the intervention. Second, we demonstrate that even when prices are exogenous, correcting for potential endogeneity via lagged-price instruments can introduce finite-sample biases due to the discrete nature of price changes during promotions. Using experimental data from Elberg, Gardete, Macera, and Noton (2019), we assess the magnitude of these effects and show they can be economically meaningful. The results provide practical guidance for researchers designing price experiments or analyzing promotional data with endogeneity corrections.

*Gardete (pedro.gardete@novasbe.pt) and Cotrim (tania.sofia.silva@novasbe.pt): Nova School of Business and Economics. Valuable input from André Trindade, Miguel Villas-Boas, and Rodrigo Belo is appreciated. This work was funded by Fundação para a Ciência e a Tecnologia (UIDB/00124/2020, UIDP/00124/2020 and Social Sciences DataLab - PINFRA/22209/2016), POR Lisboa and POR Norte (Social Sciences DataLab, PINFRA/22209/2016).

1 Introduction

The concept of demand elasticity is at the core of economic and marketing endeavors such as demand analysis, pricing, promotion design, market power estimation, simulation of the implications of mergers to competitors and consumers, etc. In empirical research, demand elasticity is a key ingredient to characterize demand, quantify market power, estimate consumer welfare, and simulate counterfactual pricing scenarios. Indeed, the variety of uses of this simple measure may partly explain the attention it has deserved by marketing scientists.[1]

Price promotions are a core component of many retailers' pricing strategies. Often retailers rely on price promotions as their default pricing method, a strategy called Hi-Lo pricing (Bell and Lattin (1998), Hitsch, Hortacsu, and Lin (2021)). Reflecting the prevalence of this strategy, Lloyd, Morgan, McCorriston, and Zgovu (2009) find that price promotions account for 40% of price variation in the UK foods market once retailer base-price differences are taken into account. This paper investigates how price promotions – an important source of the price variation used to inform demand elasticities – can be used by researchers to obtain reliable elasticity estimates. It identifies challenges associated with using promotional data and discusses methodological remedies and best practices to address them.

Demand elasticity is often considered as a single stable measure of price sensitivity of demand, but in reality it is bound to vary over time due to changes in preferences (trends and seasonal events) as well as economic, environmental and/or political shocks (see for example Villas-Boas and Winer (1999), Dhar, Chavas, and Gould (2003), Ackerberg and Rysman (2002), and Kim and Petrin (2015)). A common view is that sellers and manufacturers may find it optimal to navigate those shocks by offering price promotions over time. This paper considers the different scenario in which a seller offers promotions due to facing demand heterogeneity, and it is those promotions that generate fluctuations in the composition of demand, and in the dynamics of demand elasticity over time.

We develop a continuous-time model in which a firm faces consumers with heterogeneous valuations. While high-types are willing to buy at 'any price', consumers with low valuations are only willing to buy at a discount. When the latter group faces regular prices, however, some of its members are willing to wait to take advantage of the next promotion. In such cases, it may be optimal for the seller to offer promotions in order to clear the stock of waiting low-type consumers. This cyclical sorting process implies that the elasticity of demand is not fixed, but endogenously shaped by the history of pricing decisions and the temporal

[1] Examples include Guadagni and Little (1983), Tellis (1988), Kamakura and Russell (1989), Mela, Gupta, and Lehmann (1997), Bijmolt, Van Heerde, and Pieters (2005), Hitsch, Hortacsu, and Lin (2021), Datta, van Heerde, Dekimpe, and Steenkamp (2022), Bray, Stamatopoulos, and Sanders (2024), among many others. See also the Elasticity Open Science initiative, www.elasticity-open-science.com.

dynamics of market participation.

In equilibrium, we find that the duration of the promotion cycles generally increases with the valuation of the high-type consumers and the seller's menu costs, and decreases with the valuation and the arrival rate of low-type consumers. Moreover, as the stock of waiting consumers increases during the regular price period, so does the price sensitivity of aggregate demand. The implications for demand elasticity are not trivial. Throughout the promotion cycle, demand elasticity may change non-monotonically over time, with cycles of rising and falling sensitivity emerging naturally even in the absence of forward-looking behavior, inventory effects, or stochastic demand shocks. The model is useful to derive specific conclusions about the challenges faced by empirical researchers trying to estimate demand elasticities in promotion contexts.

The first set of challenges we address is related with experimental interventions. Because promotions endogenously sort consumers into waiting or buying states, elasticity at any given moment reflects the history of pricing decisions. If an experiment introduces a promotion out of sync with the market's natural cycle, the measured elasticity may not reflect the market's equilibrium behavior. We quantify the biases that a naive analyst will induce by failing to respect the equilibrium promotion rhythms of the market. Assuming our promotion cycles model as the data generation process, we find that introducing the experimental price promotion earlier than what would have been dictated by the equilibrium timing will underestimate price sensitivity (and vice versa). Failing to match the duration of the experimental intervention with the equilibrium one will likely generate biased estimate of elasticity. In this case, the causes underlying the sign of the bias are more complex, as it depends on the curvature of sales during the promotional period. For example, when promotional sales are strictly concave over time, an analyst who runs a promotion that is shorter than the market's natural promotions may incorrectly believe that the market is more price sensitive than what it actually is. Our results also indicate that historical data is valuable to determine whether failing to match the experiment's promotional rhythms with the market will induce biases to elasticity estimates. While it is not surprising that failing to control for promotional timings will affect elasticity estimates, implementation constraints often lead researchers to conduct experiments in a pragmatic manner and possibly fail to match the characteristics of the natural market rhythms, with implications to their findings.

The second set of challenges pertains to the use of instrumental variables. For example, in transaction data, the occurrence of price promotions may be correlated with other display activities in the store that are unobservable to the analyst, and the rationales and mechanisms of this procedure are well documented in the literature. We focus on the less investigated case of spurious instrumentation, that is, the potential bias introduced by attempting to correct for price endogeneity despite its absence. When price changes occur in discrete jumps – such as the start and end of promotions – lagged prices systematically misalign with current prices at these transition points. This mismatch generates out-of-phase observations that create finite-sample bias in instrumental variable estimates, even when price is exogenous. While the signal of the bias is unconstrained, we show that inspection of promotional sales is enough to sign the bias, and our empirical analysis and institutional knowledge suggests that the bias may amplify estimated elasticities. In addition, when the sampling resolution available to the analyst is fine enough (e.g., the analyst holds daily rather than weekly data), we show that the spurious instrumental variables bias will always lead to an overestimate of price sensitivity, vis-a-vis the uninstrumented estimate. While finite sample biases seem irrelevant in the era of big data, we note that the bias we document is introduced by the analyst lacking enough observations in the temporal dimension rather than the cross section. This means that the fact that large retail datasets feature transactions by many consumers does not attenuate the bias. Our empirical application further confirms that the spurious instrumental variables bias can be economically significant.

The topics of price promotions and price elasticity are extremely vast. We refer the reader to the review sources Tellis (1988), Bijmolt, Van Heerde, and Pieters (2005), Rao (2009), and Anderson and Fox (2019). The idea that price elasticities change over time has typically been considered in the literature in contexts different from retail. In the context of durables, Parker (1992) finds that demand elasticities (in absolute value) do not necessarily increase as the adoption process evolves, which is surprising in light of the prediction that early adopters should be less price-sensitive than later adopters. In the case of frequently-consumed goods, Simon (1979) considers the case of frequently purchased branded items who nonetheless are situated at different stages of their individual brand life cycles, and finds a U-shaped evolution of (absolute) price elasticities over time. A notable exception that links price promotions to dynamic price elasticities is the work by Fibich, Gavious, and Lowengart (2005), where dynamics are induced by reference-dependent preferences. While much of the prior literature has focused on time-varying elasticities driven by dynamic consumer preferences or forward-looking behavior, this paper shows that even with static preferences, the practice of cyclic price promotions creates endogenous variation in price sensitivity over time. This insight has important implications for empirical researchers conducting price experiments or applying endogeneity corrections in transactional data.

At the core of the analysis is a model of promotion cycles inspired by Conlisk, Gerstner, and Sobel (1984), who show that price promotions may occur in a stationary environment in terms of consumer preferences and perfect consumer information. The insight in Conlisk, Gerstner, and Sobel (1984) is that if low-type consumers are willing to wait for a sale and then buy the product in a single day, it is best for the seller to offer periodic promotions to clear them from the market, without having to lose much value due to time discounting or to sales made to a few high-types at a lower price. The mechanism underlying our model is similar to the one described above, but we consider the continuous-time case, which allows us to model purchases by low-type consumers during a non-instantaneous promotional period. In fact, the rates at which low-type consumers buy during the promotional phase is at the core of the empirical implications we derive in our analysis.

The model is also inspired by Villas-Boas and Villas-Boas (2008), and Freimer and Horsky (2012). Villas-Boas and Villas-Boas (2008) develops a continuous-time model of promotions based on Bergemann and Välimäki (2006) where consumers learn and forget about their experiences with the product over time. The model solution of the continuous-time model in Villas-Boas and Villas-Boas (2008) is obtained by taking the limit of the seller's planning horizon T to zero, at which point it is possible to characterize the fraction of the duration of each promotional phase as well as conduct comparative statics. The limit approach is a sophisticated and compelling solution strategy but less suited to our analysis, which requires a precise characterization of the equilibrium sales over the whole planning period. Freimer and Horsky (2012) consider a discrete-time model where two competing firms can offer promotions to attract customers from each other as well as from an outside option. These consumers exhibit positive purchase reinforcement from a company after buying from it, a phenomenon typically designated as 'state dependence' in the empirical literature. Freimer and Horsky (2012) find that it is optimal for sellers to offer promotions in different periods. The focus on discrete time, however, means that the results are obtained on a case-by-case basis, with Freimer and Horsky (2012) considering the case of a monopolist offering promotions up to every fourth period. In order to produce a more tractable model that allows us to consider empirical implications, we employ the assumption from Villas-Boas and Villas-Boas (2008) that the seller cares about a fixed planning period and abstract away from time discounting. For example, the seller may devise its promotion plan for each year or season. In this case, even without time discounting, the seller prefers to set promotion cycles with finite durations because this allows is to extract value from low-type consumer multiple times during the planning horizon.[2]

Finally, Erdem, Imai, and Keane (2003) and Hendel and Nevo (2006) develop empirical models in which forward-looking consumers take advantage of promotions to stockpile goods. Both studies find that accounting for forward-looking behavior significantly affects elasticity

[2] In fact, introducing time discounting would lead managers to plan different promotions of different durations over the course of the year, simply due to the time-value of money changing, which we believe may be unlikely.

estimates. In contrast, our focus is not on determining that parameter estimates can be biased due to misspecification. Instead, we focus on how analysts mistakes, namely of promotion timing and unnecessary use of instrumental variables, can introduce potentially unexpected biases, and provide guidance in each case.

The next section presents the model of promotion cycles, which guides the analyses throughout the rest of the paper. Section 3 explains the equilibrium dynamic nature of demand elasticity. Section 4 analyzes the implications of the model for analysts who fail to take into account the promotion rhythms when conducting experimental interventions. Section 5 investigates the potential for spurious instrumental variable bias when relying on lagged prices. Section 6 utilizes the promotions data from Elberg, Gardete, Macera, and Noton (2019) to ascertain the magnitude of the theoretical biases, and Section 7 offers some concluding remarks and avenues for future research.

2 A Model with Promotion Cycles

Demand structure. We develop a simple model of price promotions. Consider the case of a market with two types of consumers, high and low, exhibiting product valuations $v_H > v_L > 0$, and continuously flowing to the market. High-type consumers arrive at rate equal to 1, and low-type consumers arrive at rate μ, which captures the ratio of arrivals between low- and high-type arrivals.[3] We consider the case where high-type consumers are impatient, such that they buy immediately as long as price is less than or equal to their valuation v_H. Low-type consumers who arrive to the market buy immediately as long as the price is less than or equal to v_L. If the price is higher than v_L, they may decide to leave or wait for a promotion to occur. We define the stock of waiting (low-type) customers as $W(t)$, where t is the accumulation time of these consumers. We assume $W'(t) \geq 0$, $W''(t) \leq 0$, and $\lim_{t\to\infty} W(t) = \omega$. The concavity of $W(t)$ reflects the fact that some consumers give up waiting for a low price, possibly to take advantage of some passive outside option or not consume at all. The last limit condition implies that the stock of waiting consumers does not grow forever, but approximates a steady-state level over time. Although not a mathematical requirement, it also makes sense to assume that $W'(t) \leq \mu \, \forall t$, since the stock of $W(t)$ cannot accumulate faster than the arrival rate of the low-type consumers.

Given this demand structure, it makes sense for the firm to only practice prices v_L or v_H at any point in time. We consider the case of a seller facing a planning horizon of duration

[3] In settings where consumers are interested in buying more than one unit, this parameter may also reflect the relative difference in quantities purchased across consumer types.

T (e.g., one year), in which it can launch price promotions.[4] Suppose that at time 0 the seller just finished a period of promotional price, and is starting a period of regular price. We define t_1 as the duration of the regular price period and $t_2 - t_1$ as the duration of the subsequent promotional price period. Given consumers' preferences, we inspect the policy where the seller sets price v_H during the first phase and v_L during the second. The intuition for this pricing policy is that the seller may be better off waiting for low-type consumers to accumulate, to then sell the product to them at the promotional price. This policy may be profitable as long as the waiting low-type consumers buy the promoted product fast enough.

Regardless of whether the seller sets price v_L or v_H, high-type consumers buy at a rate equal to their arrival rate, of one. As for the low types, they only buy when the price is equal to v_L. During the promotional phase there exist two types of low-type customers: those who were already waiting and those who arrived in the meantime to find a low price. Letting t_1 be the duration of the regular price regime over which low-type consumers have accumulated, and Δ the duration of a subsequent promotional period, we define the total demand by waiting consumers as

$$D\left(t_1, \Delta\right) = W\left(t_1\right).\Gamma\left(\Delta\right), \forall t_1 > 0, \Delta > 0 \tag{1}$$

The demand by low-type consumers at price v_L depends on the accumulated stock $W\left(t_1\right)$ and on function $\Gamma\left(\Delta\right)$, the latter representing the fraction of the waiting consumers who buy during a period of duration Δ. We define $\Gamma : [0, \infty) \to [0, 1]$ to be weakly increasing, so that the total mass of waiting consumers who purchase increases as the promotional period elapses, and $\Gamma\left(0\right) = 0$. We focus on the case where the seller is better off promoting until it clears all waiting low-type consumers. This means that at the end of the promotion cycle and at all moments afterwards, say at some point in time Δ', $\Gamma\left(\Delta\right) = 1$, $\forall \Delta \geq \Delta'$. The exact shape of Γ is not restricted.

Note that the multiplicative specification of $D\left(\cdot, \cdot\right)$ is equivalent to assuming that the time to clear waiting customers is independent of their number, since function $\Gamma\left(\cdot\right)$ depends only on the time interval Δ and not on the base $W\left(t\right)$. This assumption is not crucial for the analysis but provides easier interpretability to our results.[5] It is, however, compatible with situations in which the clearance time of waiting customers is generally constant. This applies, for example, in situations where consumers' buying times are inversely proportional to the time waiting for the promotion. In such cases, consumers who have been waiting longer for a promotion may be more eager to acquire the product once it is offered on promotion.

[4] We assume the planning horizon is short enough so as to make discounting concerns irrelevant.

[5] The solution of the model is tractable without this assumption, and an analysis is available from the authors. In this case the seller optimizes both the promotion and the regular period durations.

This assumption translates to the following condition:

$$t_2 = t_1 + \gamma_1 \tag{2}$$

where γ_1 is the exogenous clearance time of the waiting consumers, such that time t_2 – the time at which the promotion cycle completes – is composed of the duration of the regular-price phase, t_1, and the duration of the promotional phase, γ_1. Constant γ_1 can be interpreted as the duration at which $\Gamma(\cdot)$ first equals one, that is, $\gamma_1 := \min \Delta : \Gamma(\Delta) = 1$.[6]

Pricing problem. We assume the firm faces menu costs (Villas-Boas and Villas-Boas (2008)), with each price change leading to it incurring cost $\frac{k}{2}$. During a full promotion cycle in which the firm charges price v_H until time t_1 and then charges $v_L$from time $t_1$to t_2, it accrues profit

$$\Pi_{cycle} = \underbrace{v_H t_1 + v_L (t_2 - t_1)}_{\text{Revenue from High Types}} + \underbrace{v_L W(t_1) \Gamma(t_2 - t_1) + v_L \mu (t_2 - t_1)}_{\text{Revenue from Low Types}}$$

Above, the firm continuously sells to high types, albeit at different prices. By the end of the promotion cycle it has also cleared the low-type consumers who decided to wait to buy at a lower price, plus the low-type consumers who arrived during the promotional phase.

In each promotion cycle, the firm incurs total menu costs of k, and during the planning period T it goes through $T \div t_2$ cycles. Hence, during its planning period, the firm accrues profit

$$\Pi^* = \max_{t_1,t_2} \frac{T}{t_2} \left(v_H t_1 + v_L (t_2 - t_1) + v_L W(t_1) \Gamma(t_2 - t_1) + v_L \mu (t_2 - t_1) - k \right) \tag{3}$$

$$s.t.\, t_2 = t_1 + \gamma_1 \tag{4}$$

Market Outcome. The optimal (interior) duration of the regular price phase t_1^* satisfies

[6] It is trivial to consider the case where the seller prefers serving a pre-specified set of waiting consumers, by redefining parameter γ_1. For the case of selecting the mass of low-type consumers to serve endogenously (e.g., the last few consumers may take 'too long' to buy and so the seller may prefer to move onto the regular price earlier), an interior solution in principle exists and is obtainable. Such cases may originate different comparative statics than the ones we find, but have no bearing on the empirical implications we derive in this paper, so we do not tackle them in our analysis.

the following condition:

$$foc(t_1): \frac{d\Pi}{dt_1} = 0 \tag{5}$$

$$\Leftrightarrow \frac{\partial \Pi}{\partial t_1} + \frac{\partial \Pi}{\partial t_2}\frac{\partial t_2}{\partial t_1} = 0 \tag{6}$$

$$\Leftrightarrow \frac{T\left(v_L W'(t_1^*) + v_H - (1+\mu)v_L\right)}{t_2^*} - \frac{T(t_1^*(v_H - v_L) + v_L\left(W(t_1^*) - \mu t_1^*\right) - k)}{(t_2^*)^2} = 0 \tag{7}$$

$$\Leftrightarrow v_H - v_L + v_L\left(W'(t_1^*) - \mu\right) - \frac{(t_1^*(v_H - v_L) + v_L\left(W(t_1^*) - \mu t_1^*\right) - k)}{(t_1^* + \gamma_1)} = 0 \tag{8}$$

The first two terms of the first-order condition (expression (8)) comprise the profits accrued from high-types from increasing the duration of the regular price phase, t_1: For each unit of time extended, the firm will be earning v_H rather than v_L from high-type consumers. The third term of expression (8), $v_L\left(W'(t_1^*) + \mu\right)$, captures the same effect for low type consumers: extending the duration of the regular price accumulates $W'(t_1^*)$ more low-type consumers, but precludes sales to μ consumers for that same period of time.

The last term of expression (8) captures the fact that extending t_1 has a direct implication on the number of promotion cycles that occur in planning period T: the longer t_1 is, the fewer promotion cycles fit in T, and so the lower the profitability from engaging in dynamic pricing regimes. The end-result of increasing t_1 slightly means running slightly fewer promotion cycles, which also means foregoing profits from high-type consumers: a value of $(v_H - v_L)$ over the fraction $t_1^* \div (t_1^* + \gamma_1)$ of the promotion cycle. Similarly, the loss from introducing slightly longer cycles means accumulating fewer waiting consumers of the low-type, but it means a gain from the remaining ones who buy as they arrive. Finally, the menu cost shows up as a gain because, by prolonging the duration of promotion cycles, the seller incurs the cost less frequently.

We take advantage of the first-order condition to analyze comparative statics on the model parameters.[7]

Proposition 1. *The duration of the regular price phase, t_1^*, increases in the menu cost k, in the valuation of the high-type consumers v_H, and in the clearing time of accumulated low-type consumers γ_1, and decreases in the valuation and in the ratio of arrival rates of low to high-type consumers, v_L and μ.*

The fact that t_1^* increases with k is intuitive, since longer cycles are associated with fewer price changes overall. Duration t_1^* also increases with v_H: When high-type consumers are

[7]All proofs are presented in the appendix.

willing to pay more, the seller is better off keeping prices high for longer. Time t_1^* also increases with γ_1, the time associated with clearing the stock of low types. To see why, consider the case of a very short clearance period. This would be a very efficient case for the seller, since it can clear a large pool of waiting customers in a short amount of time. Because $W''(\cdot) < 0$, however, the seller would have an incentive to let few low-type consumers accumulate each time to minimize lost sales to consumers who decide not to wait for the promotion. So, a smaller value of γ_1 results in shorter regular-price period t_1^* and conversely, t_1^* increases with γ_1. Finally, duration t_1^* decreases with v_L and μ because when low-types become less attractive, the seller is better off focusing on high types.

We briefly consider functions $W(\cdot)$ and $\Gamma(\cdot)$ to visualize the model with precision. Let $W(t) = \omega - \exp(-t)$ and $\Gamma(\Delta) = (1+\lambda)\left(1 - \frac{1}{1+\Delta}\right)$, which satisfy the model assumptions. The following proposition follows from the discussion above.

Proposition 2. *When $W(t) = \omega - exp(-t)$ and $\Gamma(\Delta) = (1+\lambda)\left(1 - \frac{1}{1+\Delta}\right)$, the durations of the regular and promotional price periods are given by:*

$$t_1^* = \pi_{-1}\left(f\left(v_H, v_L, \lambda, \mu, \gamma, k, \omega\right)\right) - \frac{1}{\lambda} \tag{9}$$

$$t_2^* = \pi_{-1}\left(f\left(v_H, v_L, \lambda, \mu, \gamma, k, \omega\right)\right) \tag{10}$$

where

$$f(\cdot) = \frac{1}{\lambda e^{1/\lambda}} \frac{v_H - v_L(1 + \mu + \omega\lambda) + \gamma k}{v_L}, \tag{11}$$

and $\pi_{-1}(x) := -\left(1 + L_{-1}\left(\frac{1}{e}x\right)\right)$ where L_{-1} is the '-1' branch of the Lambert W function.

In this parametrization, the duration of the promotion obtained by the difference of t_2^* and t_1^* is equal to $\frac{1}{\lambda}$, which corresponds to parameter γ_1 in the general model. Figure 1 depicts the optimal duration of the promotion cycle t_2^* as a function of $V := f(\cdot)$:

Figure 1: Optimal duration of the promotion cycle

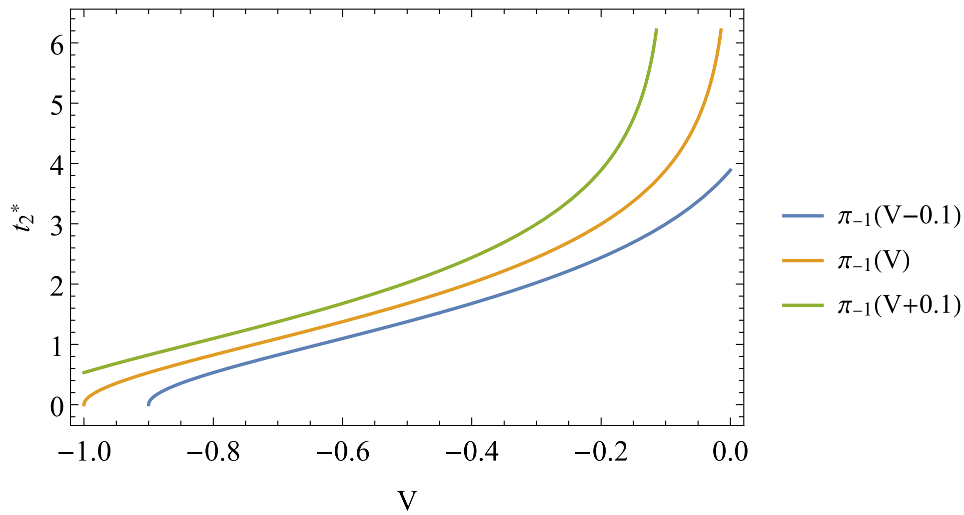

Note: Above, V stands for the argument of $\pi_{-1}(\cdot)$ in expression (10), that is, $V = \frac{1}{\lambda e^{1/\lambda}} \frac{v_H - v_L - v_L(\mu + \omega\lambda) + \gamma k}{v_L}$.

As depicted in the figure above, the duration of the promotion cycle t_2^* (and of the regular price duration, t_1^*) is increasing in V, including in parameter ω, which represents the limiting size of the pool of waiting customers.

Figure 2: Promotion Cycles

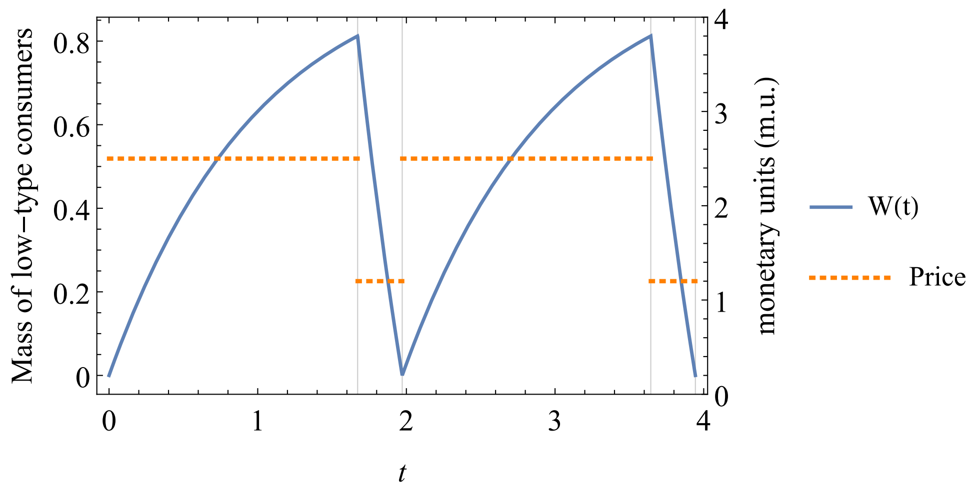

The plot above uses the following parameter values: $\alpha = 1$, $v_H = 2.5$,$v_L = 1.2$, $\gamma = 3.3$, $k = 0.5$, and $\mu = 1$.

Figure 2 shows the evolution of the stock of customers waiting for a promotion as well as the evolution of prices over time. The fact that these customers take longer to accumulate than to clear is the key for profitability of this strategy for the seller. This is also consistent with real-world patterns where most products are sold at the regular price for long periods, with comparatively short-lived promotional periods.

We now consider how the dynamics of the promotion cycles model translate to the evolution of demand elasticity.

3 Elasticity Analysis

We now specifically consider demand elasticity, that is, the relative change in demand due to a change in price in the context of a promotion. Demand elasticity is often defined as the percentage change in quantity demanded when price increases by one percent. In empirical settings, however, estimates of price elasticity are often informed by discrete increases and decreases in price. For example, researchers and analysts tend not to differentiate between price increases and decreases, by applying the formula:[8]

$$\varepsilon_t = \frac{\frac{Q_{t+1}-Q_t}{Q_t}}{\frac{p_{t+1}-p_t}{p_t}} \tag{12}$$

so that elasticity is defined by the relative change in quantity demanded divided by the relative change in price, within some temporal resolution captured by t. This elasticity measure is not time-reversible; that is, in general,

$$\frac{\frac{Q_{t+1}-Q_t}{Q_t}}{\frac{p_{t+1}-p_t}{p_t}} \neq \frac{\frac{Q_t-Q_{t+1}}{Q_{t+1}}}{\frac{p_t-p_{t+1}}{p_{t+1}}} \tag{13}$$

This means that, by construction, the elasticity measured by the introduction of a promotion is always different from the one obtained from the price variation induced by its end. It is also interesting that the inequality above holds even for the constant elastic demand function, except in the case of infinitesimal price changes. Since typically price promotions involve discrete price jumps, it is useful to define two elasticity measures, depending on whether they are obtained from introducing or concluding a price promotion:

$$\varepsilon_{p\downarrow} = \frac{\frac{Q^{promo}-Q^{regular}}{Q^{regular}}}{\frac{p^{promo}-p^{regular}}{p^{regular}}} \tag{14}$$

[8]Throughout the paper we refer to negative demand elasticities, and use the term 'price sensitivity' to refer to their absolute value.

and

$$\varepsilon_{p\uparrow} = \frac{\frac{Q^{regular} - Q^{promo}}{Q^{promo}}}{\frac{p^{regular} - p^{promo}}{p^{promo}}} \tag{15}$$

Elasticity $\varepsilon_{p\downarrow}$ is the price elasticity obtained from introducing a promotion during a period of regular price, and elasticity $\varepsilon_{p\uparrow}$ is the demand elasticity obtained from charging the regular price during a promotional period. Applying the formulas to our model of promotions, the demand elasticity from introducing a price promotion at time t for duration Δ is given by

$$\varepsilon_{p\downarrow}(t,\Delta) = \frac{\frac{\Delta + W(t)\Gamma(\Delta) + \mu\Delta - \Delta}{\Delta}}{\frac{v_L - v_H}{v_H}}, t \in (0, t_1^*) \tag{16}$$

$$= -\frac{v_H}{v_H - v_L} \frac{W(t)\Gamma(\Delta) + \mu\Delta}{\Delta}, t \in (0, t_1^*) \tag{17}$$

Elasticity $\varepsilon_{p\downarrow}(t,\Delta)$ measures the demand elasticity obtained from the price variation of a promotion introduction during a regular-price regime that started t periods ago. The reason t affects the price elasticity is because the composition of demand changes over time: as t increases, so does the mass of low-type consumers waiting for a promotion. Parameter Δ measures the time-horizon of the elasticity. When $\Delta \to 0$, for example, the elasticity measures the instantaneous or very short-run impact of introducing a price promotion, and $\Delta \to \infty$ represents the case of long-run price elasticity. The net effect on sales due to the promotion is given by $W(t)\Gamma(\Delta) + \mu\Delta$, and sources from sales to two types of low-type consumers: those who were waiting to buy and those who arrived while the promotion was taking place.

We define $\varepsilon_{p\uparrow}(t,\Delta t)$ as the demand elasticity obtained from the price variation induced by the conclusion of a price promotion:

$$\varepsilon_{p\uparrow}(t,\Delta) = \frac{\frac{\Delta - \left(\Delta + W\left(t_1^*\right)\left(\Gamma\left(t - t_1^* + \Delta\right) - \Gamma\left(t - t_1^*\right)\right) + \mu\Delta\right)}{\Delta + W\left(t_1^*\right)\left(\Gamma\left(t - t_1^* + \Delta\right) - \Gamma\left(t - t_1^*\right)\right) + \mu\Delta}}{\frac{v_H - v_L}{v_L}}, t \in (t_1^*, t_2^*) \tag{18}$$

$$= -\frac{v_L}{v_H - v_L}\left(1 - \frac{\Delta}{W(t_1^*)\left(\Gamma(t - t_1^* + \Delta) - \Gamma(t - t_1^*)\right) + (1+\mu)\Delta}\right), t \in (t_1^*, t_2^*) \tag{19}$$

When a promotion concludes, the stock of consumers who had been waiting for the promotion and have since bought is given by $W(t_1^*)\Gamma(t - t_1^*)$, where $t - t_1^*$ is the duration of the ongoing promotion. Had price not increased, the promotion would have been extended by Δ periods, so that the sales to that consumer pool would have been $W(t_1^*)\left(\Gamma(t - t_1^* + \Delta) - \Gamma(t - t_1^*)\right)$. These definitions reveal, somewhat explicitly, that elasticities depend on the counterfactual scenarios we consider, had prices followed trajectories different than the ones we observe in a given dataset. This will be relevant when we analyze the challenges involved in experimental interventions.

Specific Example. We plot the evolution of price elasticities in Figure 3, with specific functions of $W(\cdot)$ and $\Gamma(\cdot)$, for illustration purposes. So, for example, an increase in the value of the price elasticity (i.e., a 'less negative value') corresponds to a decrease in price sensitivity.

Figure 3: Specific Example: Evolution of Demand Elasticities

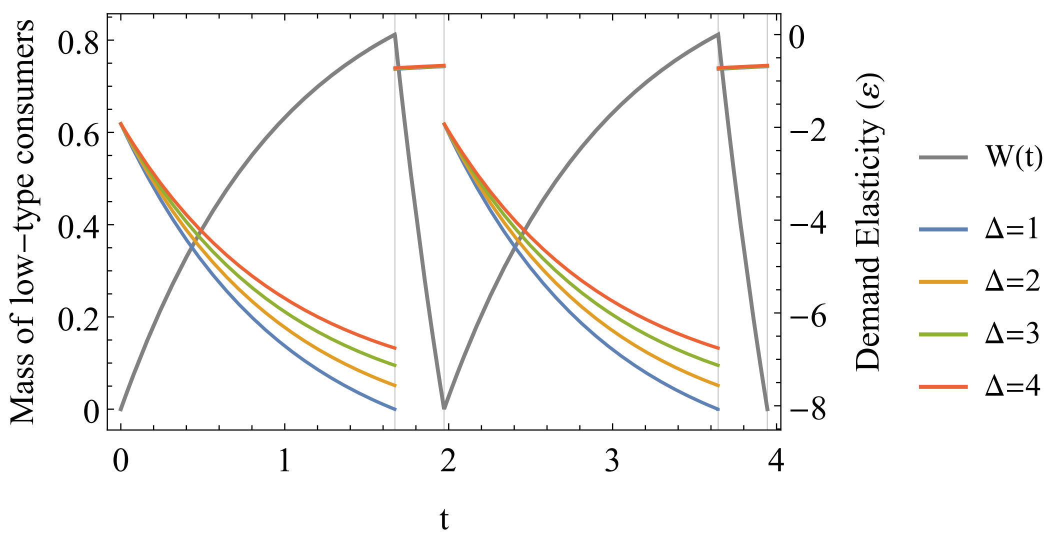

The plot above uses the following parameter values: $\alpha = 1$, $v_H = 2.5$,$v_L = 1.2$, $\gamma = 3.3$, $k = 0.5$, and $\mu = 1$.

Above, the evolution of the stock of waiting consumers is depicted in gray, together with demand elasticities, in color. The elasticity curves represent the demand elasticities induced by a price change at time t, each curve corresponding to different time horizons Δ. As explained before, in the figure elasticities are discontinuous whenever the price regime shifts, not because consumers became more price sensitive, but merely because elasticities $\varepsilon_{p\downarrow}$ and $\varepsilon_{p\uparrow}$ are different by definition.

During the first phase of the promotion cycle, in which the seller charges the regular price, introducing a promotion is associated with higher price sensitivity over time. The reason is that more and more low-type consumers accumulate during that phase, leading price sensitivity to increase. In all cases price sensitivity increases with the time horizon of the promotion, which is expected since price elasticities typically decrease as more of the impact of a price change is taken into account.

During the second part of the promotional period, price variation occurs when the seller concludes the ongoing promotion at time t. The elasticity curves – now overlapping almost perfectly – exhibit a positive slope because resuming the regular price earlier will affect more waiting consumers who still have not had the chance to buy at the promotional price. The main takeaway of Figure 3 is that demand elasticity can be dynamic even when consumer valuations are static: When promotions occur, demand becomes more price sensitive over time during regular-price periods, and becomes less sensitive over time during promotional periods. This occurs even in the absence of dynamic preferences, stockpiling incentives or common market shocks. In the next sections we investigate the question of dynamic elasticities in the general case and consider the relevance of the findings for empirical research and managerial decision-making.

4 Implications for Price Experimentation

Elasticity from introducing a promotion. We formally consider comparative statics to investigate the nature of time-varying price elasticities, in terms of the timing of the price changes as well as the time-horizon of those changes. We first focus on the elasticity obtained from introducing a price promotion.

Proposition 3. *The price sensitivity of introducing a price promotion increases with the introduction time of the promotion t. As for the time horizon, price sensitivity decreases with the time horizon of the promotion Δ if and only if $\Gamma(\Delta) - \Delta\Gamma'(\Delta) > 0$, and increases otherwise. The signs of the comparative statics of the elasticities are given by:*

$$\frac{\partial \varepsilon_{p\downarrow}(t,\Delta)}{\partial t} < 0 \tag{20}$$

$$\frac{\partial \varepsilon_{p\downarrow}(t,\Delta)}{\partial \Delta} \propto \Gamma(\Delta) - \Delta\Gamma'(\Delta) \tag{21}$$

where the proportionality symbol '$\propto$' is used to mean sign equality between the left- and right-hand sides.

Above, the comparative static $\frac{\partial \varepsilon_{p\downarrow}(t,\Delta)}{\partial t}$ is negative, such that price sensitivity of a price promotion increases with the duration of the ongoing regular-price regime, or equivalently, with the time since the previous price promotion, so that the result observable in Figure 3 applies to the general case. This is intuitive since more consumers accumulate as t increases, and the counterfactual sales (made to high-type consumers) are constant at Δ.

As for the comparative static $\frac{\partial \varepsilon_{p\downarrow}(t,\Delta)}{\partial \Delta}$, it measures the change in the demand elasticity as a function of the promotion duration Δ. The sign of this comparative static depends on the difference $\Gamma(\Delta) - \Delta\Gamma'(\Delta)$, which relates with the curvature of function $\Gamma(\cdot)$. Note that a generic linear function $\varphi(\Delta)$ with a zero intercept can be written as $\Delta\varphi'(\Delta)$, by definition, such that the difference $\Gamma(\Delta) - \Delta\Gamma'(\Delta)$ can be thought of as an accumulated deviation between function $\Gamma(\cdot)$ and a linear approximation around point Δ.

At point Δ, the deviation $\Gamma(\Delta) - \Delta\Gamma'(\Delta)$ measures whether function $\Gamma(\Delta)$ sits above or below the linear approximation $x\Gamma'(x)$ at $x = \Delta$, which can be interpreted as a cumulative measure of curvature.

Figure 4: Examples of Elasticity Dynamics

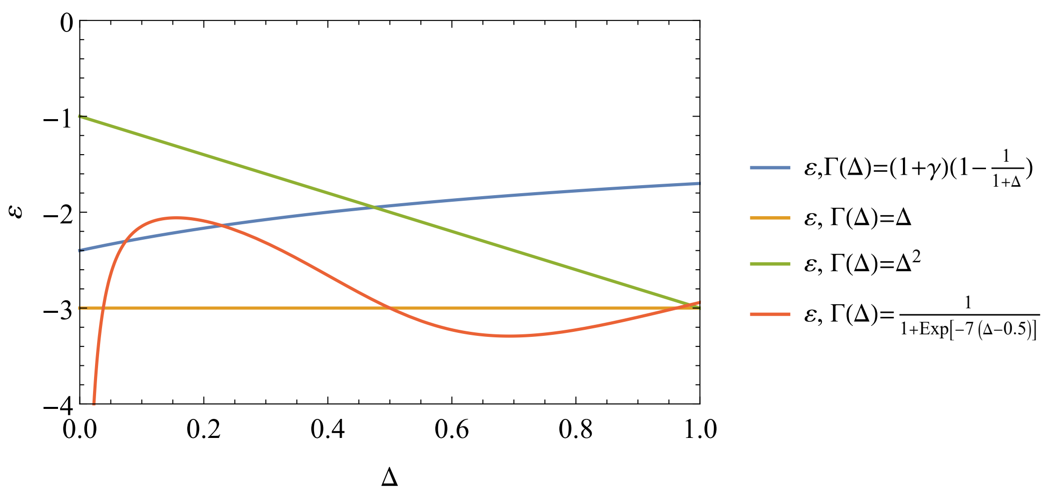

Figure 4 plots the behavior of price elasticities as a function of the time horizon Δ, for different specifications of function $\Gamma(\cdot)$. The curves represent choices of, respectively, a concave, linear, convex and S-shaped specifications. Clearly, the possibilities for the behavior of elasticities over the time horizon are richer than what was analyzed in Figure 3. As the time horizon increases, price elasticities from introducing promotions can increase, decrease, remain stable, or behave non-monotonically.

To clarify this property further, consider the case of an S-shaped $\Gamma(\cdot)$ function, which corresponds to low-type consumers buying the product at a faster rate immediately after the promotion is introduced, and then decelerating over its course. The bottom part of Figure 5 shows how the corresponding elasticity evolves when Γ is an S-shaped function $\Gamma(\Delta) = \frac{1}{1+\exp(-7(\Delta-0.5))}$ (i.e., the last case already presented in Figure 4). The top part of Figure 5 presents the actual function $\Gamma(\Delta)$ in gray, together with the two linear functions that cross the origin and are tangential to $\Gamma(\Delta)$.

Figure 5: Gamma function and Elasticity

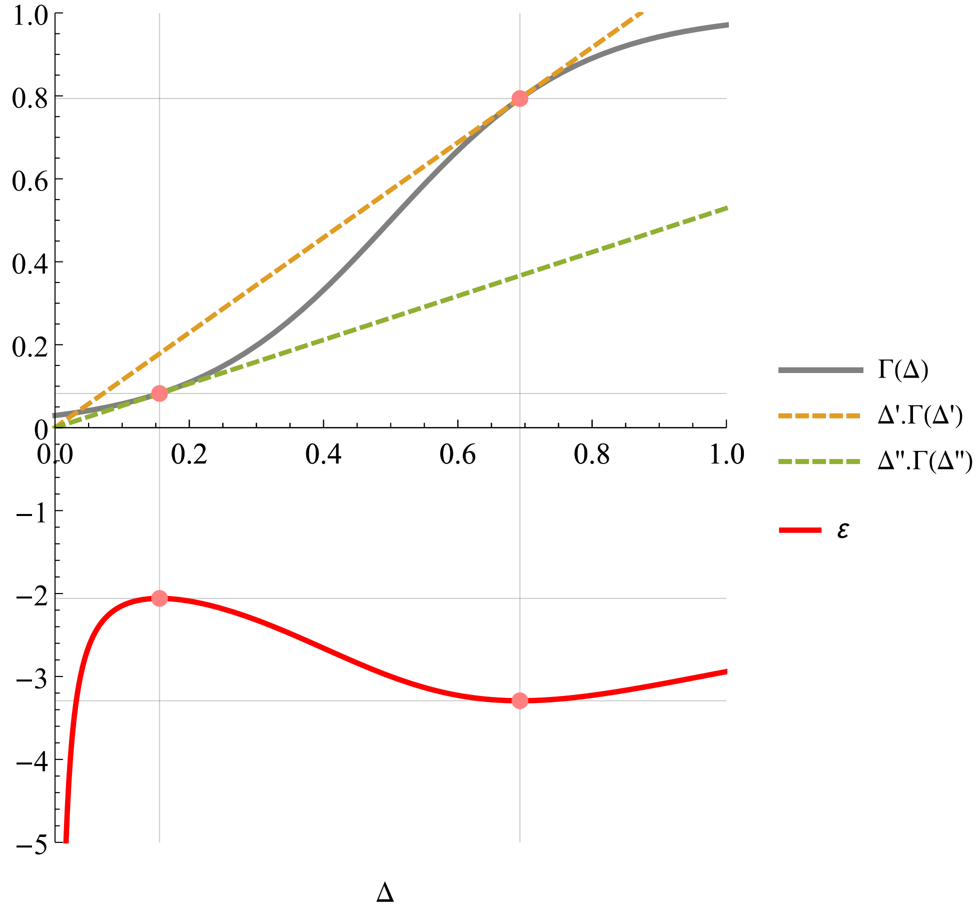

The tangency points of the top panel correspond to inflection points of the elasticity of the promotion along its duration. When function $\Gamma(\Delta)$ is cumulatively convex at Δ, the effect of introducing the promotion is such that low-type consumers bought more than if price had remained high, in which case only high types would have bought, at rate Δ. Similarly, when $\Gamma(\Delta)$ is cumulatively concave, the cumulative sales to low-types as the result of the price promotion lag behind the purchases of the high types, who would have bought at the alternative price.

The non-monotonic nature of the elasticity of a promotion along its duration is relatively counterintuitive, since in stable demand systems one typically expects that elasticity decreases with the time horizon. What this analysis shows is that what matters is the total accumulation of the purchase behaviors of the consumers who had been waiting for a promotion, vis-a-vis the consumers who are willing to buy. Finally, the S-shaped function considered above follows a logit specification (i.e., it is the c.d.f. of the logistic distribution), and may occur in real-world settings. For example, when consumers learn about the beginning of a new promotion via word of mouth, their purchases may follow an S-shaped diffusion pattern. Alternatively, low-type consumers may differ in their opportunity cost of time, with more eager consumers taking advantage of the promotion earlier and others preferring to postpone their purchase due to other activities that also compete for their time.

The results above are useful for empirical analyses especially given that function $\Gamma(\cdot)$ is observable directly from transactional data. For example, observing promotional sales with marked curvature means researchers should be aware that the elasticity estimate they obtain will depend on the duration of the promotional intervention. When the goal is to recover the equilibrium elasticity, then it may be essential to mimic the rhythm and duration of promotions observed in available transactional data. On the other hand, if promotional sales tend to be constant over time, then empirical researchers should primarily focus on other concerns (e.g., statistical power) when deciding the duration of an intervention, since the duration of the experimental promotional period is unlikely to affect their price sensitivity estimates.

Elasticity from concluding a promotion. We now consider the behavior of demand elasticity from ending a promotion at some time t, and charging the regular price for a period Δ:

Proposition 4. *The price sensitivity observable from concluding a price promotion increases in the time since its introduction if and only if* $(\Gamma'(t-t_1^*+\Delta)-\Gamma'(t-t_1^*)) > 0$*, and decreases otherwise. The price sensitivity increases with the time horizon of the regular price if and only if* $\Gamma(t-t_1^*+\Delta)-\Delta\Gamma'(t-t_1^*+\Delta)-\Gamma(t-t_1^*) < 0$ *and decreases otherwise. The signs of the comparative statics of the elasticities are given by:*

$$\frac{\partial \varepsilon_{p\uparrow}(t,\Delta)}{\partial t} \propto -\left(\Gamma'(t-t_1^*+\Delta)-\Gamma'(t-t_1^*)\right) \tag{22}$$

$$\frac{\partial \varepsilon_{p\uparrow}(t,\Delta)}{\partial \Delta} \propto \Gamma(t-t_1^*+\Delta)-\Gamma(t-t_1^*)-\Delta\Gamma'(t-t_1^*+\Delta) \tag{23}$$

Expression $t-t_1^*$ appears in both comparative statics expressions, and represents the duration of the ongoing price promotion that ended at time t with the reintroduction of the regular price. So, the relevant counterfactual to charging the regular price is to slightly extend the ongoing promotion. The sign of the first comparative static depends on a difference of slopes, that is, on the concavity/convexity of function $\Gamma(\cdot)$. For example, if function $\Gamma(\cdot)$ is convex over the period of the new regular price regime (in which case $\Gamma'(t-t_1^*+\Delta)-\Gamma'(t-t_1^*) > 0$), the effect of extending the promotion slightly on demand elasticity is negative (i.e., price sensitivity increases) because of fewer foregone sales to low-type consumers through the counterfactual extension to the promotion. Notice that function $\Gamma(\cdot)$ may be concave and convex over the span of the reestablished regular price, but the linear approximation to a very brief extension of the promotion duration depends only the slopes of $\Gamma(\cdot)$ at the endpoints.

As for the second comparative static, it measures the effect of the time horizon of the elasticity induced by reestablishing the regular price. We obtain a similar cumulative curvature measure as before, but now price elasticity increases (price sensitivity decreases) when the sales to low-types are cumulatively convex over the period of the regular price (note that $\Gamma(t - t_1^* + \Delta) - \Gamma(t - t_1^*)$ represents the fraction of low-type consumers who would have bought at the promotional price, which is relevant for the calculation of the elasticity even if they do not buy at the new price). As the time horizon of the elasticity is extended, the price elasticity increases when $\Gamma(\cdot)$ is cumulatively convex, because sales to low types would have grown faster than to high types in the foregone promotion scenario. In other words, as the time horizon grows, the sales reduction of increasing the price is largest, comparatively speaking, had the promotion been maintained.

Based on the results of the previous propositions, we draw the following conclusion about the behavior of price elasticities along their time horizons:

Corollary 1. *The demand elasticities obtained from initiating or concluding a price promotion can fluctuate up and down an arbitrary number of times over their time horizons.*

The result above implies that as one prolongs the duration of the new price regime, the induced demand elasticity will not necessarily converge to a stable value over its time horizon, as it may oscillate for as long as low-type consumers buy. While this result does not imply that one may observe infinite oscillations in demand elasticities, it is nonetheless relevant for researchers attempting to recover robust empirical estimates through experimental interventions.

Implications for experimentation. In the analysis above, we define price experimentation as analyst-driven interventions that introduce exogenous price changes to measure demand responsiveness. Our results indicate that analysts aiming to recover representative elasticity estimates must account for the dynamic nature of demand responses. For instance, assuming a constant elasticity demand function – common in both Economics and Marketing literatures – can lead to biased promotion effect estimates if the duration of the promotion is not properly accounted for (see the literature review in Bray, Stamatopoulos, and Sanders (2024), for example).

Consider an analyst who introduces a price promotion into a seller's pricing schedule and has access to a control group – perhaps through randomizing the promotion across stores or individuals. If the analyst is unaware of the underlying data-generating process, s/he may unknowingly obtain elasticity estimates that are highly time-sensitive. This is a realistic risk, as logistical constraints often prevent experiments from being aligned with the timing of equilibrium market activity. We illustrate that these effects can be economically significant in Section 6.

Focusing first on the promotion timing (while keeping the duration aligned with the market equilibrium), consider the case of an experimental promotion introduced during the regular-price period, at time $t^{'} \neq t_1^*$. By observing the change in sales through Δ periods, the analyst will observe elasticity $\widehat{\varepsilon}_{p\downarrow}\left(t^{'}, \Delta\right)$ rather than the true elasticity of the promotion, elasticity $\varepsilon_{p\downarrow}\left(t_1^*, \Delta\right)$, where $t^{'} < t_1^*$. The results laid out in this section imply the following conclusion:

Proposition 5. *Ceteris paribus, the introduction of an experiment earlier (later) than the market's equilibrium timing leads to an underestimate (overestimate) of price sensitivity.*

Hypothetically, in a world where a product's regular price may last for three or four weeks on average, introducing the experimental promotion earlier than usual (say in week 2) will produce a higher estimate of price elasticity of demand (i.e., lower price sensitivity). This result is unambiguous, and can be economically as we show in Section 6.

Now consider the case of introducing a promotion at the correct time, but with a duration different than the market's natural equilibrium. The implication for the analysis is the following:

Proposition 6. *Let bias* $= \widehat{\varepsilon}_{\downarrow}^{Duration} - \varepsilon_{p\downarrow}$ *be the difference in elasticity estimates of a promotion running for time* $\Delta^{'}$ *and the equilibrium duration* Δ^**. The sign of the bias is, ceteris paribus, equal to the sign of* $\frac{\Gamma(\Delta^*)}{\Delta^*} - \frac{\Gamma(\Delta^{'})}{\Delta^{'}}$.

The result above summarizes the effect the potential bias a naive analyst may incur for failing to take into account that elasticities evolve dynamically. The sign of the bias depends on how the per-time-unit experimental sales compare to per-time-unit sales in the market's equilibrium promotional regime. For example, suppose the promotional intervention is shorter than the equilibrium promotion duration, and that promotional sales evolve in a concave fashion. In this case, a short intervention will exhibit low per-time-unit sales, leading the researcher to obtain a positive bias, signifying an underestimation of the price sensitivity.

Beyond experimental design, dynamic promotion cycles can also create challenges when analyzing observational data. In particular, we next examine how price dynamics during promotions may affect the use of instrumental variables, even when prices are exogenous.

5 Implications for Empirical Analysis

In the analysis of transactional data, analysts often correct for the potential presence of price endogeneity. Indeed, failing to do so can generate well-understood biases in focal estimates. We consider the less well-understood case where an analyst proceeds to control for price endogeneity despite its absence. For example, the analyst could have introduced enough flexible controls to subsume the influence of omitted variables; or, faced with uncertainty about the exogeneity of price variation in the data, the analyst may prefer to 'be conservative' and control for the potential presence of endogeneity via instrumental variables. As before, we assume that the data generation process is the promotion cycles model we have developed, where price variation is indeed exogenous. In this case, the analyst would be correcting for endogeneity despite its absence.

The spurious use of instrumental variables (IV) is typically irrelevant, since at most one expects a loss of power from introducing a valid instrument to the analysis. However, in the case of price promotions, relying on lagged prices as instruments introduces two out-of-phase observations in each promotion cycle: whenever a promotion is begins or concludes, contemporaneous price jumps before lagged prices do. This effect introduces a bias in finite samples, which we investigate in this section. We also explain the conditions under which the bias disappears, and in the next section we investigate the extent to which this bias can be economically meaningful.

We focus on the typical linear specification used in two-stage least squares (2SLS) when using lagged prices as instruments, when the data generation process is the promotion cycles model. Our goal is to compare the OLS and IV estimates of the price parameter induced by the spurious IV misspecification. We also investigate the potential biases for non-linear models, such as the logit and constant-elastic demand.

Lagged prices as instruments. Consider regression equation

$$y_t = f\left(-\beta p_t\right) \tag{24}$$

where $\beta > 0$ is a measure of price sensitivity that an analyst would like to recover. If the analyst assumes a linear specification for $f\left(\cdot\right)$, she obtains the following sample estimates:

$$\widehat{\beta}_{OLS} = -\frac{Cov\left(y_t, p_t\right)}{Var\left(p_t\right)} \tag{25}$$

$$\widehat{\beta}_{IV} = -\frac{Cov\left(y_t, z_t\right)}{Cov\left(p_t, z_t\right)} \tag{26}$$

depending on whether ordinary least squares or IV estimators are employed. Both of these estimators are generally biased except when $f(\cdot)$ is linear. Our focus is not on documenting the bias introduced by the misspecification of $f(\cdot)$, which should be directly addressed by a more careful specification of function $f(\cdot)$ in the first place.[9] Rather, our goal is to characterize the difference

$$bias = \widehat{\beta}_{OLS} - \widehat{\beta}_{IV} \tag{27}$$

which the analyst may incorrectly attribute to valid endogeneity correction rather than misspecification. In our context, the case $bias > 0$ corresponds to the analyst interpreting that correcting for endogeneity led to predicting lower price sensitivity.

Consider the case of employing lagged prices as an instrument for contemporaneous ones. The rationale for this application is the assumption that demand is at most weakly correlated over time, but cost persistence or other factors may make lagged prices a powerful-enough instrument. We consider the case of an analyst who has access to one promotion cycle-worth of data generated by our model of promotion cycles, such that she observes sales y_t in each period, as well as the series of prices p_t. Let $n_1 = \frac{t_1^*}{\delta}$ and $n_2 = \frac{t_2^* - t_1^*}{\delta}$ be the number of periods during the regular and promotional price regimes, respectively, such that $\delta > 0$ is the length of each observation. We assume perfect divisibility throughout to simplify the analysis. Define the discrete-time sales as a sum of their continuous-time analogs, namely:

$$y^H := y_t, t \in \{1..n_1\} \tag{28}$$

$$y_t^L := y_{t-n_1}, t \in \{n_1 + 1, n_2\} \tag{29}$$

and

$$y_t = \int_t^{t+\delta} y'(\tau)\, d\tau \tag{30}$$

such that y^H is equal to the constant value of sales during the regular price phase, and y_t^L is equal to y_t for the promotional period in which sales generally vary, in line with the findings of the promotion cycles model. Finally, $y'(t)$ is the rate of change of sales. In this case, the spurious IV bias is given by

$$\widehat{\beta}_{OLS} - \widehat{\beta}_{IV} = -\frac{Cov(y_t, p_t)}{Var(p_t)} - \left(-\frac{Cov(y_t, p_{t-1})}{Cov(p_t, p_{t-1})}\right) \tag{31}$$

$$= \frac{Cov(y_t, p_{t-1})}{Cov(p_t, p_{t-1})} - \frac{Cov(y_t, p_t)}{Var(p_t)} \tag{32}$$

We start by calculating the denominators $Cov(p_t, p_{t-1})$ and $Var(p_t)$. Note that in the

[9] This is the focus of Erdem, Imai, and Keane (2003) and Hendel and Nevo (2006).

promotion cycles model, price p_t is equal to p_{t-1} except at the times when the promotion regime changes. In particular, $p_t = v_H$ until period $t = n_1$ and $p_t = v_L$ afterwards. Assuming $p_0 = v_L$ to balance the price cycle, it follows that:

$$Var\left(p_t\right) = \frac{1}{n}\sum_{t=1}^{n} p_t^2 - \left(\frac{1}{n}\sum_{t=1}^{n} p_t\right)^2 \tag{33}$$

$$= n_1 n_2 \left(\frac{v_H - v_L}{n}\right)^2 \tag{34}$$

and

$$Cov\left(p_t, p_{t-1}\right) = \frac{1}{n}\sum_{t=1}^{n} p_t p_{t-1} - \frac{1}{n^2}\sum_{t=1}^{n} p_t \sum_{t=1}^{n} p_{t-1} \tag{35}$$

$$= Var\left(p_t\right)\left(1 - \frac{1}{n_1} - \frac{1}{n_2}\right) \tag{36}$$

The last result shows that the serial covariance of prices can be written as the variance of price scaled by a function of the number of observations in both pricing phases. As additional observations of each pricing phase are gathered, the closer the two moments become.

Focusing now on the numerator expressions, the covariance of y_t and p_t is given by

$$Cov\left(y_t, p_t\right) = \frac{1}{n}\sum_{t=1}^{n} y_t p_t - \frac{1}{n}\sum_{t=1}^{n} y_t \overline{p} \tag{37}$$

where $\overline{p} = n_1 v_H + n_2 v_L$. We can rewrite equation (37) as

$$Cov\left(y_t, p_t\right) = \frac{1}{n}\left(\sum_{t=1}^{n_1} v_H y^H + v_L \sum_{t=1}^{n_2} y_t^L\right) - \frac{1}{n}\left(\sum_{t=1}^{n_1} y^H + \sum_{t=1}^{n_2} y_t^L\right)\overline{p} \tag{38}$$

As for the covariance between y_t and p_{t-1}, noting that $\sum_{t=1}^{n} p_t = \sum_{t=1}^{n} p_{t-1}$, we obtain:

$$Cov\left(y_t, p_{t-1}\right) = \frac{1}{n}\left(v_L y^H + \sum_{t=2}^{n_1} v_H y^H + v_H y_1^L + v_L \sum_{t=2}^{n_2} y_t^L\right) - \frac{1}{n}\left(\sum_{t=1}^{n_1} y^H + \sum_{t=1}^{n_2} y_t^L\right)\overline{p} \tag{39}$$

The new isolated terms (outside the sums) in equation (39) arise from the fact that price p_t and lagged price p_{t-1} differ at times $t = 1$ and $t = n_1 + 1$. Simplifying the expressions, we obtain the following result:

Proposition 7. *When using lagged prices as instruments, the spurious IV bias is given by:*

$$\widehat{\beta}_{OLS} - \widehat{\beta}_{IV} = \frac{n}{(n_1 n_2 - n)} \frac{y_1^L - \overline{y}^L}{v_H - v_L} \tag{40}$$

where $\overline{y}^L = \sum_{t=1}^{n_2} y_t^L$ *is the mean value of* y_t^L *during the promotional period.*

For the first factor of expression (40), $n \div (n_1 n_2 - n)$ to be positive, it suffices that $n_1, n_2 > 2$, which is likely to hold in most settings (we assume so for the discussion below). The second factor, $y_1^L - \overline{y}^L$, is the initial spike or drop in promotional sales relative to the average promotional sales. It is relatively surprising that this difference is an almost sufficient statistic for the sign of the bias. For example, when initial promotional sales are high comparatively with the rest of promotional sales, the bias from spurious instrumentation will be positive and consumers will appear less price sensitive after instrumentation.

Figure 6: Evolution and Mean Promotional Sales

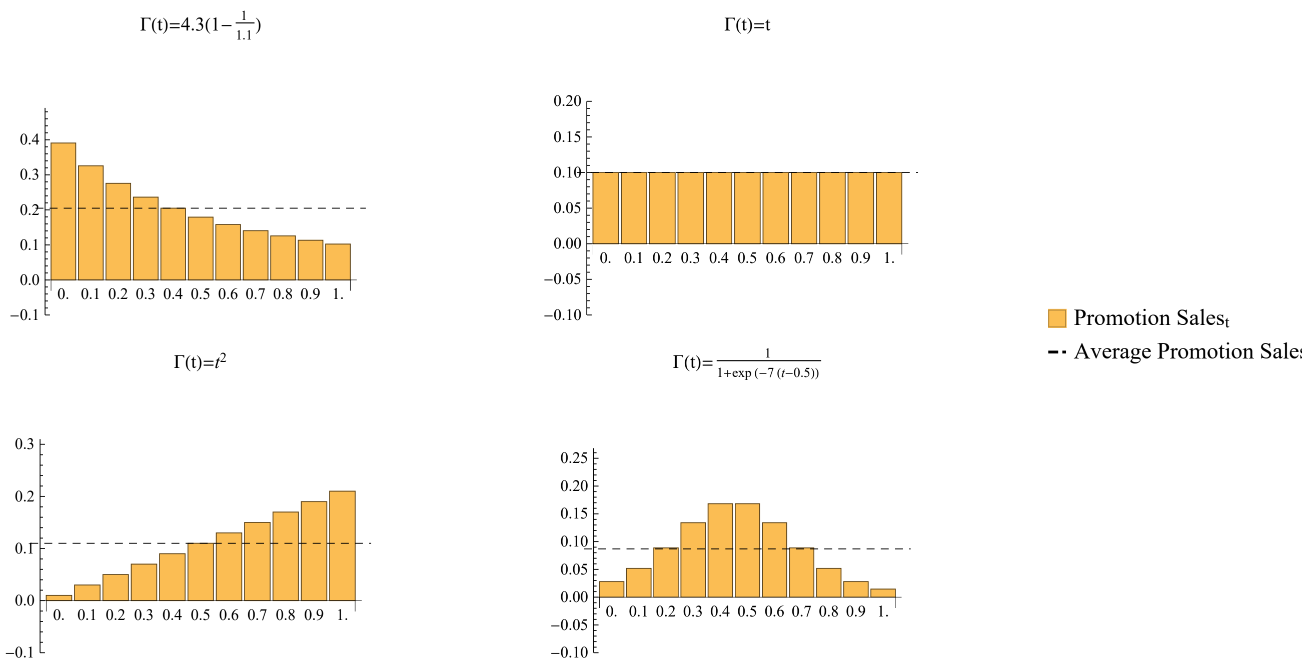

Note: The plots above denote hypothetical sales trajectories during promotional price phases.

Figure 6 illustrates promotional sales patterns for different specifications of function $\Gamma(\cdot)$. Only the top-left case, in which initial promotional sales are higher than the average promotion sales, exhibits a positive IV bias. The constant sales case (top right) would induce no bias, but it is unlikely to occur exactly in empirical settings. The cases in the bottom row both exhibit negative biases, such that the spurious IV bias may lead an analyst to believe that the endogeneity correction unveiled consumers that are more price sensitive than initially believed.

We summarize the results above in the following proposition:

Proposition 8. *When using lagged prices as instruments, the spurious IV bias $\widehat{\beta}_{OLS} - \widehat{\beta}_{IV}$ tends to zero as v_L approaches v_H and as the numbers of observations in both pricing regimes (n_1 and n_2) increase simultanously. When $n_1, n_2 > 2$, the bias is positive (negative) when the sales in the first period are higher (lower) than the mean promotional sales.*

The result above shows that the spurious IV bias introduced from misspecifying the price process in the first stage of the 2SLS procedure also depends on the difference of price levels as well as the number of observations. The fact that the bias depends on the price difference is intuitive and follows directly from the fact that the focal parameter β multiplies prices; this way, when the relative difference in prices changes, the absolute magnitude of the coefficient, and of the bias, varies accordingly. The result on the number of observations reveals that bias cannot be eliminated arbitrarily by increasing only the number of observations exclusively in one phase of the promotion cycle: the analyst needs to ensure that both periods (of regular and promotional price) feature enough observations to avoid introducing a significant bias from spurious instrumentation. Finally, the result on whether the sales function is increasing or decreasing is informative to empirical researchers. For example, when consumers are faster to take advantage of a promotion in its early days (e.g., $y_1^L > \overline{y}^L$), the misattribution bias may be positive, and in this case the IV estimator indicates a less price-sensitive demand than the linear estimator. The day on which the retailer introduces price promotions is also relevant: The fact that retailers tend to reset promotions near the beginning of the week, when there is less demand, may lead to a negative spurious IV bias. While promotional sales are unlikely to be strictly monotonic in empirical datasets, the bias depends only on the relationship between the first promotional sales and the full promotion effect, making it easy for analysts to determine the sign of the bias from direct data inspection.

The role of temporal resolution. In practical applications, analysts often have access to discrete-time data, which was generated in continuous time in the real world. Given the expression for the spurious IV bias, which specifically depends on the first-period promotion sales, it is possible that the sampling resolution available in datasets affects the results obtainable by the analyst.

Define $y(t)$ as the 'continuous-time analog' of y_1^L, such that first-period sales are given by

$$y_1^L = \int_{t_1^*}^{t_1^*+\delta} y'(t)\,dt \tag{41}$$

where δ represents the temporal window immediately preceding each time period t. Clearly, as δ approaches zero, so do the sales captured in a dataset's first promotional period. This is intuitive since as a dataset is cut in finer slices, fewer sales observations make it to each single data point. In contrast, average promotional sales $\overline{y}^L$ are not affected by data granularity. It is easy to show the behavior of the spurious IV bias $\widehat{\beta}_{OLS} - \widehat{\beta}_{IV}$ as the time resolution increases:

$$\begin{aligned}
\lim_{\delta\to 0} \widehat{\beta}_{OLS} - \widehat{\beta}_{IV} &= \lim_{\delta\to 0} \frac{n}{(n_1 n_2 - n)} \frac{y_1^L - \overline{y}^L}{v_H - v_L} \\
&= \lim_{\delta\to 0} \frac{\frac{t_2^*}{\delta}}{\underbrace{\frac{t_1^*}{\delta}\frac{t_2^*-t_1^*}{\delta} - \frac{t_2^*}{\delta}}_{\to 0}} \frac{\underbrace{y_1^L}_{\to 0} - \overline{y}^L}{v_H - v_L} \\
&= 0
\end{aligned}$$

As expected, when thinner slices of data are available, the granularity of the information makes the estimators approach each other. This result is intuitive since when many observations are used, lagged prices are equal to contemporaneous prices most of the time, leading the bias to reduce.

Now, consider the case when the sampling resolution is high (i.e., δ is low), but not infinite. In this case, the first factor of the bias (see expression 40) is positive but the second factor – related with the difference between the sales in the first period and the average promotional sales – is negative. Figure 7 illustrates the evolution of the two factors as the temporal resolution increases (i.e., δ approaches zero).

Figure 7: Gamma function and Elasticity

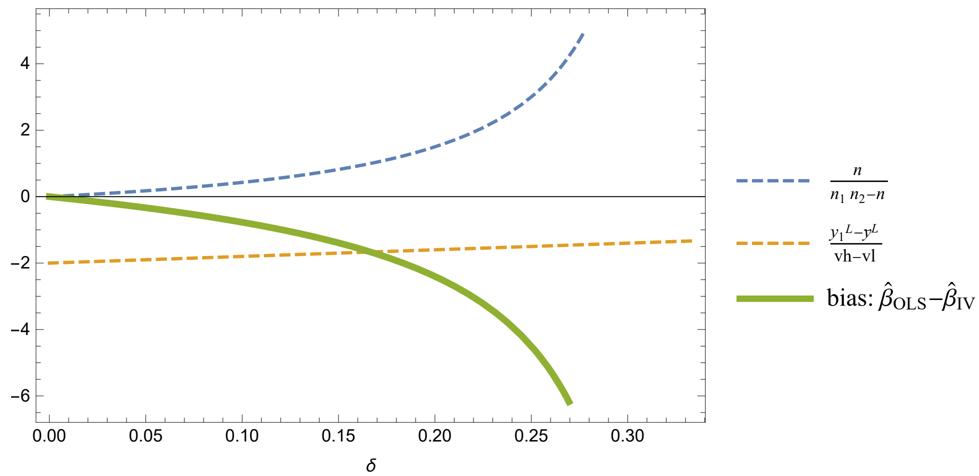

Parameter values and assumptions: $y_1^L = \delta \div (t_2^* - t_1^*)$,$\overline{y}^L = 1$, $v_H = 2$,$v_L = 1$, $t_1^* = 1$, and $t_2^* = 1.5$.

The blue dashed line in Figure 7 stems from the fact that as the sampling resolution decreases (i.e., δ increases), the number of observations decreases, leading to a partial positive bias related with the weight of mismatching observations in lagged prices on the sample. Simultaneously, a partial negative bias stemming from the difference in sales of the first period and the full promotion period takes place (dashed orange line) for any value of $\delta > 0$ (i.e., discrete time). This bias tends to zero as the number of periods increases (i.e., $\delta \downarrow$). Since the latter bias is always strictly negative when $\delta \simeq 0$, it is possible to obtain the following result, represented by the green line in the figure:

Proposition 9. *Keeping the number of observations constant, there always exists a high-enough sampling resolution (i.e., low δ) that induces a strictly negative spurious IV bias.*

In practical terms, this result implies that when the data-frequency is high, it is both true that the spurious IV bias will be low and that its sign will be negative.

Linearizable Models. The analyses above rely on interpreting y_t as sales and p_t as prices. In practical applications, analysts often transform their data to be able to estimate non-linear models via linear regression. Two classical examples are the logit and constant-elastic models. Let

$$y_t := \Psi\left(sales\left(t\right)\right) \tag{42}$$

$$p_t := \Upsilon\left(price\left(t\right)\right) \tag{43}$$

be the required transformations of sales and prices to linearize the models in question. For example, in the homogeneous logit model, a linear model is obtained from defining

$$\Psi\left(sales\left(t\right)\right)=\log\left(\frac{sales\left(t\right)}{M-sales\left(t\right)}\right) \tag{44}$$

$$\Upsilon\left(price\left(t\right)\right)=price\left(t\right) \tag{45}$$

where M is the potential market size either known or assumed by the analyst.[10] Similarly, one may estimate the constant-elastic model by employing linear regression on the transformed variables

$$\Psi\left(sales\left(t\right)\right)=\log\left(sales\left(t\right)\right) \tag{46}$$

$$\Upsilon\left(price\left(t\right)\right)=\log\left(price\left(t\right)\right) \tag{47}$$

As before, we are interested in investigating the spurious bias $\widehat{\beta}_{OLS}-\widehat{\beta}_{IV}$. The result below shows that it is possible to formalize the bias for a class of non-linear models that includes the logit and constant-elastic specifications:

Proposition 10. *When $\Psi\left(\cdot\right)$ and $\Upsilon\left(\cdot\right)$ are strictly increasing, then the spurious IV bias is given by*

$$\widehat{\beta}_{OLS}-\widehat{\beta}_{IV}=\frac{n}{\left(n_1n_2-n\right)}\frac{\Psi\left(sales\left(1\right)\right)-\frac{1}{n_2}\sum_{t=1}^{n_2}\Psi\left(sales\left(t\right)\right)}{\Upsilon\left(v_H\right)-\Upsilon\left(v_L\right)} \tag{48}$$

The result above shows that it is possible to calculate the spurious IV bias for a large class of models that are amenable to estimation via linear regression after data transformation.

Implications for empirical analysis. The Marketing field has long debated the need to control for price endogeneity in the context of retail. Our analysis broadens the discussion by contemplating the finite-sample bias introduced by correcting for endogeneity using lagged-prices as instruments in the context of price promotions. The bias arises because, in each promotion cycle, two observations are systematically out of phase: when a promotion is introduced and when it is concluded, regular prices naturally jump up and down before lagged prices. Regardless of the possible presence of endogeneity, the occurrence of discrete jumps may introduce an unexpected spurious instrumentation bias. While lagged prices

[10] The well-known BLP model (see Berry (1994), Berry, Levinsohn, and Pakes (1995)) also fits into this analysis, as it applies instrumental variables in a linear form by inverting market shares to obtain linear mean utilities.

continue to be useful instruments to correct for demand endogeneity, our analysis suggests that researchers will benefit from analyzing the patterns of promotional prices and correct the obtained estimates accordingly to prevent side effects when correcting for endogeneity.

In the era of big data, it is tempting to assume that in most applications analysts possess enough observations to minimize the bias. However, this is not the case when the large number of observations occurs in the cross-section dimension rather than the temporal one. For example, a dataset comprising many consumers' data allows one to reduce the uncertainty about each period's sales, but each consumer's observations feature the same number of price and lagged-prices mismatches, i.e., cases where $p_{it} \neq p_{it-1}$. So, increasing the number of consumers reduces uncertainty about each period sales, but does not necessarily contribute to reduce the bias we investigate. One way to alleviate the problem is to ensure that the sampling frequency or that the number observations in each phase of the promotion cycle are sufficiently high.

6 Empirical Assessment

We assess the economic relevance of some of the challenges analyzed above by analyzing data from promotion experiments in Chile in 2013 that tracked the purchases of 234,063 loyalty program cardholders, described in Elberg, Gardete, Macera, and Noton (2019). In their analyses, Elberg, Gardete, Macera, and Noton (2019) exogenously manipulate price promotions of 170 products across 17 categories, typically introducing promotions to pairs of sku's in each category every week, over a period of 10 weeks. In the first five weeks, promoted products are assigned 10% or 30% discounts depending on the store (control or treatment), and in the last five weeks discount depths are constant and equal to 10%.

Timing of introduction of price promotion. We first consider the effect of the timing of introduction of promotions. We take advantage of the fact that the timing of weekly promotions in the data were randomized across products. According to Proposition 3, introducing promotions later unambiguously increases price sensitivity. We focus on the first five weeks of sales of 140 sku's from categories in which price manipulation was validated by Elberg, Gardete, Macera, and Noton (2019).[11] We run regression

$$\log\left(y_{jct}\right) = \alpha_c + \beta \log\left(p_{jct}\right) + \lambda \text{first\_promo\_week}_j + \gamma \text{first\_promo\_week}_j \times \log\left(p_{jct}\right) + \eta_{jct} \tag{49}$$

[11]Limiting the time period of analysis to five weeks reduces potential contamination from dynamic effects, as we explain later.

where j stands for sku, c for its category, and t for promotion day. Coefficient α_c is a fixed effect for each product category, so that β and γ capture average static and dynamic price elasticity components. Table 1 shows the regression output.

Table 1: Effect of Promotion Timing on Price Elasticity

Parameters	Estimates
β	-0.344**
	(0.131)
λ	0.134**
	(0.027)
γ	-0.3**
	(0.033)
R^2	0.334
N:	1,773

Note: Standard errors in parentheses. Fixed effects introduced at the category level. ** - 1% significance level.

The results above predict that the average price elasticity starts at -0.644 ($-0.344 - 0.3$ for $t = 1$), and reduces further by -0.3 each week that the promotion introduction is postponed, which is an economically significant effect. Products that were promoted for the first time in week 5 on average exhibit an elasticity of -1.844 for the same promotion length, which is much lower than -0.644 for the first week. The estimate of γ should be caveated by the fact that, by virtue of the data intervention, sales of promoting a product in a given week are affected by having promoted a rival product in the previous week. However, this feature of the data should reduce promotion effectiveness of products whose promotions were launched later in the intervention, and so the estimate of γ is likely a higher bound for the true effect of the timing of promotions. The analysis above supports the idea that the timing at which promotions are introduced can significantly affect elasticity estimates.

Duration of promotion. The promotion durations in Elberg, Gardete, Macera, and Noton (2019) respect the seller's promotion cycles, from Tuesdays through Mondays. We can nonetheless investigate the effect of failing to account for the correct promotion durations by estimating price elasticities on promotional windows shorter than one week. We consider the following regression specifications:

$$\log(y_{jct}) = \alpha_{ck} + \beta_k \log(p_{jct}) + \eta_{jct},\ t \leq k$$

where each regression indexed by k utilizes only the observations occurring in the first k days to estimate demand elasticity. So, $k = 1$ considers the elasticity β_k obtained by promoting products for a single day, and so on, while the observations in which the product sold a the regular price are kept constant.

Table 2: Effect of Promotion Timing on Price Elasticity

Parameters	Estimates	Parameters	Estimates
β_1	-1.289**	β_5	-1.04**
	(0.275)		(0.13)
β_2	-1.092**	β_6	-1.023**
	(0.199)		(0.121)
β_3	-1.079**	β_7	-1.02**
	(2.676)		(0.111)
β_4	-1.067**		
	(0.143)		
N:	244 through 1,773		

Note: Standard errors in parentheses. Fixed effects introduced at the category level. ** - 1% significance level.

Table 2 reveals that the estimated price sensitivity reduces over time. This result is in line with intuition that price sensitivity tends to lower with the time horizon of price elasticity. The evolution of β_k is consistent with $\Gamma(\cdot)$ being cumulatively concave, and the differences across the estimates are economically significant.

These findings underscore the importance of accounting for the natural cadence of promotions when designing price experiments. Because price sensitivity evolves over the course of the promotion cycle, experiments that fail to align with the market's typical timing can yield elasticity estimates that do not reflect the market's typical consumer responsiveness. In settings where promotions serve to dynamically sort consumers, the observed elasticity depends critically on how many consumers have accumulated in the waiting pool at the time of intervention. This highlights the value of analyzing historical promotional rhythms before implementing experimental designs.

Spurious IV Bias. We now assess the real-world magnitudes of the spurious IV bias. We focus on two sku's that, due to logistical constraints, were promoted in isolation in their category.[12] We document the sales of the selected sku's, in the Cold Cuts and Milk categories, during three 10%-discount promotions. For each of the products, we calculate the spurious

[12]There exist six additional products in the dataset satisfying these conditions, but they exhibit lower sales volumes and sometimes sell zero units on a given day. The sku's we present are bought by a large number of customers, each observation exhibit low noise.

Table 3: Spurious IV Bias

Sku	n_1	n_2	y_1^L	$\overline{y}^L$	v_H	v_L	$\widehat{\varepsilon}$	bias	$\widehat{\varepsilon}_{IV}$	bias (FE)
Cold Cuts sku	411	21	8	46.8	5.36	5.08	-0.86	-1.73	-2.59	-2.26
Milk sku	411	21	29	51.2	1.42	1.31	-1.29	-0.37	-1.66	-0.21

Note: Above, elasticity (column $\widehat{\varepsilon}$) obtained from employing a constant-elastic regression of daily sales across price conditions. $\overline{y}^R$ denotes the average daily sales at the regular price.

IV bias to illustrate the extent to which it can be meaningful in real world applications.

Figure 8 presents three sales curves along the promotional cycle. The three promotions were implemented in different stores, and customers who ever bought in more than one store were removed from the analysis. The figure also presents the sales average across the three promotional campaigns for each sku.

The cold cuts sku (top chart) exhibits lower sales at the beginning and end of the promotional cycle, achieving its peak near the weekend. The milk sku also exhibits lower sales on the first day of promotion, but overall steadier sales than the cold cuts category. The fact that initial sales are lower than the average allows us to assess that the spurious IV bias is negative in both cases, that is, instrumentation will lead to an amplification of the estimated price coefficient. Finally, note that the sales patterns are likely to strongly depend on the day of the week the seller decides to launch price promotions. Had the seller decided to launch price promotions every Saturday, for example, it is likely that initial promotional sales would surpass the average and flip the sign of the spurious IV bias. This is relevant since retailers tend to introduce promotions near the beginning of the week, at which point sales are relatively lower than on weekends.

Table 3 presents the spurious IV bias for each case. For each selected sku, we include 411 days (across multiple stores) of sales made at the regular price.

The daily price elasticities of the sku's are both near -1 – as obtained from constant-elastic regression specifications at the daily level – and the magnitudes of the spurious IV biases are economically significant in both cases, with the estimated elasticities being affected noticeably despite the relatively high number of days at which the products are sold at the regular price. An inspection of Figure 8 would have immediately informed researchers of the signs of the biases, but the results in Table 3 are informative in terms of quantitative impact: the IV bias tripled the estimated elasticity in the case of the Cold Cuts sku, and amplified the price elasticity of the Milk sku by approximately 29%.

The biases obtained in the analysis above may be due to the market's gradual response to promotions as well as simply from the regular weekly sales cycle (e.g., most consumers buy on the weekend regardless of promotional activity). While the biases we document are correct

Figure 8: Evolution and Mean Promotional Sales From Data

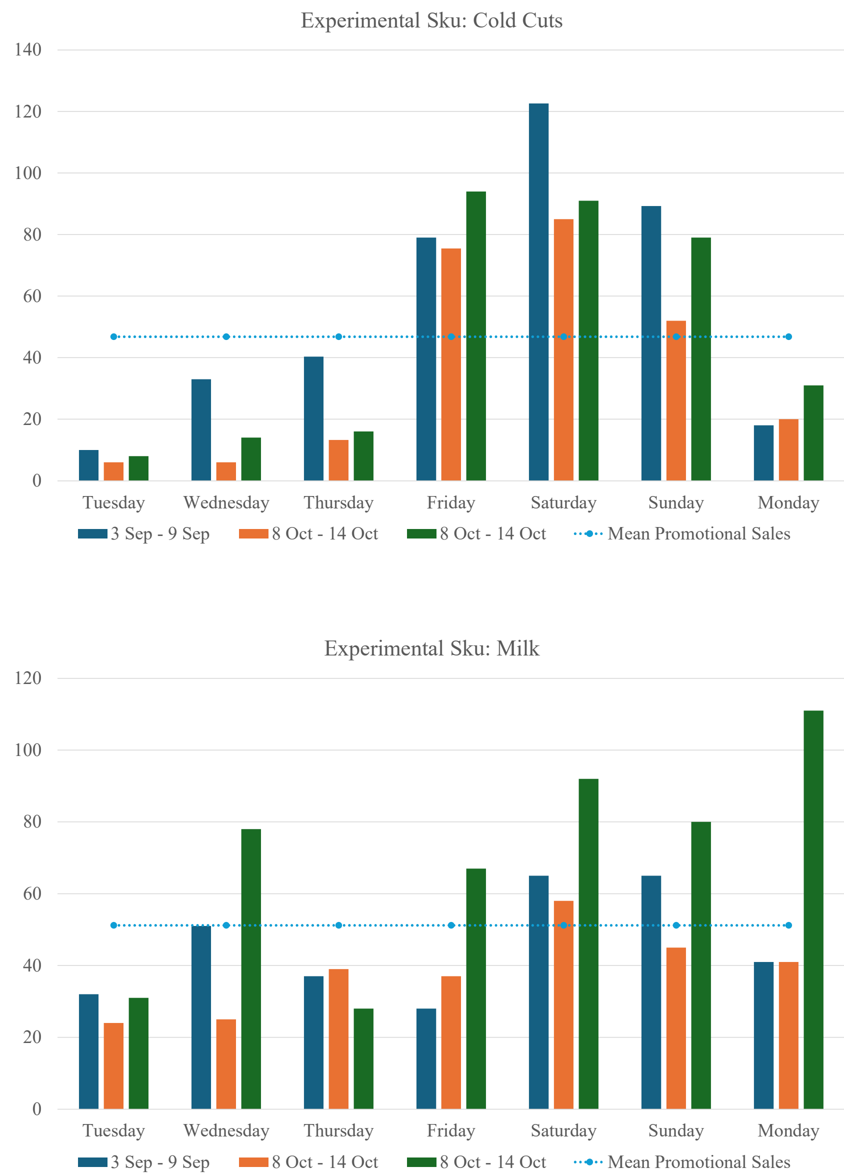

Note: Above, seller's promotional cycles start on Tuesdays and end on Mondays. Bars represent sales in three stores, at different times, at a 10% price promotion. The blue dotted line depicts the overall average of promotional sales across.

regardless of the underlying mechanism, it is useful to note that researchers can control for the second mechanism (e.g., within-week sales patterns) by introducing day-of-week fixed effects in the first-stage IV procedure. By controlling for within-week effects, the remaining IV bias will be due to the market's natural response to promotions. We recalculate the bias by transforming the data according to

$$y_t^{\prime L} = y_t^L - y_t^H + \overline{y}^R \tag{50}$$

where y_t^L are the promotional sales on day t, y_t^H are the sales at the regular price on the same day (at a different store), and $\overline{y}^R$ captures the average sales at the regular price across days. We apply the same transformation to prices, in the spirit of fixed effects procedures. The FE-robust bias figures are presented on the last column of Table 3. The biases remain economically significant, with the specific bias of the Cold Cuts sku increasing further. This means that regardless of the day of the week the retailer sets its promotion schedule, natural sales responses induce an economically significant spurious IV bias on estimates.

The results on spurious IV bias emphasize that researchers should exercise caution when applying lagged prices as instruments in promotional settings. The discrete nature of price jumps systematically creates misalignment between current and lagged prices at the start and end of promotions, introducing a finite-sample bias even when prices are exogenous. Importantly, this bias does not diminish with larger cross-sectional datasets, but rather depends on the number of observed promotion cycles and the temporal resolution of the data. In many practical applications, where relatively short promotions occur over the study period, this bias can be economically meaningful.

7 Conclusion

The goal of this paper is to investigate the implications to empirical analyses of settings in which price elasticities are dynamic as a result of price promotions. The theoretical model we develop serves as a benchmark to derive those implications, both for the case of experimental interventions and for the analysis of transactional data. In terms of best practices, we have found that analysts will do best by first considering the patterns of promotional sales available in their data. Whenever promotional sales are not constant over time, additional care is needed to ensure that both experimental and observational analyses correctly deliver the statistics of interest to researchers. In the case of experimental methods, it is possible to analyze ex ante the sign of bias induced by not matching the rhythm of equilibrium promotions during an experimental intervention. In the case of observational methods, care may be needed to ensure that instrumentation does not, by itself, bias the estimates of interest.

Any researcher versed in experimental methods knows that the goal of exogenous interventions is to facilitate a ceteris paribus environment, in which all but one factor or interest are held constant. It is therefore not surprising that measurement biases may be introduced when the equilibrium promotion timings are not matched during experimental interventions. It is, however, extremely challenging to always do so. For example, Elberg, Gardete, Macera, and Noton (2019) introduce price experiments across 17 product categories in 10 stores of a major retailer in Chile. Their analyses emphasize the comparability of treatment and control subjects of experimental price promotions, as well as the matching between the experimental promotion depths and the ones practiced by the retailer. However, the timing of the experimental promotion interventions is held constant across categories, which is likely to introduce category-specific biases depending on the equilibrium promotion frequency in each one. Similarly, Bray, Stamatopoulos, and Sanders (2024) compare the elasticities of transactional and experimental price promotions conducted at a total of 116 stores of a large midwestern grocery retailer. Like Elberg, Gardete, Macera, and Noton (2019), the duration of their transactional promotions (median price duration of 26 days) does not match the median duration of the experimental ones (median price duration of 7 days, cfr. Table 1). In markets with dynamic price elasticities, it is possible that these mismatches introduce biases in the results.

The results we explore also show that simple A/B tests may fail to capture meaningful estimates of price sensitivity. For example, a market featuring promotion cycles generates price variation at very specific points in time. Without multiple experimental interventions that vary promotion duration or an explicit dynamic model, it is impossible to reliably extrapolate price sensitivity estimates to a large set of counterfactual scenarios. Future contributions analyzing paths of promotional sales and elasticities may shed light on the magnitude of some of the challenges investigated in our analysis. Finally, while we show that elasticity may be dynamic as a result of a single seller's promotional activity, it would also be interesting to investigate whether competition in a market with several products brings new dynamics and challenges to empirical analysis.

8 Appendix

8.1 Proofs

Proposition 1.

We use the implicit differentiation of the first order condition w.r.t. t_1 to derive the results:

$$\mathrm{d}(foc) = \frac{\partial foc}{\partial t_1} dt_1 + \frac{\partial foc}{\partial x} dx \bigg|_{t_1 = t_1^*} = 0 \tag{51}$$

The expression above is used to to determine the sign of each comparative static of t_1^* w.r.t. variable x. The specific results in the proposition follow from the following expressions:

- $\frac{dt_1^*}{dk} = -\frac{1}{v_L(\gamma_1 + t_1^*)W''(t_1^*)} > 0$, since $W''(t_1^*) < 0$.
- $\frac{dt_1^*}{dv_L} = \frac{\gamma_1(1+\mu) + W(t_1^*) - (\gamma_1 + t_1^*)W'(t_1^*)}{v_L(\gamma_1 + t_1^*)W''(t_1^*)} < 0$, since the denominator is negative and numerator, when added to the first-order condition divided by v_L (which is equal to zero in equilibrium) equals a positive number. Specifically,

$$\begin{aligned} &\gamma_1(1+\mu) + W(t_1^*) - (\gamma_1 + t_1^*)W'(t_1^*) + \frac{foc(t_1^*)}{v_L} \\ =&\frac{\gamma_1 v_H + k}{v_L} > 0. \end{aligned}$$

- $\frac{dt_1^*}{dv_H} = -\frac{\gamma_1}{v_L(\gamma_1 + t_1^*)W''(t_1^*)} > 0$, since $W''(t_1^*) < 0$.
- $\frac{dt_1^*}{d\gamma_1} = -\frac{v_L W'(t_1^*) + v_H - v_L(1+\mu)}{v_L(\gamma_1 + t_1^*)W''(t_1^*)} > 0$, since $W''(t_1^*) < 0$ and $v_H > (1+\mu)\, v_L$. The last expression is necessary for an interior solution in t_1, since otherwise the seller would be better off selling always to all consumer types at price v_L.
- $\frac{dt_1^*}{d\mu} = \frac{\gamma_1}{(\gamma_1 + t_1^*)W''(t_1^*)} < 0$, since $W''(t_1^*) < 0$.

Proposition 2.

Directly obtained by replacement of functions $W(t) = \omega - \exp(-t)$ and $\Gamma(\Delta) = (1+\lambda)\left(1 - \frac{1}{1+\Delta}\right)$ in the first-order condition.

Proposition 3.

In the main text we show that:

$$\varepsilon_{p\downarrow}(t,\Delta) = -\frac{v_H}{v_H - v_L}\frac{W(t)\,\Gamma(\Delta) + \mu\Delta}{\Delta}, t \in (0, t_1^*) \tag{52}$$

The comparative statics that inform Proposition 3 are:

$$\frac{\partial \varepsilon_{p\downarrow}(t,\Delta)}{\partial t} = -\frac{v_H}{v_H - v_L}\frac{W'(t)\Gamma(\Delta)}{\Delta} < 0 \tag{53}$$

$$\frac{\partial \varepsilon_{p\downarrow}(t,\Delta)}{\partial \Delta} = \frac{v_H}{v_H - v_L}\frac{W(t)\,(\Gamma(\Delta) - \Delta\Gamma'(\Delta))}{\Delta^2} \tag{54}$$

$$\propto \Gamma(\Delta) - \Delta\Gamma'(\Delta) \tag{55}$$

When a promotion is introduced slightly later, at time, the effect on the elasticity is strictly negative (i.e., price sensitivity increases). As for the effect of the duration of the promotion, it depends on the sign of $\Gamma(\Delta) - \Delta\Gamma'(\Delta)$, as discussed in the main text.

Proposition 4.

In the main text we show that:

$$\varepsilon_{p\uparrow}(t,\Delta) = -\frac{v_L}{v_H - v_L}\left(1 - \frac{\Delta}{W(t_1^*)\,(\Gamma(t - t_1^* + \Delta) - \Gamma(t - t_1^*)) + (1+\mu)\,\Delta}\right), t \in (t_1^*, t_2^*) \tag{56}$$

The comparative statics that inform Proposition 4 are:

$$\frac{\partial \varepsilon_{p\uparrow}(t,\Delta)}{\partial t} = -\frac{v_L}{v_H - v_L}\frac{\Delta W(t_1^*)\,(\Gamma'(\Delta + t - t_1^*) - \Gamma'(t - t_1^*))}{[(1+\mu)\Delta + W(t_1^*)(\Gamma(\Delta + t - t_1^*) - \Gamma(t - t_1^*))]^2} \tag{57}$$

$$\propto -\left(\Gamma'(\Delta + t - t_1^*) - \Gamma'(t - t_1^*)\right) \tag{58}$$

$$\frac{\partial \varepsilon_{p\uparrow}(t,\Delta)}{\partial \Delta} = -\frac{v_L}{v_H - v_L}\frac{W(t_1^*)\,[\Delta\Gamma'(t - t_1^* + \Delta) - \Gamma(t - t_1^* + \Delta) + \Gamma(t - t_1^*)]}{[\Delta(1+\mu) + W(t_1^*)(\Gamma(t - t_1^* + \Delta) - \Gamma(t - t_1^*))]^2} \tag{59}$$

$$\propto \Gamma(t - t_1^* + \Delta) - \Gamma(t - t_1^*) - \Delta\Gamma'(t - t_1^* + \Delta)$$

The rationale of these results is the same as in Proposition 3.

Proposition 5.

Follows directly from a finite-difference analog of Proposition 3.

Proposition 6.

Follows directly from a finite-difference analog of Proposition 3. Specifically,

$$\varepsilon_{p\downarrow}\left(t,\Delta^{'}\right)-\varepsilon_{p\downarrow}\left(t,\Delta^{*}\right)=\frac{v_H}{v_H-v_L}\frac{W'(t)\left[\Delta^{'}\Gamma(\Delta^{*})-\Delta^{*}\Gamma(\Delta^{'})\right]}{\Delta^{*}\Delta^{'}} \tag{60}$$

$$\propto\frac{\Gamma(\Delta^{*})}{\Delta^{*}}-\frac{\Gamma(\Delta^{'})}{\Delta^{'}} \tag{61}$$

Proposition 7.

Proved in the main text. The final steps are given by:

$$\widehat{\beta}_{OLS}-\widehat{\beta}_{IV}=\frac{Cov\left(y_t,p_{t-1}\right)}{Cov\left(p_t,p_{t-1}\right)}-\frac{Cov\left(y_t,p_t\right)}{Var\left(p_t\right)} \tag{62}$$

$$=\frac{n(n_2y_1^L-\sum_{t=1}^{n_2}y_t^L)}{n_2(n_1n_2-n)(v_H-v_L)} \tag{63}$$

$$=\frac{n}{(n_1n_2-n)}\frac{y_1^L-\overline{y}^L}{v_H-v_L} \tag{64}$$

Proposition 8.

Obtained directly from expression (64).

Proposition 9.

Proved in the main text.

Proposition 10.

Follows from Proposition 7, replacing y_t, p_t, and p_{t-1} by the corresponding transformations.

References

Ackerberg, D., and M. Rysman (2002): "Unobserved product differentiation in discrete choice models: Estimating price elasticities and welfare effects," .

Anderson, E. T., and E. J. Fox (2019): "How price promotions work: A review of practice and theory," *Handbook of the Economics of Marketing*, 1, 497–552.

Bell, D. R., and J. M. Lattin (1998): "Shopping behavior and consumer preference for store price format: Why "large basket" shoppers prefer EDLP," *Marketing Science*, 17(1), 66–88.

Bergemann, D., and J. Välimäki (2006): "Dynamic pricing of new experience goods," *Journal of Political Economy*, 114(4), 713–743.

Berry, S., J. Levinsohn, and A. Pakes (1995): "Automobile Prices in Market Equilibrium," *Econometrica*, 63(4), 841–890.

Berry, S. T. (1994): "Estimating discrete-choice models of product differentiation," *The RAND Journal of Economics*, pp. 242–262.

Bijmolt, T. H., H. J. Van Heerde, and R. G. Pieters (2005): "New empirical generalizations on the determinants of price elasticity," *Journal of marketing research*, 42(2), 141–156.

Bray, R., I. Stamatopoulos, and R. E. Sanders (2024): "Observational price variation in scanner data cannot reproduce experimental price elasticities," *Working Paper*.

Conlisk, J., E. Gerstner, and J. Sobel (1984): "Cyclic pricing by a durable goods monopolist," *The Quarterly Journal of Economics*, 99(3), 489–505.

Datta, H., H. J. van Heerde, M. G. Dekimpe, and J.-B. E. Steenkamp (2022): "Cross-national differences in market response: line-length, price, and distribution elasticities in 14 Indo-Pacific Rim economies," *Journal of Marketing Research*, 59(2), 251–270.

Dhar, T., J.-P. Chavas, and B. W. Gould (2003): "An empirical assessment of endogeneity issues in demand analysis for differentiated products," *American Journal of Agricultural Economics*, 85(3), 605–617.

Elberg, A., P. M. Gardete, R. Macera, and C. Noton (2019): "Dynamic effects of price promotions: Field evidence, consumer search, and supply-side implications," *Quantitative Marketing and Economics*, 17, 1–58.

Erdem, T., S. Imai, and M. P. Keane (2003): "Brand and quantity choice dynamics under price uncertainty," *Quantitative Marketing and economics*, 1, 5–64.

Fibich, G., A. Gavious, and O. Lowengart (2005): "The dynamics of price elasticity of demand in the presence of reference price effects," *Journal of the Academy of Marketing Science*, 33(1), 66–78.

Freimer, M., and D. Horsky (2012): "Periodic advertising pulsing in a competitive market," *Marketing Science*, 31(4), 637–648.

Guadagni, P. M., and J. D. Little (1983): "A logit model of brand choice calibrated on scanner data," *Marketing science*, 2(3), 203–238.

Hendel, I., and A. Nevo (2006): "Measuring the implications of sales and consumer inventory behavior," *Econometrica*, 74(6), 1637–1673.

Hitsch, G. J., A. Hortacsu, and X. Lin (2021): "Prices and promotions in US retail markets," *Quantitative Marketing and Economics*, 19(3), 289–368.

Kamakura, W. A., and G. J. Russell (1989): "A probabilistic choice model for market segmentation and elasticity structure," *Journal of marketing research*, 26(4), 379–390.

Kim, K. i., and A. Petrin (2015): "Tests for price endogeneity in differentiated product models," *Journal of Econometric Methods*, 4(1), 47–69.

Lloyd, T. A., C. W. Morgan, S. McCorriston, and E. Zgovu (2009): "Do Sales Matter? An Exploration of Price Discounting in UK Food Retailing," Discussion paper, International Association of Agricultural Economists.

Mela, C. F., S. Gupta, and D. R. Lehmann (1997): "The long-term impact of promotion and advertising on consumer brand choice," *Journal of Marketing research*, 34(2), 248–261.

Parker, P. M. (1992): "Price elasticity dynamics over the adoption life cycle," *Journal of Marketing Research*, 29(3), 358–367.

Rao, V. R. (2009): *Handbook of pricing research in marketing*. Edward Elgar Publishing.

Simon, H. (1979): "Dynamics of price elasticity and brand life cycles: An empirical study," *Journal of Marketing Research*, 16(4), 439–452.

Tellis, G. J. (1988): "The price elasticity of selective demand: A meta-analysis of econometric models of sales," *Journal of marketing research*, 25(4), 331–341.

Villas-Boas, J. M., and R. S. Winer (1999): "Endogeneity in brand choice models," *Management science*, 45(10), 1324–1338.

Villas-Boas, S. B., and J. M. Villas-Boas (2008): "Learning, forgetting, and sales," *Management Science*, 54(11), 1951–1960.