Selling a Cheaper Mousetrap: Wal-Mart's Effect on Retail Prices

Emek Basker*
University of Missouri

First Draft: December 2003
This Draft: March 2005

Abstract

I quantify the price effect of a low-cost entrant on retail prices using a case-study approach. I consider the effect of Wal-Mart entry on average city-level prices of various consumer goods by exploiting variation in the timing of store entry. The analysis combines two unique data sets, one containing opening dates of all US Wal-Mart stores and the other containing average quarterly retail prices of several narrowly-defined commonly-purchased goods over the period 1982-2002. I focus on 10 specific items likely to be sold at Wal-Mart stores and analyze their price dynamics in 165 US cities before and after Wal-Mart entry. An instrumental-variables specification corrects for measurement error in Wal-Mart entry dates. I find robust price effects for several products, including shampoo, toothpaste, and laundry detergent; magnitudes vary by product and specification, but generally range from 1.5-3% in the short run and four times as much in the long-run.

JEL Numbers: L130, L810, E310

Keywords: Wal-Mart, Competition, Prices, Market Size

---

*Comments welcome to: emek@missouri.edu. A previous version of this paper was presented at Illinois State University, UC-Berkeley, University of Nevada - Las Vegas, Federal Trade Commission, the 2004 Missouri Economic Conference and the 2004 International Industrial Organization Conference (Chicago). I thank Saku Aura, Emin Dinlersoz, Bengte Evenson, Luke Froeb, Eric Gould, Saul Lach, Mark Lewis, Dave Mandy, Mike McCracken, Chuck Moul, Peter Mueser, Mike Noel, David Parsley, Dan Rich, John Stevens, Chad Syverson, Ken Troske, Pham Hoang Van, and seminar participants for helpful conversations and comments. I also thank Sean McNamara and David Parsley for sharing ACCRA data with me, and Aliya Kalimbetova and Loren Risker for research assistance. All remaining errors are my own.

"Wal-Mart's mania for selling goods at rock-bottom prices has trained consumers to expect deep discounts everywhere they shop, forcing competing retailers to follow suit or fall behind."
– *Washington Post*, November 6, 2003

"Even if you don't shop at Wal-Mart, the retail powerhouse increasingly is dictating your product choices – and what you pay – as its relentless price cutting helps keep inflation low."
– *USA Today*, January 29, 2003

1 Introduction

In most models of imperfect competition, entry of a lower-cost competitor reduces output prices. The effect is larger the smaller the initial number of firms and the higher are cross-price elasticities of demand. In this paper, I quantify the price effect of a low-cost entrant on retail markets using a case study of Wal-Mart entry, and show that Wal-Mart's price impact can be quite large. The analysis combines data on the opening dates of all US Wal-Mart stores with average city-level retail prices of several narrowly-defined commonly-purchased goods over the period 1982-2002. I focus on 10 specific items likely to be sold at Wal-Mart stores and analyze their price dynamics in 165 US cities before and after Wal-Mart entry. I find price declines of 1.5%-3% for many products in the short run, with the largest price effects occurring for aspirin, laundry detergent, toothpaste and shampoo. Long-run price declines tend to be much larger, and in some specifications range from 7-13%. These effects are driven mostly by relatively small cities, which have high ratios of retail establishments to population.

Wal-Mart's low labor costs and the retail chain's logistics and distribution innovations make it the prototypical low-cost entrant. Broadly, there are two mechanisms by which Wal-Mart's expansion could have affected retail prices and consumer inflation rates: an aggregate mechanism and a market-specific mechanism. The aggregate mechanism works through Wal-Mart's interactions with both suppliers (manufacturers and importers) and other large retail chains. This mechanism can lower prices in communities not served by Wal-Mart if it leads to lower costs for other retailers.[1] The market-specific mechanism works through competition (and

[1] The argument for this mechanism is as follows. By demanding lower prices from suppliers, Wal-Mart forces manufacturers to cut costs, possibly by relocating overseas. Competing retail chains (notably Target, but also many smaller chains) also increase efficiency by emulating Wal-Mart's innovations in logistics and distribution (McKinsey Global Institute [22]). The result is lower prices in chain stores across the country, some in locations that have no Wal-Mart stores.

possibly learning) at the local level.

The focus of this paper is on the second mechanism. Wal-Mart's entry into a given market (city or town) can lower prices by increasing the competitive pressure incumbents (and future entrants) face. This is the prediction of most standard imperfect-competition models, such as differentiated-product Bertrand competition and a spatial-competition model, and also of many models with equilibrium price dispersion (such as Reinganum [25]).[2]

Although these theories make consistent predictions about the price impact of entry, very little empirical work has been done to quantify these effects. I test these predictions on average prices of 10 specific goods such as toothpaste, Coke, and jeans, by exploiting exogenous variation in the timing of store entry in different markets. I combine two unique data sources on Wal-Mart store locations and retail prices in 165 US cities over a 20-year period, 1982-2002. The Wal-Mart data include store locations and opening dates of all U.S. Wal-Mart stores. Price data from the American Chamber of Commerce Research Association (ACCRA) consist of average retail prices of 10 products across multiple establishments in each city.

The methodology follows Basker [3], which examines the employment effects of Wal-Mart entry, with two innovations. In Basker [3], I consider the effect of Wal-Mart entry on *county-level* employment in the retail and wholesale sectors, using 1749 counties (slightly more than half of all US counties). Because price data are available at the city (or town) level, rather than the county level, I disaggregate the Wal-Mart data to the city level for this study. In addition, because price data are collected quarterly, I perform the analysis using quarterly, rather than annual, data.

I define the "effect" of Wal-Mart entry broadly. For example, if Wal-Mart entry induces the exit of an incumbent drugstore, the long-run price effect I isolate combines the effect of both the entry and the exit. By estimating separate short- and long-run price effects, I attempt to separate these issues. If Wal-Mart entry spurs other entries or leads to increased differentiation among incumbents, this too is incorporated in the net effect.

I find that the price effect of Wal-Mart entry differs by product and city size. For several products, including toothpaste, shampoo, aspirin, and laundry detergent, Wal-Mart entry reduces average retail prices by an economically large and statistically significant 7-13% in the long run. These results have real implications: if the market basket of low-income shoppers

[2] The products whose prices I consider are homogeneous, but the shopping experience, including location and service, may be better described as differentiated.

declined by this full amount, the income effects could be very large. Results are driven by Wal-Mart's effects in smaller cities, with more retail establishments per capita, where retail establishments tend to be smaller and the retail environment is less competitive.

The remainder of this paper is organized as follows. Section 2 provides background information on Wal-Mart. Section 3 describes the data. The empirical methodology is outlined in Section 4, and results are presented in Section 5. Section 6 concludes with a discussion of the magnitude of the price effect and its implications.

2 Wal-Mart Background

The first Wal-Mart store opened in Rogers, Arkansas in 1962. By the time the company became publicly traded in 1969 it had 18 stores throughout Arkansas, Missouri, and Oklahoma. Wal-Mart slowly expanded its geographical reach, building new stores and accompanying distribution centers further and further away from its original location, and continued, at the same time, to build new stores in areas already serviced. Figure 1 shows maps of the 48 contiguous states with locations of Wal-Mart stores over time (in 1982, 1993 and 2002) to illustrate this point. In the beginning of the period, Wal-Mart's sales accounted for approximately 0.2% of all US retail sales; by the end of it, Wal-Mart had approximately 2900 stores in all 50 states and about 800,000 employees in the United States, and accounted for approximately 5% of all US retail sales.

Wal-Mart is extremely efficient even compared with other "big-box" retailers. Lehman Brothers analysts have noted Wal-Mart's "leading logistics and information competencies" (Feiner [12]). The *Financial Times* has called Wal-Mart "an operation whose efficiency is the envy of the world's storekeepers" (Edgecliffe-Johnson [11]). Wal-Mart's competitive edge is driven by a combination of conventional cost-cutting and sensitivity to demand conditions and by superior logistics and distribution systems. The chain's most-cited advantages over small retailers are economies of scale and access to capital markets, whereas against other large retail chains the most commonly cited factor is superior logistics, distribution, and inventory control.[3] Employing exclusively non-union workers may be another source of cost-savings relative to other retailers.[4]

[3]Details on Wal-Mart's operations can be found in Harvard Business School's three case studies about Wal-Mart (Ghemawat [15], Foley and Mahmood [13], and Ghemawat and Friedman [16]).

[4]On the other hand, Wal-Mart is said to match the union wage in markets where it competes directly with

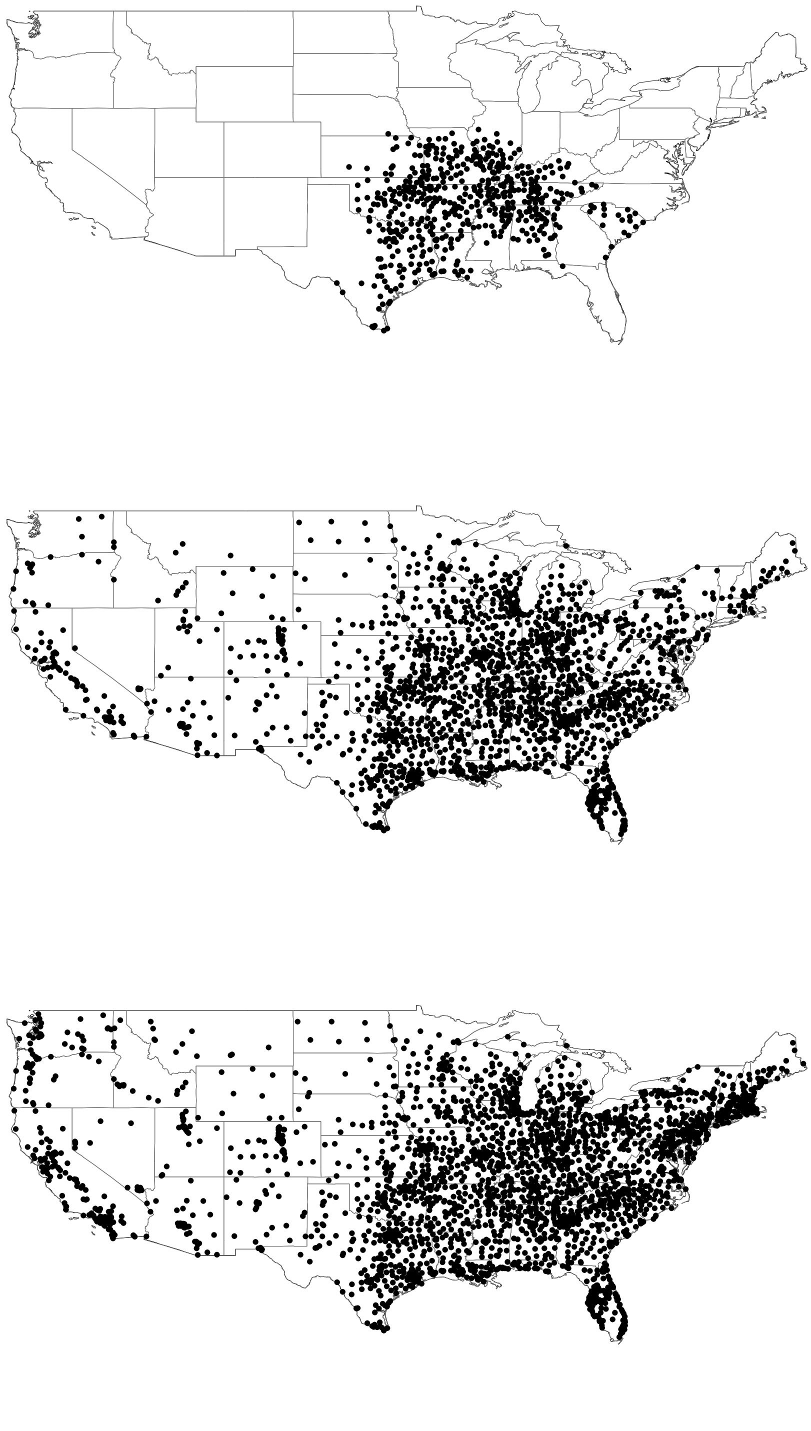

Figure 1: Locations of Wal-Mart Stores, 1982, 1993, 2002

3 Data

To assess Wal-Mart's effect on retail prices, I combine data on the locations and opening dates of Wal-Mart stores with retail price data.

3.1 Retail Prices

Retail prices are obtained from the American Chamber of Commerce Research Association. I use quarterly prices of 10 products in 165 cities over 21 years, from 1982-2002.

The American Chamber of Commerce Research Association (ACCRA), through local Cambers of Commerce, surveys 5-10 retail establishments in the first week of each quarter in participating cities. Participating cities vary from quarter to quarter, with some cities moving in and out of the sample frequently, while others are included more regularly; 250-300 cities are surveyed each quarter during the sample period. Of these, I selected the 165 cities for which data from at least 50 quarters were available (including both current price and lagged price). Each quarter between 100-140 cities are in the sample.

The prices collected by ACCRA cover approximately 50 goods and services. From the list of items, I selected 10 goods that were homogeneous or nearly so, and likely to be sold at most Wal-Mart stores (I exclude groceries and alcoholic beverages because most Wal-Mart stores over the sample period do not include a grocery section). The selected products are listed in Table 1. I begin my sample period in 1982 because most of these products were introduced into the ACCRA price list in that year.[5]

While these data are the best available local price data that span such a long time-series as well as cross-section of cities, covering specific prices, they have some drawbacks. One problem is that the individual stores from which prices are collected cannot be identified; if they include Wal-Mart stores, we will not be able to distinguish the direct impact of Wal-Mart (e.g., due to charging lower prices than other stores) from the indirect impact, through competitive pressure, on other stores. Somewhat mitigating this concern, the ACCRA manual urges its price collectors nationwide to "[s]elect only grocery stores and apparel stores where professional and managerial households normally shop. Even if discount stores are a majority of your overall market, they

unionized retailers (Saporito [26]).

[5] For some products, there is a change in the brand or/and size of the product during the time period, but since quarter dummies will be controlled in the analyses below, changes in the price due to such shifts are accounted for. We assume that the price response to Wal-Mart entry is the same across these alternatives.

shouldn't be in your sample at all unless upper-income professionals and executives really shop there" (ACCRA [1], p. 1.3). If price surveyors follow this instruction, we may estimate a *lower bound* of Wal-Mart's impact, because Wal-Mart and its most direct competitors are unlikely to be included.

Another potential problem is measurement error in prices. Measurement error may come from several sources. The most likely are errors made by price surveyors in the tallying of individual prices (e.g., copying prices incorrectly from store shelves); the average price reported by ACCRA will be sensitive to such errors. This type of error is likely to be idiosyncratic, uncorrelated over time or across products. Differences across cities or surveyors in the types of establishments surveyed or the size of the geographic area covered may also be treated as errors; these are less innocuous, because they are likely to be correlated over time within a city. I discuss this issue and the problems it creates, along with possible remedies, in the methodology section.

3.2 Wal-Mart Stores

The Wal-Mart store data are described in detail in Appendix 1 of Basker [3]. I briefly review the data sources.

I collected data on the locations and opening dates of 2,382 Wal-Mart stores in the United States from Wal-Mart annual reports, Wal-Mart editions of Rand McNally Road Atlases and annual editions of the *Directory of Discount Department Stores*.[6] The available data include store location (by city) and store number. Opening years of individual stores can be inferred by comparing lists of existing stores from consecutive years.[7] Table 2 (reproduced from Basker [3]) summarizes these sources.[8]

From these lists, I create the indicator variable $\mathbf{WMopen_{jt}}$, which takes on the value of 1 if the directories and store lists suggest that a Wal-Mart store exists in city j in quarter t, and zero otherwise. While this assignment is simple in the calendar years preceding or following the year of entry, the exact quarter of entry is not known from these lists, so quarterly values

[6]Requests for store-opening data from Wal-Mart corporation were denied. I have also tried to obtain opening dates from local newspapers, but only succeeded in a handful of cases.

[7]Vance and Scott [30] list store entries to 1969, the year the company became publicly traded.

[8]For the years 1990-1993, in which no satisfactory store list allowing identification of Wal-Mart stores in a city exists (the *Directory of Discount Department Stores* was not updated during those years), I assign opening dates to stores according to a probabilistic algorithm that uses information on the number of stores opening in each *state* each year. State-level data are obtained from Wal-Mart annual reports [31]. This procedure is detailed in Basker [3], Appendix 1.

of $\mathbf{WMopen_{jt}}$ have to be imputed for the year of entry. I obtain the empirical distribution of entry dates over the year from Wal-Mart's web site, which provides exact opening dates for all store openings over the three-year period July 2001 - June 2004. (Quarterly probabilities are approximately 29%, 26%, 26%, and 19%.) So, for a store that opened in 1986, I assign the variable **WMopen** a value of 0 up to and including the first quarter of 1986 (since the first-quarter observation gives prices in the first week of January), 0.29 in the second quarter, 0.55 in the third quarter, 0.81 in the fourth quarter, and 1 from the first quarter of 1987 onwards.[9]

I also construct a set of identifiers of Wal-Mart *planned* entry dates, using a combination of company-assigned store numbers (available from the Rand McNally atlases) and the net change in the number of stores each year (from company annual reports). Wal-Mart assigns store numbers roughly in sequential order, with store #1 opening first, followed by store #2, and so on; store numbers appear to be assigned early in the store planning process. Following this practice, I assign planned entry dates to stores sequentially, based on their store numbers, holding fixed (at actual levels) the total number of new stores to open each year. For example, since there were 18 stores at the end of 1969, and 20 new stores opened in 1970, I assign 20 stores (numbers 19-38) the planned entry year 1970. This assignment assumes that the number of planned entries in 1970 was the same as the actual number of entries that year.

Within a calendar year, I assign the first 29% of stores to the first quarter (so they are assumed to be open by the second quarter), the next 26% to the second quarter, another 26% to the third quarter, and the remainder to the fourth quarter (so they are open by the first quarter of the following year).

This assignment method provides a good approximation to the likely order in which the stores were planned. Aggregating these store-level entry dates to the city-quarter level, I construct the indicator variable $\mathbf{WMplan_{jt}}$, which equals 1 if city j would have had a Wal-Mart store in quarter t had the stores opened in the order in which they were planned. In the most specifications below, I use only the opening dates (both **WMopen** and **WMplan**) for the *first* Wal-Mart store to open in each city, but I also estimate two models that try to get at the effect of subsequent Wal-Mart entries.[10]

[9]The resulting variable $\mathbf{WMopen_{jt}}$ is equal to either 0 or 1 in 96% of the observations. Most of the observations in which $\mathbf{WMopen_{jt}}$ takes on values strictly between 0 and 1, representing the probability that a Wal-Mart store exists, occur during the years (1990-1993) when store directories were not updated (see footnote 8). Replacing this imputation with an indicator variable which equals 1 if the probability of a Wal-Mart store exceeds 50%, and zero otherwise, does not meaningfully change any of the results in the paper.

[10]An alternative to imputing quarterly data for Wal-Mart is to aggregate the price data to annual frequency.

Table 1: Products

	Product	Description	Price*
Drugstore Products	Aspirin	100-tablet bottle, Bayer brand (to 1994:3)	$ 4.96
		0.5oz Polysporin ointment (from 1994:4)	$ 4.56
	Cigarettes	Carton, Winston, king-size (85mm.)	$19.69
	Coke	2-liter Coca Cola, excluding deposit	$ 1.56
	Detergent	49oz Tide/Bold/Cheer laundry detergent (to 1991:3)	$ 3.39
		42oz Tide/Bold/Cheer laundry detergent (1991:4-1996:3)	$ 3.97
		60oz Cascade dish washing powder (from 1996:4)	$ 3.02
	Kleenex	200-count Kleenex tissues (to 1983:4)	$ 1.46
		175-count Kleenex tissues (from 1984:1)	$ 1.34
	Shampoo	11oz bottle, Johnsons Baby Shampoo (to 1991:2)	$ 4.21
		15oz Alberto VO5 (from 1991:3)	$ 1.29
	Toothpaste	6-7oz Crest or Colgate	$ 2.40
Clothing	Shirt	Mans dress shirt, Arrow	$30.19
	Pants	Levis 501/505 jeans, rinsed/washed/bleached size 28-36 (to 1999:4)	$33.28
		Mens Dockers "no wrinkle" khakis size 28-36 (from 2000:1)	$35.72
	Underwear	3 boys cotton briefs, Fruit of the Loom (to 1983:4)	$ 7.29
		3 boys cotton briefs, size 10-14, cheapest brand (from 1984:1)	$ 5.36

* Average price over entire sample period in 2000 dollars

Table 2: Directory Sources for Wal-Mart Opening Dates

Years	Source
1962-1969	Vance and Scott (1994)
1970-1978	Wal-Mart Annual Reports
1979-1982	*Directory of Discount Department Stores*
1983-1986	*Directory of Discount Stores*
1987-1989	*Directory of Discount Department Stores*
1990-1993	Imputed (see footnote 8)
1994-1997	*Rand McNally Road Atlas*

3.3 Sample Cities

The sample cities, determined by price-data availability, are shown in Figure 2, and some summary statistics are shown in Table 3. The average city in the sample had approximately 200,000 residents in 2000 (the median city had approximately half as many residents). The large apparent decrease in the number of establishments between 1987 and 1997 is due to a change in Census industrial coding from SIC to NAICS. (More information about the switch from SIC to NAICS is available at `http://www.census.gov/epcd/www/naics.html`.[11]) The right-skew implied by deviation between median and mean is comparable to the one obtained from a census of US cities.

Of the 165 sample cities, 25 had at least one Wal-Mart store at the beginning of the sample period. All but three cities (Denver, South Bend, and Tacoma) had experienced entry by the end of the sample period, and three cities experienced exit.

4 Empirical Methodology

4.1 Ordinary Least Squares (OLS) Regressions

I estimate the following regression by product:

$$\begin{aligned}\mathrm{p}_{kjt} = \alpha_k + \beta_k \mathrm{p}_{kj,t-1} + \theta_k \mathrm{WMopen}_{jt} \\ + \sum_j \gamma_{kj} \mathrm{city}_j + \sum_t \delta_{kt} \mathrm{quarter}_t + \sum_j \tau_{kj} \mathrm{trend}_t + \varepsilon_{kjt} \quad (1)\end{aligned}$$

where $\mathbf{p_{kjt}}$ is the natural log of the price of product k in city j in quarter t, $\mathbf{quarter_t}$ is a quarter indicator (where t ranges from 1982q2 to 2002q4 – a total of 83 quarter indicators), $\mathbf{city_j}$ is a city indicator, $\mathbf{trend_t}$ is a linear trend (with coefficient τ_{kj}, which is specific to city j) and $\mathbf{WMopen_{jt}}$ is the Wal-Mart indicator: it equals 1 if city j has a Wal-Mart store in quarter t. (From now on, I suppress the k subscript; it is always implied, since the regressions are estimated one product at a time.)

The quarter fixed effects are intended to capture macroeconomics price fluctuations, changes in product definitions or cost of production that are common across all cities. City fixed effects

This was done in a previous version of the paper, with qualitatively similar results.

[11]Wal-Mart entry does not lead to sharp decreases in the number of retail establishments; on average at the county level, 5 retail establishments close within 5 years of Wal-Mart entry (Basker [3]).

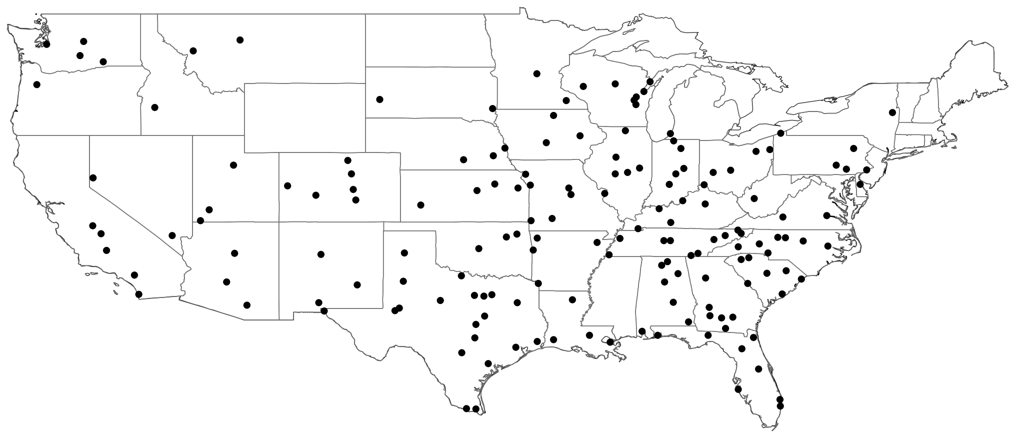

Figure 2: Locations of Sample Cities

Table 3: Sample Cities Summary Statistics

		Percentiles		
	Mean	25th	50th	75th
Population (1980)	164,322	40,450	71,626	172,140
Population (1990)	179,344	44,972	84,021	187,217
Population (2000)	199,797	52,894	91,565	217,074
Retail Establishments, SIC (1982)*	1,780	537	906	1,736
Retail Establishments, SIC (1987)**	2,132	641	1,069	2,069
Retail Establishments, NAICS (1997)†	1,037	373	515	1,106
Per-Capita Income (1979)‡	16,640	15,041	16,351	17,810
Per-Capita Income (1989)‡	17,911	16,052	17,346	19,592

Source: Author's calculations from *County and City Data Book*, Various Years

* Mean calculated using only 142 of 165 cities; percentiles include all cities

** Mean calculated using only 145 of 165 cities; percentiles include all cities

† Mean calculated using only 146 of 165 cities; percentiles include all cities

‡ Real 2000 dollars

capture average (long-run) price differences across cities, for example due to differential land or labor costs. City-specific time trends capture changes over time in relative prices for various exogenous reasons: trends in average income, income inequality, land and labor costs, and the level of competition.[12] If Wal-Mart entry is correlated with the sign and degree of these trends – for example, if Wal-Mart enters cities with declining prices earlier (or later) than cities with rising prices – the estimated coefficient will be biased in regressions that omit this trend. (City-level time trends are jointly extremely statistically significant; F statistics are almost always above 10,000.) Put together, these control variables capture three things: average differences in prices across cities; different trends in prices across cities; and macroeconomic shocks, as well as changes in product definitions, that shift these trends.

The coefficient on lagged price, β, is a relative "stickiness" parameter, indicating to what extent relative prices in city j in quarter t – adjusting for differences in their long-run trends – are unchanged from one quarter to the next. If $\beta = 0$, then the price (of shampoo or cigarettes or laundry detergent) in city j in quarter t is an i.i.d. draw, with expected value that depends only on city j's price trend, and any aggregate factors captured by the calendar quarter (and possibly the presence of Wal-Mart). If $\beta > 0$, however, deviations of price from this deterministic path last more than just one quarter, so that last period's price of a product has predictive power for this period's price. The interpretation of the coefficients on the city fixed effects and the city-specific trends depends critically on having $\beta < 1$, i.e., on the price series being stationary. My estimates of β vary by product and specification, but mostly lie in the interval $[0.7, 0.8]$.[13] Thus, β captures the fact that price deviations from a city's "trending average" (relative to all other cities in the sample) last more than one period, but decay over time.

The coefficient θ captures the *instantaneous* effect of Wal-Mart entry on prices. Assuming, for expositional purposes, that $\theta < 0$ and $\beta < 1$, if prices fall by θ in the quarter of entry, there will be an additional effect of $\beta\theta$ the following quarter, as lagged prices are now θ lower than they were before; and an additional $\beta^2\theta$ the following quarter. The long-run (asymptotic) effect of Wal-Mart entry on prices is $\frac{\theta}{1-\beta}$.[14] I estimate both the instantaneous and asymptotic

[12]The inclusion of both city fixed effects and city-specific trends allows the average price in each city to take a separate (linear) path. This is a reduced-form way of controlling for variables that change at low frequencies (like income and population) without including them directly.

[13]Parsley and Wei [23] use ACCRA price data for 48 cities over the period 1975-1994 and find strong evidence against unit roots in prices, and in favor of price convergence.

[14]As with the city-level coefficients, this expression depends critically on $\beta < 1$. As note above, this condition

effects, and separately test for their significance. (I use a Wald test for the non-linear test.)

Figure 3 illustrates the interpretation of the coefficients. Wal-Mart entry occurs at time t_0. The first panel shows the case with no price trends; the second and third panels show that the interpretation of the coefficients does not change when city-specific linear time trends are included in the model.

Because the short- and long-run effects estimated by these models are mechanically related (through the coefficient β), I also test to see whether there is an additional long-run effect of Wal-Mart. This could occur for several reasons: as some incumbents exit following Wal-Mart's entry, the remaining stores' prices may increase over time; alternatively, they may decrease if consumers are relatively unresponsive to price differences across locations in the short-run, but become more responsive over time. I allow a two-step adjustment process to Wal-Mart entry: an instantaneous effect θ, and a second effect ϕ five years (20 quarters) later:[15]

$$\begin{aligned} \mathrm{p}_{jt} = \alpha + \beta \mathrm{p}_{j,t-1} + \theta \mathrm{WMopen}_{jt} + \phi \mathrm{WMopen}_{j,t-20} \\ + \sum_j \gamma_j \mathrm{city}_j + \sum_t \delta_t \mathrm{quarter}_t + \sum_j \tau_j \mathrm{trend}_{tj} + \varepsilon_{jt} \end{aligned} \tag{2}$$

Figure 4 shows this specification graphically, with θ and ϕ "price shocks" T periods apart. (For simplicity, trends are omitted from the figures.) The long-run effect is given by $\frac{\theta+\phi\cdot\beta^{T+1}}{1-\beta}$. Depending on the relative magnitudes of θ and ϕ, and on the sharpness of "discounting" implied by β, the long-run effect may be larger or smaller than the short-run effect of Wal-Mart, and may even be of the opposite sign (as in the second panel). This is in contrast with estimates from Equation (1) in which, as long as $\beta < 1$, long-run estimates are always of the same sign as short-run effects, and always larger in absolute terms.

4.2 Measurement Error in Prices

ACCRA price data are subject to measurement error, as discussed in Section 3.1. If price were only on the left-hand side of the regression, measurement error would increase the variance of the error term without affecting any coefficient estimates. But both specifications (1) and (2) include lagged price on the right-hand side of the regression, to account for the fact that prices

holds for all specifications.

[15]The choice of a five-year window is arbitrary, but Basker [3] shows that within five years the retail sector has completed its adjustment.

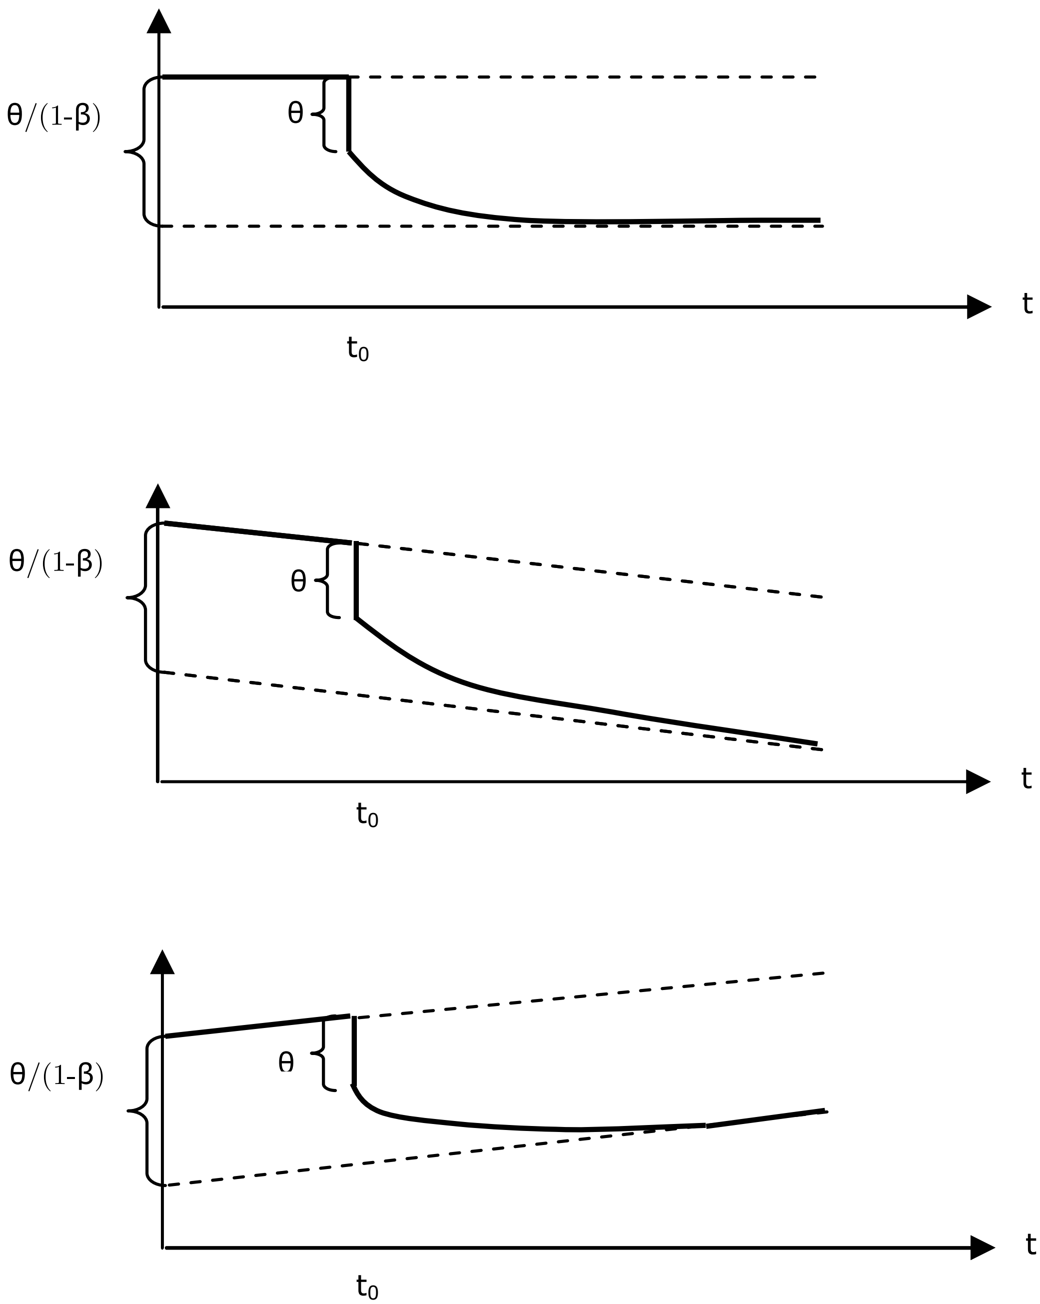

Figure 3: Interpreting Regression Coefficients

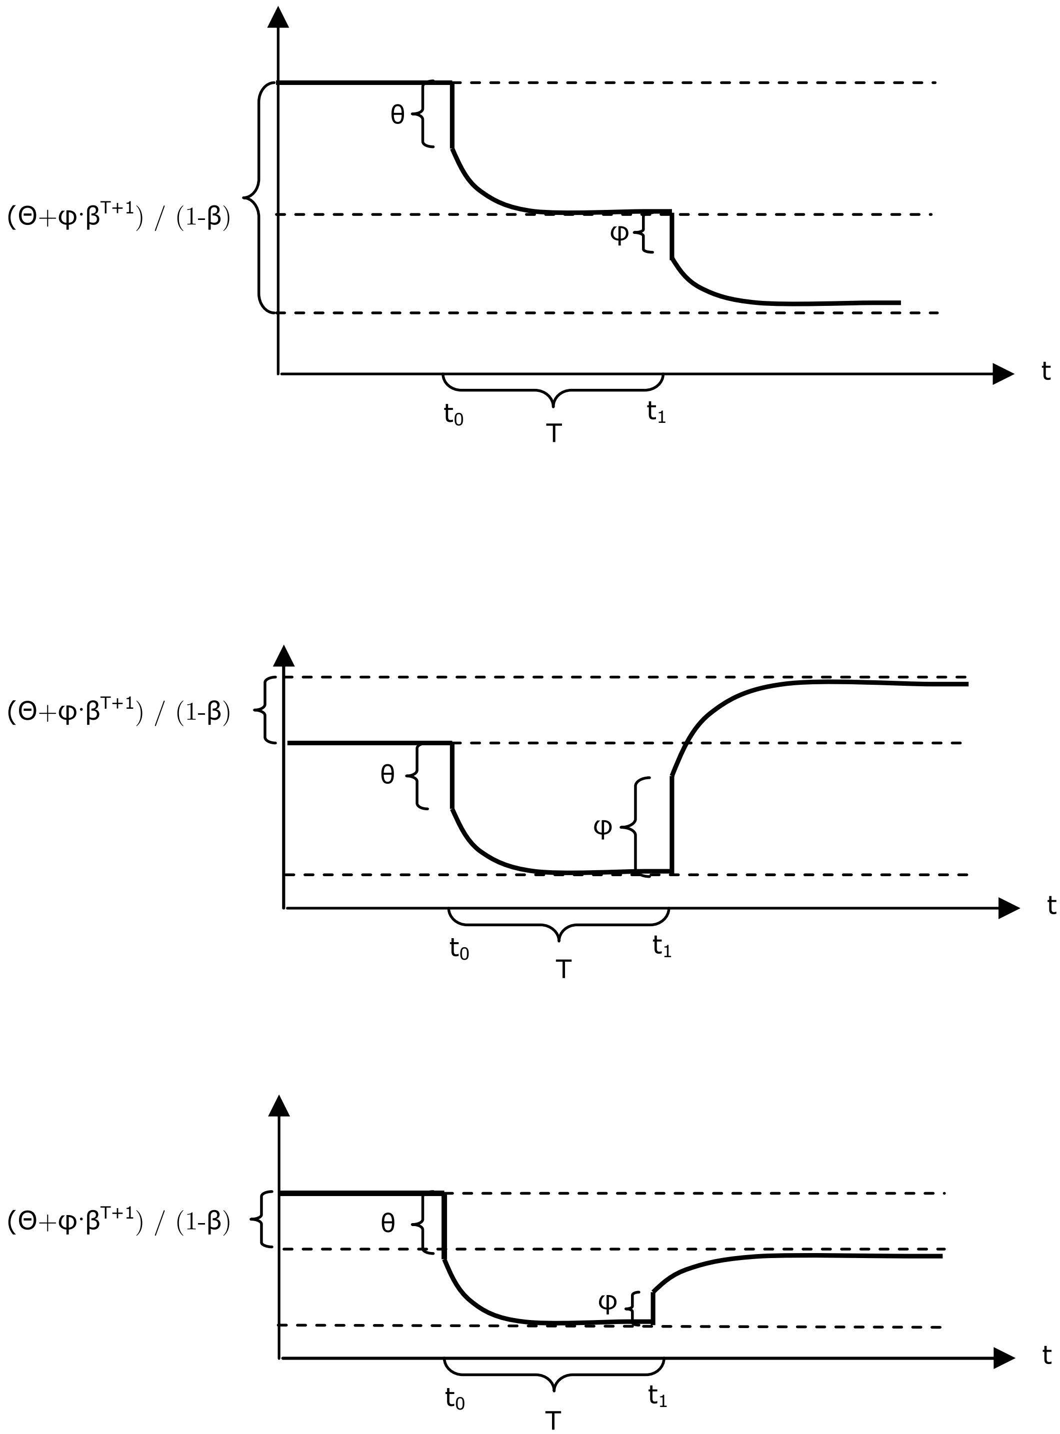

Figure 4: Interpreting Regression Coefficients with Two Wal-Mart Coefficients

move slowly, and do not return to their expected (trend-adjusted) level for several quarters following any "shock".

In the presence of measurement error in prices, the estimate of β is attenuated (biased towards zero), and other point estimates, including the estimate of θ, are also biased. If the measurement error is solely, or primarily, due to errors in the process of collecting prices – e.g., writing down a price of 0.79 when the true price is 0.97 – so that the errors are uncorrelated over time, we can address this problem by using $\mathbf{p}_{kj,t-2}$, the second lag of price, to instrument for $\mathbf{p}_{kj,t-1}$. Unless otherwise noted, all the estimates shown below use this IV strategy to eliminate (or at least reduce) the problem of measurement error.

If measurement error is due to other causes – such as non-representative selection of retail establishments in the price survey – twice-lagged price may not be a valid instrument. For this reason I experiment with a functional form that omits lagged price:[16]

$$\mathrm{p}_{jt} = \alpha + \theta \mathrm{WMopen}_{jt} + \sum_j \gamma_j \mathrm{city}_j + \sum_t \delta_t \mathrm{quarter}_t + \sum_j \tau_j \mathrm{trend}_t + \varepsilon_{jt} \tag{3}$$

Including lagged prices on the right-hand side of the regression is problematic for another reason: lagged price may be endogenous. Endogeneity of lagged price could arise because the error term $\varepsilon_{\mathbf{jt}}$ is not i.i.d. – for example if it follows an AR or MA process – in which case, the error term and lagged price will be correlated. Like measurement error, endogeneity of lagged price would also yield biased estimates of Wal-Mart's effect in these regressions.[17]

Instrumenting for lagged price with twice-lagged price, as described above, is a valid correction for endogeneity under some – but not all – conditions. Specifically, if the error term is MA(1), an IV specification would be valid; but if the error term is AR(1), it would not: in that case twice-lagged price (as well as any longer lag) would still be correlated with the error term. The alternative specification with no lagged price may be preferable under these conditions. (A completely exogenous instrument would be the best solution to this problem, but none is available, as noted in footnote 16.) While neither specification is perfect, taken together, the

[16]There is no available alternative instrument for lagged price. A valid instrument needs to be correlated with last quarter's average price of a given good in a city, after controlling for all the covariates, but be uncorrelated with the error in measurement of lagged price.

[17]The same problem would result if there is an omitted determinant of price, which is correlated with lagged price, and is not captured by any of the other regressors in the model. (A possible example could be entry of another large retailer, if it depresses prices.) Since the regression includes both city and time (quarter) fixed effects, to induce bias, this omitted variable must vary across both time and cities, and not follow a linear trend at the city level.

two specifications tell a more complete story than either one tells alone.

4.3 Measurement Error in Wal-Mart Data

I use instrumental-variables estimates to correct for two problems in the Wal-Mart variable: measurement error and endogeneity. Inaccuracies in the directories and store lists, combined with the probabilistic method of assigning opening dates across quarters (and across years over the period 1990-1993), lead to measurement error in the variable $\mathbf{WMopen_{jt}}$, causing attenuation bias in estimates of θ. Endogeneity may also a problem, even after controlling for city-level trends, if Wal-Mart entry is more likely in times of high (or low) prices, relative to the city's long-run trend. In that case, mean reversion alone will cause prices to fall (rise), and the estimated coefficient θ may be spuriously negative (positive).

Measurement error in the Wal-Mart entry variable $\mathbf{WMopen_{jt}}$ takes a particular form: while the entered cities are correctly identified, the *timing* of entry may be incorrectly measured due to errors in the directories. The variable $\mathbf{WMplan_{jt}}$ is also measured with error, by construction: it represents the number of stores that would be open had stores always opened in the order in which they were planned, with the total number of stores opening each year is held at its actual level, and the distribution of opening dates across quarters assumed constant – at 2001-04 levels – over the 20 year period.

As long as the measurement errors in the two variables is classical and uncorrelated, we can use $\mathbf{WMplan_{jt}}$ to instrument for $\mathbf{WMopen_{jt}}$. The second assumption appears to be correct; it would be violated if, for example, stores in rural areas are completed faster than stores in urban areas and also appear later in directories, but this does not appear to be the case.[18] But measurement error is not classical: the actual number of Wal-Mart stores in city j in quarter t differs from the measured number by a number whose expected value is correlated with the measured number. (For example, when the directories report zero Wal-Marts in town, the expected number of stores is some (small) positive number. When the reported number of Wal-Marts is one, the actual number is usually either zero or one, so the expected number is less than one.) This induces a slight bias in the instrumental-variables results reported here.[19]

[18] There is no correlation between observable location characteristics and the time lag between the "planning date" and the opening date of a store.

[19] Kane, Rouse and Staiger [20] suggest a GMM estimator to address this problem. Unfortunately, due to the size of the panel, their solution is not computationally feasible in this setting. While the bias could be large in theory, in the static returns-to-schooling example of Kane, Rouse and Staiger, the bias is approximately 5% of the IV coefficient estimate. The sign of this bias is indeterminate.

The first-stage regression associated with Equation (1) is

$$\text{WMopen}_{jt} = \widetilde{\alpha} + \widetilde{\beta}\text{p}_{j,t-1} + \sum_t \widetilde{\delta}_t \text{quarter}_t + \widetilde{\theta}\text{WMplan}_{jt} + \sum_j \widetilde{\tau}_j \text{trend}_{tj} + \widetilde{\varepsilon}_{jt}.$$

Although the first-stage regression has a binary dependent variable, I estimate the first-stage by Ordinary Least Squares rather than a nonlinear model such as probit or logit. Angrist and Kreuger [2] caution against using a nonlinear first-stage model because, unless it is exactly right, it will generate inconsistent second-stage coefficients. A first-stage OLS specification, in contrast, yields consistent second-stage estimates even if it is not exactly correct.

This IV strategy can also correct for possible endogeneity of Wal-Mart's entry decision. Concerns about endogeneity have two dimensions: Wal-Mart may select the cities it enters non-randomly, and it may choose the timing of entry non-randomly.

The cross-sectional dimension (choice of cities) is very plausible; for example, Wal-Mart may prefer to enter cities with less-competitive retail markets (hence higher pre-entry prices) or with a larger fraction of search-savvy lower-middle income families (whose presence leads to *lower* average retail prices; see Frankel and Gould [14]). This concern is greatly mitigated by the fact that Wal-Mart entry is observed in 162 of the 165 sample cities, and the sample cities are diverse with respect to all standard economic and demographic variables (see Table 3).

The timing dimension is important if Wal-Mart can schedule its entry to coincide with high retail prices. Under reasonable assumptions, using $\mathbf{WMplan_{jt}}$ to instrument for $\mathbf{WMopen_{jt}}$ can correct not only for measurement error but also for endogeneity in the timing of entry. Of course, if there are omitted variables – if Wal-Mart's entry is correlated with other shocks to the local retail market, such as the entry or exit of other stores – in which case OLS estimates may confound Wal-Mart's effect with price changes that happen for other reasons. (Any omitted city-specific variable that is either unchanged over this period, or moves linearly, will be captured by the city fixed effects and/or the city-specific trends in the model.)

The identification strategy assumes that Wal-Mart plans its store entries well in advance of entry and cannot accurately forecast exact market conditions (prices) at the time for which entry is planned. This assumption seems reasonable: it can take a year, and often longer, to get a property zoned, acquire all the necessary building permits, build and open a store for business. While Wal-Mart may want to enter any market at a time when prices are relatively high – when potential profits, as well as apparent benefit to consumers, are highest – it cannot predict the timing of high prices with any accuracy if the delay between planning and entry is long enough. The company may fine-tune opening dates based on current market conditions, but if the planning occurs sufficiently in advance – and prices are sufficiently flexible – then planning dates can be treated as exogenous to retail prices at the time of Wal-Mart entry. So, while the *actual* date of entry may be manipulated – moved forwards or delayed – to correspond with "favorable" (high) price conditions, the planned date of entry will be largely free of those biases, at least once we account for long-run differences in average prices across cities.[20]

5 Results

5.1 OLS Results

Before proceeding to the IV results, I show some results from OLS regressions of Equation (1). Unless otherwise noted, the specifications below use twice-lagged price as an instrument for lagged price; I use the term "OLS" to mean that the Wal-Mart variable, $\mathbf{WMopen_{jt}}$, is not instrumented.

Table 4 shows results from 10 separate OLS regressions. Each row represents a separate regression, and shows the estimated coefficients θ (the instantaneous effect of Wal-Mart entry) and β (the mean-reversion/price stickiness parameter), as well as the estimated long-run effect of entry, $\frac{\theta}{1-\beta}$. Significance levels for the latter are given by the p-value from a (two-sided) Wald test for $\mathrm{H}_0 : \frac{\theta}{1-\beta} = 0$. So, for example, the average price of detergent is estimated to decline by 0.78% in the quarter following Wal-Mart entry, and reach a decline of 3.86% asymptotically.

OLS estimates of the price effect of Wal-Mart are negative for 9 of the 10 products (cigarettes are the exception), and significant for three products – detergent, shampoo and toothpaste (although the latter is significant only at the 10% level). For these three products, the estimated short-run effect has a magnitude of 0.7-0.8%, and the long-run magnitude of 3.5-4%.

The distinction between short- and long-run effects of Wal-Mart depend on the inclusion of

[20] A possible exception to this argument is if Wal-Mart's entry is correlated with the entry (or exit) of other firms. If both the planning and the implementation of Wal-Mart entry coincides with that of other retailers – for example, if entry tends to be associated with the building of new "strip malls" that bring additional stores – coefficient estimates of the impact of Wal-Mart incorporate the impact of the ancillary entries. By attributing the joint causal effect to Wal-Mart, we are assuming that the ancillary entries would not have occurred without Wal-Mart's. If this assumption is invalid – that is, if other entries occur independently of, but concurrently with, Wal-Mart's – the IV strategy described here would no longer be valid. Anecdotal evidence suggests that other stores enter (and exit) in response to Wal-Mart, so that the instrument is valid. Unfortunately, data on the timing of entry of other large retailers are not available, so I cannot control directly for these.

lagged price in the regression, and on the validity of the IV correction for measurement error in lagged price (and possibly for endogeneity as well). If lagged price is omitted from these regressions, the estimated effect of Wal-Mart (not shown) is negative for 8 of the 10 products, and statistically different from zero in four cases. The statistically-significant coefficients range from a 2.2% estimated decline in the price of shampoo (significant at the 10% confidence level) to a 3.6% estimated decline in the price of detergent (significant at the 1% confidence level).

5.2 IV Results

Table 5 show instrumental variables estimates corresponding to Table 4. As expected where measurement error is present, almost all point estimates of θ are (absolutely) larger than the OLS estimates. Point estimates of Wal-Mart's effect show statistically significant short-run declines of 1.5-3% in the prices of aspirin, detergent, Kleenex and toothpaste (with long-run declines of 8-13%), and marginally significant declines, with similar magnitudes, for shampoo. Only estimates of Wal-Mart's impact on the prices of cigarettes and pants are positive, and neither is significantly different from zero.

Table 6 shows IV estimates of Equation (3), which omits the lagged price variable. Qualitatively, the results are extremely similar to the results in Table 5. The point estimates of θ lie between the short- and long-run point estimates from the full model, near their midpoint. Since results are qualitatively similar, we can infer that uncorrected endogeneity of lagged price is not seriously biasing the estimates of Wal-Mart's impact. In these specifications, the test statistic for joint significance of the city-level trends increases by a factor of 10 or more, exceeding 500,000 for all regressions (and reaching 50,000,000 for cigarettes); this suggests that there is a strong relationship between price observations one quarter apart. The model with lagged price captures this relationship explicitly. It also allows us to distinguish between the short- and long-run effects of entry, which this restricted model does not.

Table 7 shows IV estimates of Equation (2), which allows for different effects of Wal-Mart on prices in the short- and medium-runs (up to five years after entry) and the long-run. The coefficient ϕ represents the additional (positive or negative) change in price five years after Wal-Mart entry. Estimates of ϕ are generally small relative to θ; it is positive, large (2%) and significant for one product (toothpaste). With this one exception, the estimates are broadly similar to the ones shown in Table 5, in which only a single coefficient was estimated. Long-run effects are negative for all products except underwear.

Table 4: OLS Estimates

Product	θ	β	$\frac{\theta}{1-\beta}$ [a]
Aspirin	-0.0025	0.7721***	-0.0111
	(0.0041)	(0.0189)	
Cigarettes	0.0015	0.8350***	0.009
	(0.0014)	(0.0153)	
Coke	-0.0035	0.7981***	-0.0175
	(0.0039)	(0.0416)	
Detergent	-0.0083***	0.7946***	-0.0406***
	(0.0032)	(0.0206)	
Kleenex	-0.0030	0.8043***	-0.0152
	(0.0024)	(0.0248)	
Pants	-0.0038	0.7556***	-0.0156
	(0.0040)	(0.0238)	
Shampoo	-0.0090**	0.7932***	-0.0437**
	(0.0035)	(0.0232)	
Shirt	-0.0057	0.7194***	-0.0203
	(0.0045)	(0.0269)	
Toothpaste	-0.0067*	0.7899***	-0.0319*
	(0.0040)	(0.0225)	
Underwear	0.0006	0.7575***	0.0026
	(0.0057)	(0.0235)	

Robust standard errors in parentheses (clustered by city)
Regressions include city and quarter FE and city-specific price trends;
F statistics from joint test for city level trends range form 20,000-250,000
* significant at 10%; ** significant at 5%; *** significant at 1%
[a] Significance level from Wald test for $H_0 : \frac{\theta}{1-\beta} = 0$

Table 5: IV Estimates

Product	θ	β	$\frac{\theta}{1-\beta}$ [a]
Aspirin	-0.0244**	0.7685***	-0.1053**
	(0.0114)	(0.0197)	
Cigarettes	0.0019	0.8349***	0.0112
	(0.0034)	(0.0153)	
Coke	-0.0034	0.7981***	-0.0169
	(0.0092)	(0.0415)	
Detergent	-0.0193**	0.7904***	-0.0920**
	(0.0085)	(0.0211)	
Kleenex	-0.0156**	0.8015***	-0.0786**
	(0.0064)	(0.0247)	
Pants	0.0015	0.7570***	0.0062
	(0.0108)	(0.0232)	
Shampoo	-0.0163*	0.7920***	-0.0784*
	(0.0093)	(0.0233)	
Shirt	-0.0008	0.7191***	-0.0028
	(0.0109)	(0.0271)	
Toothpaste	-0.0286**	0.7846***	-0.1327***
	(0.0111)	(0.0228)	
Underwear	-0.0016	0.7575***	-0.0065
	(0.0142)	(0.0235)	

Robust standard errors in parentheses (clustered by city)
Regressions include city and quarter FE and city-specific price trends
F statistics from joint test for city level trends range from 25,000-450,000
* significant at 10%; ** significant at 5%; *** significant at 1%
[a] Significance level from Wald test for $H_0 : \frac{\theta}{1-\beta} = 0$

This finding effectively rules out many possible scenarios, among them the possibility that, while Wal-Mart's entry initially lowers prices, prices rise again once competitors are driven out of the market. Instead, Wal-Mart entry is associated with one relatively large price shock, which reverberates through the auto-regressive component.[21]

5.3 Differential Impact of First and Later Entries

In the preceding regressions, I consider only the effect of the *first* Wal-Mart store to enter a city. But many cities have more than one Wal-Mart store. In 43% of sample cities – 70 of 164 – two or more store were open by the end of the sample period. (Of these, 36 cities had two stores, and 14 had three. One city – Houston – had nine stores by the end of the sample period.)

To test whether the second Wal-Mart store has the same effect as the first, I estimate

$$\begin{aligned} \mathrm{p}_{jt} = \alpha + \beta \mathrm{p}_{j,t-1} + \theta_1 \mathrm{WMopen}_{jt} + \theta_2 \mathrm{WMopen2}_{jt} \\ + \sum_j \gamma_j \mathrm{city}_j + \sum_t \delta_t \mathrm{quarter}_t + \sum_j \tau_j \mathrm{trend}_{tj} + \varepsilon_{jt} \end{aligned} \tag{4}$$

where $\mathbf{WMopen_{jt}}$ is as defined above, and $\mathbf{WMopen2_{jt}}$ equals 1 if there are at least two stores in city j in quarter t.[22] Estimating $\boldsymbol{\theta_1}$ and $\boldsymbol{\theta_2}$ separately allows us to test directly whether the second store has the same impact as the first.

Estimation results are shown in Table 8. First-store effects are all of the same sign as in the baseline specification (Table 5, with the statistically-significant effects (for aspirin, detergent, Kleenex, shampoo and toothpaste) all negative, larger in absolute terms than in the baseline specification. The effect of the second store varies by product; it is generally of a similar magnitude, but opposite sign, to the first store's effect. In the case of aspirin and toothpaste, it is positive and statistically significant (at 5% and 1%, respectively).

Since $\boldsymbol{\theta_1}$ estimates a once-and-for-all effect, the joint (cumulative) effect of the first and second stores is a function of both parameters. To a first approximation (ignoring the time lag between the first and second store openings), the joint effect of two stores is the sum of the two coefficients θ_1 and θ_2. (The average lag between the first and second store openings in my

[21] Estimates of this specification without lagged price are qualitatively similar, but magnitudes of the long-run effect, specified as $\theta + \phi$, are slightly smaller in absolute terms.

[22] I also define $\mathbf{WMplan2_{jt}}$, which equals 1 if at least two Wal-Mart stores were planned to have opened in city j by quarter t, and use it as an instrument along with $\mathbf{WMplan_{jt}}$.

sample, for cities with more than one store, is three years.) This sum is generally very small, and tends to be negative but statistically insignificant.

In specifications omitting the lagged dependent variable, estimates of both θ_1 and θ_2 are larger. The sum $\theta_1 + \theta_2$ is statistically different from zero for two products – detergent and Kleenex (at 10% and 5% significance, respectively); in both cases the sum is negative.

I also estimate a model in which the effect of Wal-Mart is linear in the number of stores:

$$\mathrm{p}_{jt} = \alpha + \beta \mathrm{p}_{j,t-1} + \theta \mathrm{WMsopen}_{jt} + \sum_j \gamma_j \mathrm{city}_j + \sum_t \delta_t \mathrm{quarter}_t + \sum_j \tau_j \mathrm{trend}_{tj} + \varepsilon_{jt} \qquad (5)$$

where $\mathbf{WMsopen_{jt}}$ is the number (0-10) of Wal-Mart stores open in city j in quarter t (and $\mathbf{WMsplan_{jt}}$, the number of Wal-Mart stores planned to have opened in city j by quarter t, serves as the instrument). Across all products, estimated θ coefficients (not shown) are small (below 0.5% in absolute terms), and statistically indistinguishable from zero; seven of the ten are negative.

5.4 City Size Effect

In this section, I test the hypothesis that Wal-Mart's effect will be larger in cities with less competitive retail markets. Bresnahan and Reiss [4] and Campbell and Hopenhayn [5], among others, note that most models of imperfect competition predict smaller cities will have less competitive retail markets, but this prediction has been hard to test empirically.

There is not a single obvious measure of city size. The number of retail establishments in a city is strongly correlated with population size; in my sample, the correlation coefficient between the 1980 city population and the number of retail establishments the city had in 1982 is 0.98. The number of establishments does not increase one-to-one with population, however; larger cities tend to have a larger absolute number of retail establishments, but fewer establishments per capita. Small cities have disproportionately many small stores (Campbell and Hopenhayn [5]), which are likely to be less efficient – and more expensive.[23]

In Figure 5 I plot the number of retail establishments in 1982 against the number of retail establishments per capita.[24] Outlier cities are individually labeled. The figure highlights

[23] Dinlersoz [10] notes another difference between the organization of retail markets in small and large cities: larger cities have relatively fewer chain stores and more stand-alone stores, which are likely to behave more competitively. Counts of chains and stand-alone stores are not available for my sample.

[24] Only the 142 sample cities with 1980 population above 25,000 are included in the figure. I use the 1980 Census

Table 6: IV Estimates without Lagged Price

Product	θ	Product	θ
Aspirin	-0.0666*	Pants	0.0274
	(0.0341)		(0.0318)
Cigarettes	0.0153	Shampoo	-0.0362
	(0.0119)		(0.0250)
Coke	0.0110	Shirt	0.0194
	(0.0282)		(0.0291)
Detergent	-0.0589**	Toothpaste	-0.0863***
	(0.0274)		(0.0310)
Kleenex	-0.0325	Underwear	-0.0058
	(0.0205)		(0.0404)

Robust standard errors in parentheses (clustered by city)
Regressions include city and quarter FE and city-specific price trends; F statistics from joint test for city level trends exceed 500,000 in all regressions
* significant at 10%; ** significant at 5%; *** significant at 1%

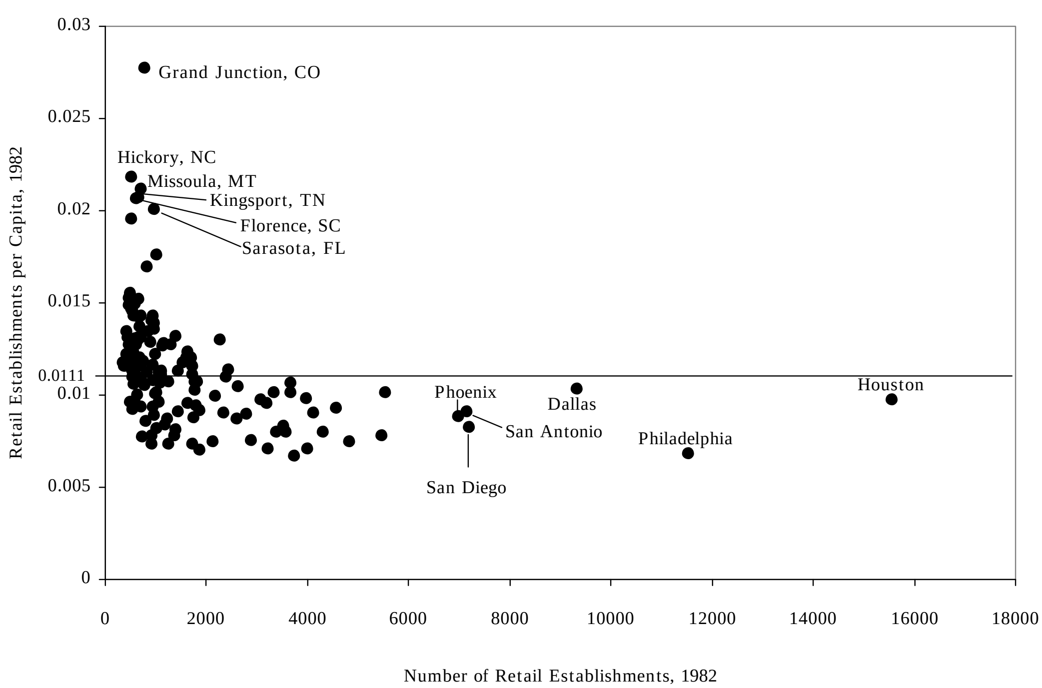

Figure 5: Total Retail Establishments vs. Establishments per Capita (1982)

Table 7: IV Estimates with Different Short- and Long-Run Effects

Product	θ	ϕ	β	$\frac{\theta+\phi\cdot\beta^{21}}{1-\beta}$ [a]
Aspirin	-0.0235**	0.0011	0.7685***	-0.1016**
	(0.0103)	(0.0081)	(0.0198)	
Cigarettes	-0.0001	-0.0030	0.8359***	-0.0013
	(0.0031)	(0.0029)	(0.0159)	
Coke	-0.0011	0.0032	0.7958***	-0.0051
	(0.0084)	(0.0067)	(0.0419)	
Detergent	-0.0148*	0.0062	0.7979***	-0.0728**
	(0.0077)	(0.0053)	(0.0213)	
Kleenex	-0.0129**	0.0040	0.8001***	-0.0645**
	(0.0061)	(0.0033)	(0.0246)	
Pants	-0.0056	-0.0109	0.7503***	-0.0226
	(0.0101)	(0.0070)	(0.0249)	
Shampoo	-0.0180**	-0.0028	0.7888***	-0.0854**
	(0.0087)	(0.0068)	(0.0233)	
Shirt	-0.0009	-0.0002	0.7191***	-0.0034
	(0.0103)	(0.0078)	(0.0272)	
Toothpaste	-0.0147	0.0203***	0.7883***	-0.0689
	(0.0094)	(0.0073)	(0.0229)	
Underwear	0.0012	0.0043	0.7565***	0.0051
	(0.0125)	(0.0090)	(0.0237)	

Robust standard errors in parentheses (clustered by city)
Regressions include city and quarter FE and city-specific price trends;
F statistics from joint test for city level trends range from 25,000-120,000
* significant at 10%; ** significant at 5%; *** significant at 1%
[a] Significance level from Wald test for $H_0 : \frac{\theta+\phi\cdot\beta^{21}}{1-\beta} = 0$

Table 8: IV Estimates with Different First- and Second-Store Effects

Product	θ_1	θ_2	β
Aspirin	-0.0355**	0.0382**	0.7660***
	(0.0153)	(0.0182)	(0.0202)
Cigarettes	0.0033	-0.0050	0.8336***
	(0.0045)	(0.0055)	(0.0159)
Coke	-0.0051	0.0057	0.7983***
	(0.0124)	(0.0167)	(0.0415)
Detergent	-0.0249**	0.0199	0.7899***
	(0.0108)	(0.0133)	(0.0213)
Kleenex	-0.0204**	0.0171	0.8032***
	(0.0084)	(0.0107)	(0.0246)
Pants	0.0037	-0.0078	0.7579***
	(0.0137)	(0.0162)	(0.0232)
Shampoo	-0.0209*	0.0161	0.7921***
	(0.0121)	(0.0148)	(0.0234)
Shirt	-0.0075	0.0237	0.7164***
	(0.0147)	(0.0203)	(0.0268)
Toothpaste	-0.0453***	0.0566***	0.7787***
	(0.0166)	(0.0209)	(0.0236)
Underwear	-0.0040	0.0085	0.7568***
	(0.0176)	(0.0226)	(0.0235)

Robust standard errors in parentheses (clustered by city)
Regressions include city and quarter FE and city-specific price trends;
F statistics from joint test for city level trends range from 40,000-500,000
* significant at 10%; ** significant at 5%; *** significant at 1%

the strong negative relationship between the number of retail establishments and their number relative to population (the correlation coefficient is -0.37): Grand Junction, Colorado (1980 population: 28,000) has the highest number of retail establishments per capita, while Philadelphia (1980 population: 1.7m) has the second-lowest (the lowest is in New Orleans, with 1980 population of 560,000). The horizontal line at 0.0111 indicates the median for this sample.

While establishment breakdowns are not available at the city level, approximately 10% of retail establishments in the US in 1982 were clothing stores (SIC 5600) and 3.8% of retail establishments were "drug stores and proprietary stores" (SIC 5910). (Both fractions declined over the 1980s and 1990s, each by approximately one percentage point.) Therefore, a city with 987 retail establishments (the median in my sample in 1982) had approximately 37 drugstores – well within the intermediate range (between quintopolies and urban markets) considered by Bresnahan and Reiss [4].

I estimate

$$\begin{aligned} \mathrm{p}_{jt} = \alpha + \beta \mathrm{p}_{j,t-1} + \theta_1 \mathrm{WMopen}_{jt} \cdot \mathrm{low}_j + \theta_2 \mathrm{WMopen}_{jt} \cdot \mathrm{high}_j \\ + \sum_j \gamma_j \mathrm{city}_j + \sum_t \delta_t \mathrm{quarter}_t + \sum_j \tau_j \mathrm{trend}_{tj} + \varepsilon_{jt} \quad (6) \end{aligned}$$

where **low** and **high** are indicator variables for cities with number of retail establishments per capita in 1982 below and above the median (0.0111), respectively.

Table 9 shows estimates of the coefficients of interest, θ_1, θ_2, and β. For nine of the ten products, θ_2 is estimated to be negative (and statistically different from zero), with Wal-Mart's short-run effects estimated to be between 2-4%; θ_1 is never statistically different from zero. The difference between these coefficients is statistically significant for shampoo at the 10% significance level, for shirt at the 5% significance level, and for detergent and Coke at the 1% level. Long-run effects are magnified by β, and range from 10-20% for small cities (cities with many retail establishments per capita); long-run effects for large cities (with few establishments per capita) are never statistically different from zero.[25]

population, rather than 1982 population estimates, to compute the per-capita number of retail establishments.

[25]When lagged price is omitted from the regressions, estimates of θ_1 are negative for 6 of the 10 products, and estimates of θ_2 are negative for 7 of the 10 products. In contrast to the main specification, however, θ_2 is statistically significant in only three cases: for detergent, the estimated price decline for cities with many retail establishments per capita is 11% (significant at the 1% level); for toothpaste, it is 9.5% (significant at the 5% level) and for shampoo, it is 6% (significant at the 10% level). The estimate of θ_1 is significant (at the 10% level) in one case, for the price of a shirt, where it implies a price increase of 6% due to Wal-Mart entry.

As a robustness check, I also estimate a model in which I interact the number of retail establishments per capita in 1982 (a continuous variable) with the Wal-Mart variable. (The equation also includes a main Wal-Mart effect; city fixed effects prevent estimation of a main 1982-establishments-per-capita effect.) The interaction term is negative in 9 of 10 regressions (not shown), again indicating that Wal-Mart's negative price effect is larger in cities with many establishments per capita. For three products – Coke, shampoo, and shirt – the interaction term is statistically significant at the 10% level; and it is significant at the 5% level for detergent.

Interestingly, when I separate the cities in the sample into "small" and "large" using alternative criteria – such as 1980 city population, or 1982 absolute number of retail establishments – the two coefficients (θ_1 and θ_2) are never statistically different from one another.[26] This suggests that the number of establishments per capita may be a better measure of competitiveness than the absolute number of retail establishments in a city. However, the strong correlation between population, absolute number of retail establishments, and number of establishments per capita prevents us from drawing any definitive conclusions on the main cause of the variation in Wal-Mart's impact: when I allow for *simultaneous* interactions of Wal-Mart entry with log population and retail establishments per capita, none of the interaction coefficients are statistically different from zero.

6 Discussion and Conclusion

This paper tests the hypothesis that Wal-Mart entry leads to lower average retail prices in the markets it enters. Using a unique panel data set that combines average retail prices for 10 specific goods and a complete time-series of Wal-Mart store locations, I estimate the effect of Wal-Mart entry on prices. I find that for many items typically sold in drugstores, such as aspirin and shampoo, average prices decline following Wal-Mart entry. This decline is economically large – 1.5-3% in the short run, and four times as much in the long run – and statistically significant.

These findings are in line with other price surveys. In April 2002 UBS Warburg collected prices of 100 grocery and non-grocery items in 4-5 grocery stores in each of four large markets: Sacramento, a city with no Wal-Mart presence; and Las Vegas, Houston and Tampa, each of

[26] In a model with continuous interaction term with log population, Wal-Mart's effect varies very little with population; only in two cases is the coefficient different from zero at the 5% confidence level.

which had at least one Wal-Mart Supercenter. Their study found that Wal-Mart's prices were 17-39% lower than competitors' prices in the three "Wal-Mart cities," and that average prices at other grocery stores were 13% lower in the Wal-Mart cities than in Sacramento (Currie and Jain [9]). I repeated Currie and Jain's analysis using a subset of 24 drugstore products from their data set comparable to the ACCRA products: Tylenol, Pepto Bismal, shampoo, deodorant, feminine hygiene items, soap, toothpaste, detergent and Coke. For these items, Wal-Mart's prices were 23% lower on average than competitors' prices in the Wal-Mart cities. Competitors' prices in Wal-Mart cities were lower than Sacramento prices for most, but not all, items; on average, drugstore prices were 15% lower in Wal-Mart cities. Of course, these numbers should be interpreted cautiously, due to the cross-sectional nature of the data, the small number of cities included, and the exclusive reliance of the sample on grocery store prices.

Since figures on profit margins for Wal-Mart competitors in the entered markets are not available, I use Census data to assess the economic significance of these estimates. The Census Bureau's annual series, "Current Business Reports: Annual Benchmark Report for Retail Trade and Food Services," provides data on average gross margins (roughly, the retail price markup over wholesale price) by retail sub-sector from 1986-2001 (U.S. Census Bureau [28] and [29]). Gross margins do not account for any other costs, such as labor, rent, or capital inputs, and should not be mistaken for profit margins. Apparel stores have gross margins above 40% on average, meaning that wholesale costs account for under 60% of the retail prices charged by apparel stores, on average. Drugstores have margins of 25-30% on average.[27] Using this benchmark, a decline of 5% in the price of some drugstore items can be interpreted as a decline of 15-20% in the gross margin of drugstores.[28]

In interpreting the point estimates of Wal-Mart's effect, several caveats are in order.

First, the average-price effect masks a large amount of intra-market variation in competitive response to Wal-Mart entry. Because the ACCRA data cannot be disaggregated to the store

[27] The exact definition of these industries changes halfway through the sample due to the shift from SIC to NAICS. For the period 1993-1998, when both SIC and NAICS-based figures are available, the implied margins are very similar.

[28] An alternative, more detailed, benchmark is provided by the University of Chicago Graduate School of Business study of Dominick's Finer Foods (DFF) in the Chicago metropolitan area. Over the period 1989-1994, the GSB collected scanner data, including gross margins, from Dominick's for thousands of individual items. These data are available on-line at
`http://gsbwww.uchicago.edu/kilts/research/db/dominicks/index.shtml`
and are described in detail in Peltzman [24] and Chevalier, Kashyap and Rossi [8]. Several of the items included in this analysis are also included in the DFF database. The average gross margin on items classified as "drugstore products" in this study is 12%.

level, it is impossible to estimate the distribution of responses. But theory suggests that stores selling the closest substitutes to Wal-Mart – for example, those that are located near Wal-Mart, or that are similar on other dimensions – will have the most elastic price responses to Wal-Mart entry. Stores located far from Wal-Mart are likely to have very small price responses, because their clienteles' cross-price elasticity of demand will be low.

A related problem is that Wal-Mart stores may be included in ACCRA's surveys. If Wal-Mart's prices are lower, on average, than other stores' (Currie and Jain [9], Hausman and Leibtag [17]), then it is impossible to distinguish, in my results, between Wal-Mart's direct effect on average prices, and the more interesting indirect effect due to competitive pressures on other stores. ACCRA's explicit instruction to sample only retailer that cater to the upper quintile of the income distribution mitigates this concern, because it reduces the probability that Wal-Mart will be sampled. It also reduces the probability that many of Wal-Mart's direct competitors – other discount retailers – will be in the sample. The absence of discount stores and, more generally, the over-sampling of stores that cater to the upper quintile, biases my estimates against finding any effect of Wal-Mart.[29]

Third, ACCRA weights all establishments equally in computing average prices. If Wal-Mart's prices are indeed lower than their competitors' prices, and consumers respond to Wal-Mart entry by shifting demand from incumbent establishments to Wal-Mart, the effect on the unweighted price average, estimated here, will be lower than the effect on the effective average price paid by consumers (properly weighted). This issue is important, but require additional data – on quantities purchased – to resolve. [30]

Pricing strategies that vary by type of store may bias the result in the opposite direction. Many drug- and grocery stores use so-called "High-Low" pricing, so that an item alternates between a high "regular" and a low "sale" price, while Wal-Mart uses "every day low pricing," which is lower on average than competitors' prices but is rarely the lowest price in the market. For non-perishable goods such as the ones considered in this paper, consumers' ability to buy in bulk during sales means that the unweighted average price computed by ACCRA under-weights sale prices. If Wal-Mart's entry pushes down the "high" (regular) price charged by incumbents

[29]This effect is probably strongest for clothing, where there is a wide variance in store attributes, even if the product is relatively homogeneous. This may explain why my estimates of Wal-Mart's effect on clothing prices are consistently nil.

[30]Hausman and Leibtag [17] raise similar points in their discussion of Wal-Mart's effect on aggregate inflation; they use household-level scanner data to estimate these biases in the CPI, and find them to be quite large.

but does not affect either the probability of a sale or the sale price, the estimated impact of Wal-Mart will be much larger than its properly-weighted effect. Quantity data are needed in order to address this issue as well.

Finally, the products selected by ACCRA are not a random sample. For the purpose of price comparison across cities, ACCRA selected well-known national brands for its price survey; these items' prices may not be representative of "typical" drugstore and clothing prices. If the difference between Wal-Mart's and other stores' prices is greater for national brands than for local and "store" brands – perhaps because national brands have higher margins – the estimated effect of Wal-Mart will be biased away from zero. We do not observe that margins are higher on national brands in the Dominick's price data; if anything, the brands considered here have lower margins than average in their categories (e.g., margins for Colgate and Crest toothpaste are lower than the average toothpaste margin).

Despite these caveats, the estimated effect of Wal-Mart on the prices of several products considered here are strong and robust. Wal-Mart's effect is strongest for products traditionally sold in drugstores, and weakest, or absent, for cigarettes and Coke (sold in many outlets, including convenience stores) and clothing. This result is intuitively appealing, since (with the relatively recent exception of grocery stores) Wal-Mart competes most directly with drugstores. The estimated effects are also strongest for cities with a large number of retail establishments per capita, consistent with the intuition that Wal-Mart's entry brings lower prices to consumers in relatively small cities, where establishments tend to be smaller, and retail environments less competitive, than in large cities.

Table 9: IV Estimates by Retail Establishments per Capita

Product	θ_1	θ_2	$\theta_1 = \theta_2$ [a]	β	$\frac{\theta_1}{1-\beta}$ [b]	$\frac{\theta_2}{1-\beta}$ [b]	$\frac{\theta_1}{1-\beta} = \frac{\theta_2}{1-\beta}$ [a]
Aspirin	-0.0207	-0.0234	0.0202	0.7625***	-0.0873	-0.0986	0.0204
	(0.0141)	(0.0159)	(0.8871)	(0.0225)			(0.8868)
Cigarettes	0.0006	-0.0010	0.0714	0.8324***	0.0036	-0.0059	0.0715
	(0.0047)	(0.0045)	(0.7897)	(0.0164)			(0.7895)
Coke	0.0174	-0.0296**	9.5615	0.8095***	0.0912	-0.1553**	7.4456
	(0.0114)	(0.0121)	(0.0024)	(0.0431)			(0.0072)
Detergent	-0.0003	-0.0458***	9.2884	0.7814***	-0.0015	-0.2093**	10.054
	(0.0106)	(0.0122)	(0.0028)	(0.0228)			(0.0019)
Kleenex	-0.0063	-0.0221**	2.2595	0.7964***	-0.0311	-0.1085**	2.2755
	(0.0073)	(0.0093)	(0.1350)	(0.0254)			(0.1337)
Pants	0.0139	-0.0042	1.1312	0.7651***	0.0590	-0.0179	1.1490
	(0.0136)	(0.0131)	(0.2893)	(0.0264)			(0.2856)
Shampoo	-0.0020	-0.0306**	3.4004	0.7927***	-0.0098	-0.1474**	3.4106
	(0.0125)	(0.0119)	(0.0673)	(0.0241)			(0.0669)
Shirt	0.0178	-0.0290*	5.7983	0.7056***	0.0606	-0.0986*	6.0740
	(0.0128)	(0.0169)	(0.0173)	(0.0304)			(0.0149)
Toothpaste	-0.0177	-0.0360**	0.9560	0.7826***	-0.0816	-0.1654**	0.9763
	(0.0145)	(0.0142)	(0.3299)	(0.0245)			(0.3248)
Underwear	-0.0233	0.0170	2.5669	0.7643***	-0.0988	0.0722	2.7226
	(0.0197)	(0.0184)	(0.1114)	(0.0250)			(0.1012)

Robust standard errors in parentheses (clustered by city)
Regressions include city and quarter FE and city-specific price trends;
F statistics from joint test for city level trends range from 20,000-200,000
* significant at 10%; ** significant at 5%; *** significant at 1%
[a] F statistic (p-value) from test for $H_0 : \theta_1 = \theta_2$
[b] Significance level from Wald test for $H_0 : \frac{\theta}{1-\beta} = 0$

References

[1] American Chamber of Commerce Research Association, ACCRA Cost of Living Index Manual (2000).

[2] J.D. Angrist, A.B. Krueger, Instrumental Variables and the Search for Identification: From Supply and Demand to Natural Experiments, Journal of Economic Perspectives 15 (2001) 69-85.

[3] E. Basker, Job Creation or Destruction? Labor-Market Effects of Wal-Mart Expansion, Review of Economics and Statistics 87 (2005, forthcoming).

[4] T.F. Bresnahan, P.C. Reiss, Entry and Competition in Concentrated Markets, Journal of Political Economy 99 (1991) 977-1009.

[5] J.R. Campbell, H.A. Hopenhayn, Market Size Matters, Federal Reserve Bank of Chicago Working Paper 2003-12 (2003).

[6] Chain Store Guide, Directory of Discount Department Stores, Business Guides, Inc., New York 1979-1982, 1987-1993.

[7] Chain Store Guide, Directory of Discount Stores, Business Guides, Inc., New York 1983-1996.

[8] J.A. Chevalier, A.K. Kashyap, P.E. Rossi, Why Don't Prices Rise During Periods of Peak Demand? Evidence from Scanner Data, American Economic Review 93 (2003) 15-37.

[9] N. Currie, A. Jain, Supermarket Pricing Survey, UBS Warburg Global Equity Research publication 2002.

[10] E.M. Dinlersoz, Firm Organization and the Structure of Retail Markets, Journal of Economics and Management Strategy 13 (2004) 207-240.

[11] A. Edgecliffe-Johnson, A Friendly Store from Arkansas, Financial Times, June 19, 1999.

[12] J.M. Feiner, A.H. D'Andraia, J. Black, C.E. Jones, R.J. Konik, Wal-Mart Encyclopedia X: Building a Global Brand, Lehman Brothers, New York 2001.

[13] S. Foley, T. Mahmood, Wal*Mart Stores, Inc., Harvard Business School Case Study 9-794-024 (1996).

[14] D.M. Frankel, E.D. Gould, The Retail Price of Inequality, Journal of Urban Economics 49 (2001) 219-239.

[15] P. Ghemawat, Wal-Mart Stores' Discount Operations, Harvard Business School Case Study 9-387-018 (1989).

[16] P. Ghemawat, G. Friedman, Wal-Mart in 1999, Harvard Business School Case Study N9-799-118 (1999).

[17] J. Hausman, E. Leibtag, CPI Bias from Supercenters: Does the BLS Know that Wal-Mart Exists?, National Bureau of Economic Research Working Paper 10712 (2004).

[18] J. Hopkins, Wal-Mart's Influence Grows, USA Today, January 29, 2003.

[19] H. Hotelling, Stability in Competition, Economic Journal 39 (1929) 41-57.

[20] T.J. Kane, C.E. Thomas Rouse, D. Staiger, Estimating Returns to Schooling When Schooling is Misreported, Princeton Industrial Relations Section Working Paper 419 (1999).

[21] The Lex Column, The Wal-Mart Way, Financial Times, February 19, 2003.

[22] McKinsey Global Institute, US Productivity Growth 1995-2000: Understanding the Contribution of Information Technology Relative to Other Factors, MGI, Washington DC 2001.

[23] D.C. Parsley, S.-J. Wei, Convergence to the Law of One Price without Trade Barriers or Currency Fluctuations, Quarterly Journal of Economics 111 (1996) 1211-1236.

[24] S. Peltzman, Prices Rise Faster than They Fall, Journal of Political Economy 108 (2000) 466-502.

[25] J.F. Reinganum, A Simple Model of Equilibrium Price Dispersion, Journal of Political Economy 87 (1979) 851-858.

[26] B. Saporito, Can Wal-Mart Get Any Bigger?, Time, January 13, 2003.

[27] U.S. Census Bureau, County and City Data Book, US Government Printing Office, Washington, DC various years a.

[28] U.S. Census Bureau, Current Business Reports: Annual Benchmark Report for Retail Trade, Washington, DC various years b.

[29] U.S. Census Bureau, Current Business Reports: Annual Benchmark Report for Retail Trade and Food Services, Washington, DC various years c.

[30] S.S. Vance, R.V. Scott, Wal-Mart: A History of Sam Walton's Retail Phenomenon, Twayne Publishers, New York 1994.

[31] Wal-Mart Stores, Inc., Annual Report, 1971-2001.