Do Volume Discounts Contribute to Cyclicality in Prices?

June 2016

Duncan Simester

MIT Sloan School of Management

Volume discounts are widely used in industrial markets. They allow for price variation without intervention by the seller and have the potential to contribute to cyclicality in prices. We might think volume discounts would lead to higher prices during recessions, because customers place smaller orders. We show that the reverse is true. When unemployment increases, customers in industrial markets place larger orders that qualify for larger volume discounts. As a result, volume discounts contribute to falling prices during recessions.

We measure the contribution to falling prices by focusing on industrial sellers' profit margins. We show that profit margins for almost 80,000 industrial components fall by 1.29% when unemployment increases by 1%. Consolidation of orders allows industrial purchasers to qualify for larger discounts. We estimate that this consolidation contributes approximately a third of the reduction in the seller's profit margins.

Our findings are obtained using two datasets describing industrial transactions. One contains $2 billion of sales by a component manufacturer to its global distributors. The second contains $1.3 billion in onward sales from the five largest distributors to their customers (product manufacturers).

We also replicate the key findings using retail transaction data. Volume discounts are available for some consumer products, with retailers generally charging lower unit prices for larger package sizes. We show that consumers purchase a higher proportion of large sizes when unemployment is high, and this contributes to both lower prices and lower retail profit margins.

1. Introduction

There has been extensive recent investigation of variation in retail prices. However, there is a notable paucity of research using data from industrial markets. Perhaps as a result, little attention has been given to the role of volume discounts, and how they contribute to monetary policy. Discounts for larger order quantities are widely available in industrial settings and they have the potential to contribute to cyclicality in prices. We might expect that when economic conditions deteriorate industrial purchasers will respond by placing smaller orders, which could contribute to an increase in prices during a recession. We show that the reverse is true. Adverse shocks result in industrial customers consolidating their orders so that they qualify for larger discounts. As a result, volume discounts contribute to lower prices during recessions. Notably this contribution does not rely upon intervention by the seller and so is not dependent upon the absence of price frictions.

Our primary findings come from two large samples of industrial transactions. The data was provided by an industrial manufacturer that sells products through distributors. The first data describes transactions between this manufacturer and its global distributors. The second dataset describes sales by the distributors to their customers, who are generally product manufacturers that incorporate the components into their products. These include products in a wide variety of industries including consumer electronics, marine, automotive and aerospace. Although the use of volume discounts is a lot more prevalent in industrial markets, volume discounts also exist in retailer settings.[1] Therefore, we also replicate the key findings using a third dataset containing retail transactions provided by a retailer of consumer packaged goods in the United States.

The evidence that industrial purchasers organize their orders so that they qualify for larger discounts when economic conditions deteriorate raises the question as to why they do not organize their orders in the same way at other times. Qualifying for larger volume discounts is presumably always desirable. The simplest answer is managerial inattention (Maćkowiak and Wiederholt 2009; Zbaracki et al. 2006). Managers have to prioritize activities, and in the face of adverse demand shocks they allocate greater priority to managing costs.

To help understand how distributors respond to adverse macro-economic demand shocks we interviewed industrial distributors in several industries. Because many readers will be unfamiliar with industrial components, we will instead illustrate our arguments using an example from another industry (that is not represented in the data). A Vice President at one of New England's largest alcohol distributors described how his firm responded during the 2008-09 recession. They consolidated their orders around their "core items", which are the items that traditionally contribute the bulk of their sales. They continued to promote other items outside this core, as they anticipated that these promotional investments would yield competitive differentiation once financial conditions improved. However, instead of featuring several non-core items a quarter, they featured just one non-core item. This explanation is consistent with the findings that we report. During quarters with higher

[1] For example, in a retail setting Tide laundry detergent pods are sold in different package sizes, with lower unit prices common on the larger package sizes.

unemployment distributors were less likely to order non-core items. However, the items they did order were ordered in larger volumes.

The paper is related to other research studying the behavior of prices from a macroeconomic perspective. The recent empirical research in this area can be loosely categorized into those that study customers' actions (buyer behavior) and those that study retailers' decisions (seller behavior). Examples of research investigating retailer behavior include Bils and Klenow (2004), who confirm that there is considerable variation in retail prices in response to aggregate shocks. Nakamura and Steinsson (2008) demonstrate that much of this price flexibility is due to temporary sales rather than changes to "regular" prices. Eichenbaum et al. (2011) present evidence that most price changes due to sales are associated with a change in wholesale prices, suggesting that sales prices are an important source of price flexibility. More recently, Anderson et al. (2016) present evidence suggesting the opposite; retailers respond to wholesale price changes primarily through their regular prices rather than their sale prices.

Chevalier and Kashyap (2011) argue that when measuring price flexibility it is not sufficient to focus solely on posted prices. If households respond to economic conditions by reallocating their expenditure then the "effective" price flexibility may differ from posted price flexibility. Coibion et al. (2015), compare how posted prices and prices paid vary with regional unemployment rates in the US. Their findings show significant cyclicality in prices paid by consumers, but relatively little cyclicality in prices posted by retailers. They present evidence that the difference reflects reallocation of household expenditures.

In this paper we also focus on buyer behavior and prices actually paid, rather than prices posted by sellers. However, rather than studying retail price data, our primary focus is industrial markets. This is not the first paper to study pricing in industrial markets. For example, in response to deficiencies in BLS aggregate price data, Stigler and Kindahl (1970) collected transaction data from 11 different product groups, ranging from steel and truck motors to plywood and household appliances. The same data was later used by Carlton (1986) who documented nine characteristics of rigidity in industrial prices. A more recent example is Zbaracki et al. (2004), who invested in an extensive data-collection program to obtain data describing the price adjustment costs faced by an industrial manufacturer and its customers. The relative paucity of research studying industrial prices reflects the difficulty of obtaining well-organized industrial transaction data of a scale and breadth that facilitates careful empirical investigations. Perhaps in part because of the difficulties in obtaining detailed datasets describing industrial prices, we have not been able to find any previous research studying how volume discounts contribute to the cyclicality of prices.[2]

The paper proceeds in Section 2, where we describe the three datasets used in the paper. In Section 3 we focus on distributors' purchases from the component manufacturer, and in Section 4 we study sales by the distributors to their customers. In Section 5 we replicate the findings using retail data and the paper concludes in Section 6.

[2] Previous studies of volume discounts in the IO and marketing literatures include theoretical explanations for why volume discounts are optimal, and empirical studies of consumer responses to nonlinear prices, primarily in utility markets (including water, electricity, natural gas, telecommunications and Internet access). Recent examples include Ito (2013) and Yao et al. (2012).

2. Data

The industrial data includes separate datasets describing sales between multiple levels of industrial markets: (a) sales from a component manufacturer to its distributors, and (b) sales from the distributors to their customers. We will discuss each of these datasets in turns, and will then introduce the retail data.

Industrial Transactions: Manufacturer Sales to Distributors

This dataset is a complete record of every sales transaction from a components manufacturer to its global distributors between October 1, 2011 and March 28, 2014. The dataset contains 1.43 million transactions, which contribute $2 billion in revenue. The median order size is 1,600 units, the median unit price is $0.55 (the data treats orders for multiple items as separate orders for each item). In many ways the transaction data resembles retail scan data (such as the Dominick's panel data). The data includes the quantity ordered and the price paid. It also includes the cost of producing each unit. These unit costs vary across time and across countries for the same item, but they do not vary across customers or across order sizes within a country.

A feature of industrial transactions is that manufacturers sometimes rebate a portion of the transaction price back to their distributors. For example, this may occur when the distributor is supplying products to an OEM (original equipment manufacturer, such as General Motors or General Electric), and the OEM has a preferred price that it has negotiated directly with the components manufacturer. Consider for example an automotive component manufacturer that supplies an automobile manufacturer through a distributor. The component manufacturer may provide rebates to the distributor to fund discounts that the automobile manufacturer has negotiated directly with the components manufacturer. Although the data does not link the rebates to individual transactions, we can identify which items and countries distributors receive rebates on. We remove all of these observations from our estimation sample (identifying them at the distributor x item x country level).

The dataset also contains a small number of entries with negative order quantities. The negative order quantities reflect full or partial product returns or corrections of inaccurate data entries. Unfortunately the data does not identify which original order the return or correction modifies. This introduces the risk of error in calculating prices and order sizes. For this reason, we also remove items with negative order quantities (identifying them at the distributor x item x country level).

Like retail prices, industrial markets have "posted prices", in the sense that there is a list price for each item. However, unlike retail prices, these list prices often vary for different customers. Moreover, these list prices are not posted publicly, and so industrial customers typically do not know what other customers are paying for the same item. The prices are also generally not a single price. Instead, they are often a schedule of prices, with volume discounts for larger quantities. The transaction data only describes the price paid for each order. It does not describe the price at other order volumes, or any price information in periods that the distributor did not order the item. Therefore, we supplement the transaction data with a second data set describing the "standard price list" for each item as of January 2015. The standard price list provides a full schedule of quantity discounts by item and geographic

region (we list the countries by region in the Appendix). For example, one (arbitrarily chosen) component has the following price list:

Order Quantity	Price per 100 Units
Up to 1000 units	$650.00
Between 1,000 and 5,000 units	$585.00
Between 5,000 and 25,000 units	$487.50
Between 25,000 and 50,000 units	$422.50
Between 50,000 and 100,000 units	$390.00
Between 100,000 and 250,000 units	$370.50
Between 250,000 and 500,000 units	$357.50
Between 500,000 and 750,000 units	$344.50
Over 750,000 units	$329.75

In the standard price list, 34.9% of the price schedules are a single price, while the remaining 65.1% include volume discounts for larger order volumes. In our analysis we will focus on the 65.1% of items that offer volume discounts, but will also report findings for the complete sample of items. The volume discount schedules include an average of 5.1 pricing levels (for different order volumes), with a maximum of 12 levels on some items. The maximum volume discount (at the highest volume threshold) averages 42%, calculated as a percentage of the undiscounted unit price.

However, few distributors experience price variation as large as 42% on the same item in the same country. As a preliminary comparison we used the transaction data to compare the prices on the largest and smallest orders for the same item in the same country by the same distributor. The average unit price of the largest order is 6.8% less than the average unit price of the smallest order. A simple explanation for the difference between 42% and 6.8% is that distributors do not place orders for the same item at both the minimum order quantity and at quantities that qualify for the largest volume discounts. Instead, there is less variation in distributors' order sizes than the variation in the volume discount schedule.

In our analysis we first use the standard price list to calculate the discount level that transactions qualify for, and then later analyze the actual costs and actual prices paid. Although the standard price list varies by item and region, volume discounts may also vary over time and across distributors within a region. Therefore, these standard discounts should be interpreted as a noisy measure of the size of the volume discounts available to each distributor on each item.

Investigation reveals that the use of volume discounts largely varies by industry. In the energy, enterprise networking, and automotive industries none (or close to none) of the items have volume discounts. We speculate that the difference could reflect bargaining power differences. In industries in which the manufacturer has less bargaining power, prices are closer to the manufacturing cost, and so there is less room to offer discounts. This is consistent with the data that reveals that the manufacturer

earns considerably larger profit margins on products that have volume discounts than on products that do not offer volume discounts.[3]

The manufacturing process suggests that there are relatively few opportunities for scale economies in manufacturing. The raw materials are commodities for which the unit price is largely independent of manufacturing volumes. Many different products are manufactured on the same machines, and there is relatively little set up time required to shift between products. Machine time is therefore approximately proportional to the number of units produced. There may be some small savings from shipping larger quantities. However, these are largely negated by minimum order quantities that ensure that shipping savings are obtained even for relatively small orders. Notably, the unit costs that the manufacturer attributes to each transaction reflect constant unit costs, without reductions for larger quantities. These observations reinforce our conjecture that volume discounts are used primarily for demand-side reasons rather than supply-side reasons.

To measure the impact of aggregate demand shocks we will exploit cross-sectional variation in unemployment rates across countries and time. The unemployment rates in different countries were obtained from the Federal Reserve Economic Data (FRED) site. The data describes seasonally adjusted unemployment rates by quarter for all persons aged 15-64 in each country.[4]

Summary statistics for all of the variables used in the paper are provided in the Appendix, together with a list of the countries represented in the analysis. We replicate the findings from this first dataset of industrial transaction using a second dataset also describing industrial transactions. We discuss this second dataset next.

Industrial Transactions: Distributor Sales to End Customers

The second dataset is a record of onward sales by the five largest distributors (measured in terms of purchases from the manufacturer) to their end customers. These end customer are generally OEM manufacturers that include the components in manufactured products. The transactions involve the same items as the previous dataset.

The data is provided by the distributors to the manufacturer as a condition of doing business with the manufacturer. The manufacturer uses the data to monitor the activities of the distributors. While this monitoring role could create incentives for the distributor to distort the data, the manufacturer does not believe that this occurs. It can monitor the accuracy of the data by asking end customers to disclose their purchase volumes from specific distributors.

The data records transactions between the distributors and their customers and extends from October 1, 2010 to March 25, 2014. It includes over 10.3 million orders and totals $1.3 billion in revenue. The median order size is 15 units and the median unit price is $1.17 (this dataset also treats orders for multiple items as separate orders for each item).

[3] The actual profit margins are confidential and so we are unable to present additional details about this difference.

[4] The data describes the unemployment rate at the start of each quarter and so for each quarter we averaged the unemployment rate at the start and the end of the quarter (the start of the next quarter).

Like the manufacturer sales data, the data describing the transactions between distributors and their customers does not reveal the price in periods without a transaction, or the prices at alternative order volumes. Recall that for the manufacturer sales data we used a regional "standard" price list as a measure of the available volume discounts. To measure the volume discounts offered by the distributors we used a different approach: we scraped the complete price lists from the distributors' websites. All five distributors have websites through which customers can purchase. These websites reveal the complete price schedule at different order volumes. We scraped the data from each distributor's US website, and use this as an indication of the percentage discounts available at different order volumes in every country. However, it is possible that distributors vary their volume discount schedules by region.

In the (scraped) price schedules, 13.5% of the price schedules are a single price, while the remaining 86.5% include volume discounts for larger order volumes. The volume discount schedules include an average of 4.6 pricing levels (for different order volumes), with a maximum of 10 levels on some items. The maximum volume discount (at the highest volume threshold) averages 32%, calculated as a percentage of the undiscounted unit price.

The distributors' sales typically experience much smaller price variation than 32%. We again used the transaction data to compare the prices on the largest and smallest orders for the same item in the same country. In the distributors' sales to their customers the average unit price of the largest order is 11.0% less than the average unit price of the smallest order. This suggests that the distributors' customers generally do not place orders for the same item at both the minimum order quantity and at quantities that qualify for the largest volume discounts.

The possibility that the manufacturer rebates a portion of the transaction price back to their distributors potentially affects the cost of goods sold calculation. Therefore, we again exclude any items for which the manufacturer ever pays a rebate to the distributor (identifying them at the distributor x item x country level). For the reasons described earlier we also remove items with returns.

To measure the impact of aggregate demand shocks we use the same unemployment data that we used to analyze the transactions between the manufacturer and its distributors.

Retail Data

The retail data used in the study was provided by a United States retailer of consumer packaged goods. The retailer sells products in the grocery, health and beauty and general merchandise categories. The data summarizes 195 weeks of store transactions at 102 of the retailers' stores. The stores are located in thirteen states on the East Coast, and the data period extends from the start of 2006 through the end of the third quarter of 2009. The data was previously used by Anderson et al. (2016) to study how the retailer responded to changes in unemployment (in contrast we focus on the customer response).[5] The data reveals the aggregate number of units sold (for each "SKU" or "stock keeping unit") together with the total revenue received. The data excludes price variation due to employee discounts and manufacturer coupons.

[5] Different data from this retailer has also been used in other published studies, including: Ailawada et al. (2007), Anderson et al. (2015) and McShane et al. (2016).

We focus on a sample of 536 items that are sold in two package sizes. These products were identified by manually comparing the SKU descriptions, category descriptions and size descriptions for each SKU. For 376 items we have complete product category descriptions, and these 376 items come from 137 different product categories, ranging from groceries, to personal health care, beauty care, and general merchandise (e.g. toilet tissue). Not all stores sell every item. Across the 536 items and 102 stores there are 33,209 item x store combinations that offer a volume discount, defined as a lower unit price on the larger package size. In addition, there are 52 item x store combinations offer the same unit prices in each package size, and 3,390 combinations that offer a quantity premium, with higher unit prices in the larger package size. We caution that some consumers may not be aware that larger sizes do not always offer unit discounts. This retailer does not print unit prices on its shelf stickers, unless required by law.[6]

A strength of the data is that it also reports the cost of goods sold for each SKU each week. As Anderson et al. (2016) discuss, this measure excludes trade deals and is a "marginal cost", rather than an "average acquisition cost". This cost varies over time but does not vary across geographic regions (stores). Unlike the industrial manufacturing costs where unit costs do not vary according to the order size, the retailer's unit cost often varies across package sizes, with lower unit costs (wholesale prices) on larger package sizes.[7] These cost differences help to reduce the cost to the retailer of offering volume discounts.

The data has two important limitations. First, like the industrial data, we only observe the price of a SKU in a store in weeks that the SKU has sales. The second limitation is that we only observe weekly store sales, rather than sales at the transaction level. This means that we cannot replicate all of the transaction-level analysis conducted with the industrial data.

We match the US retail data with CBSA-level unemployment rates obtained from the Bureau of Labor Statistics Local Area Unemployment Statistics program. In cases where the stores in our main dataset are located in rural areas that are not part of a CBSA for which an unemployment rate is available, we manually match the store with the closest CBSA, and use the unemployment rate for that CBSA.

We begin the analysis in the next section, where we analyze transactions between the industrial components manufacturer and its global distributors.

3. Distributors' Purchases from the Components Manufacturer

In this section we investigate whether distributors place smaller or larger orders in quarters in which unemployment is high. Recall that this analysis is motivated by two conflicting predictions. Higher unemployment could dampen demand leading to smaller orders that qualify for smaller volume discounts. Alternatively, unfavorable economic conditions could motivate distributors to more carefully

[6] Eleven United States territories and states currently have mandatory unit pricing provisions. They are: Connecticut, Maryland, Massachusetts, New Hampshire, New Jersey, New York, Oregon, Puerto Rico, Rhode Island, Vermont and the Virgin Islands. This variation in regulation might suggest an opportunity to investigate how unit pricing requirements affect how customers respond to changes in regional unemployment. Unfortunately the 102 stores in our dataset are in only one of these states.

[7] Here and elsewhere, when we refer to a unit cost we refer to the cost of each unit inside the package (rather than the cost per package). For example a $1.60 price for a 16oz packet of backing soda would have a unit cost of 10 cents per ounce. A $1.20 price for an 8oz packet of baking soda would have a unit cost of 15 cents per ounce.

organize their orders so that they qualify for larger volume discounts. The outcome determines how volume discounts contribute to cyclicality in prices.

The analysis proceeds in three steps. We first measure how changes in unemployment are associated with whether a distributor ordered an item in that calendar quarter. Second, we measure changes in the size of the orders that are placed, including whether the order qualified for a volume discount and how large the volume discount was. Finally, we measure the extent to which volume discounts contribute to changes in the manufacturer's profit margins when unemployment changes.

Whether an Ordered Occurred

Orders for many items are relatively sparse and so our initial response measures describe whether a distributor ordered an item in a quarter:

Any Discounted Order — Equals 1 if distributor *d* ordered item *j* in country *c* in quarter *t* at a volume that qualified for a volume discount; and zero otherwise.

Any Undiscounted Order — Equals 1 if distributor *d* ordered item *j* in country *c* in quarter *t* at a volume that did not qualify for a volume discount; and zero otherwise.

We use the standard regional price list to evaluate which orders qualify for volume discounts.

These variables are constructed to provide an initial evaluation of whether adverse demand shocks prompt distributors to reduce their order sizes and forgo larger volume discounts, or to consolidate their orders to qualify for larger volume discounts. The measures are dependent variables in the following weighted OLS model, where the unit of analysis is a country *c* x item *i* x distributor *d* x quarter *t*:

$$Y_{cidt} = \boldsymbol{\beta}\, \boldsymbol{\Theta}_{cid} + \boldsymbol{\beta}\, \mathbf{Quarter}_t + \beta_1\, Unemployment_{ct} + \beta_2\, Price\ Index_{cdt} + \varepsilon_{cidt} \qquad (3.1)$$

The $\boldsymbol{\Theta}_{cid}$ and $\mathbf{Quarter}_t$ terms identify country x item x distributor and quarter fixed effects (respectively). These fixed effects play an important role. They ensure that the relationship between unemployment and the dependent measures cannot be explained by cross-sectional variation across countries as these are controlled for by the country x item x distributor fixed effects. The quarter fixed effects also ensure that the *Unemployment* coefficient is not identified by common temporal shocks.

The coefficient of interest is β_1, which measures how the dependent variables change as the country-level unemployment rate changes. The inclusion of the fixed effects means that this coefficient is identified by temporal changes in the cross-sectional variation in the unemployment rate. Coibion et al. (2015) use an analogous identification to investigate how retailers respond to regional variation in US unemployment rates (see also Anderson et al. 2016).

The model includes a price index to control for the impact that the manufacturer's price changes had on distributors' use of volume discounts in each country. As a benchmark we first calculate the average price (across the entire data period) that a distributor paid for the same item in the same country at the same discount level. For each transaction we then compare the price paid with this benchmark price,

which yields a measure of how different the transaction price is from the average (benchmark) price. We then aggregate these price change measures across items to create an index for each distributor x country x quarter. Because we only observe prices when there is a transaction, we restrict attention to items for which the distributor placed at least one order every quarter (in that country). This ensures that we use a balanced panel of observations to construct the price index each quarter and prevents the possibility of bias if we do not observe prices when prices are high. Notice also that because we evaluate which orders qualify for discounts using a static price schedule, variation in the dependent variables cannot be explained by the manufacturer changing the volume discount schedule.

The estimation sample is balanced: we include an observation for every quarter for every item x country x distributor that has at least one transaction. The final sample includes 2,341,980 observations, representing 10 quarters (Q4 2011 through Q1 2014), 95 distributors, 26 countries (listed in the Appendix), and 78,051 items. Coefficient estimates are reported in Table 1, where we omit the fixed effects coefficients and report standard errors clustered at the country x quarter level (the level at which the *Unemployment* measure varies). For completeness we also report standard errors clustered at the quarter and country levels in the Appendix.

Table 3.1. Probability of An Order in a Quarter

	Any Discounted Order	Any Undiscounted Order
Unemployment	-0.132% (0.653%)	-0.886%** (0.174%)
Price Index	-12.575%** (4.221%)	-6.924%† (4.049%)
R^2	0.3023	0.2834

The table reports the coefficients from estimating Equation 3.1 using each dependent variable. The unit of analysis is an item sold in a country to a distributor in a quarter. The models include fixed effects but these are omitted from the table. Standard errors clustered at the country x quarter level are reported in parentheses. The sample size in all three models is 2,341,980.

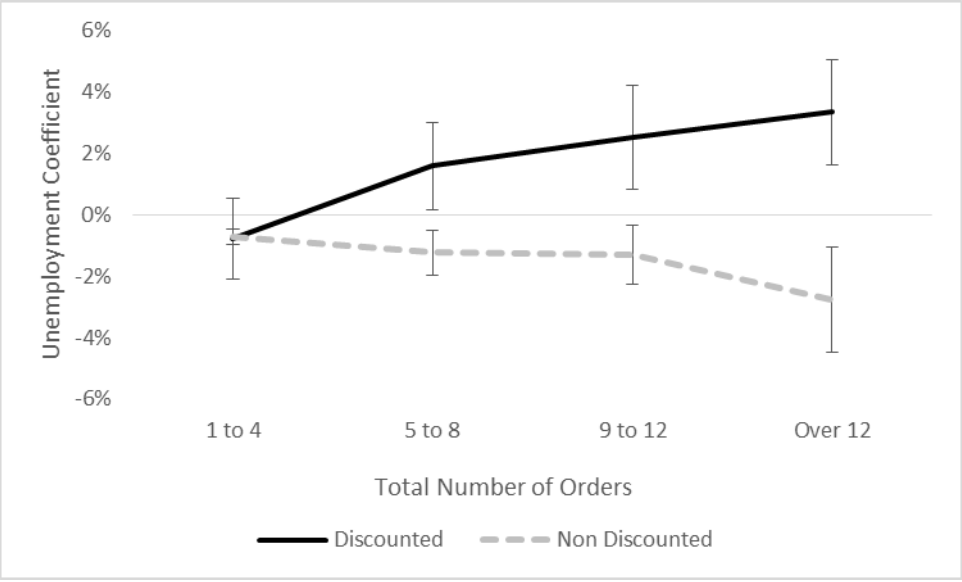
The findings reveal that an increase in unemployment coincides with a significant reduction in the probability of an undiscounted order. A 1% increase in unemployment is associated with a 0.89% reduction that a distributor will place an order in that quarter that does not qualify for a volume discount. However, a 1% increase in unemployment is not associated with a significant change in the probability of a discounted order. This represents initial evidence that in response to adverse macro-economic shocks distributors consolidate their orders to avoid placing orders that do not qualify for volume discounts.

Recall the example of the wine distributor in the Introduction. How did this distributor respond to adverse demand shocks in order to increase its access to volume discounts? It reduced orders for non-core items and consolidated its demand on core items. This is the type of story that we also heard from other distributors. When market conditions are adverse, they focus on their strengths and avoid products in which they are less strong. To investigate this interpretation we will next compare whether distributors increase orders of items they order frequently and decrease orders of items they order infrequently in quarters that unemployment is high (relative to quarters in which it is low).

How do Distributors Change Their Orders When Unemployment Increases?

We re-estimate Equation 3.1 separately for different groups of observations. In particular, we calculate how many orders each distributor placed for each item in each country across the entire data period (*Nbr Orders*$_{cid}$). We then group the observations according to this measure: 1 to 4 orders, 5 to 8 orders, 9 to 12 orders, over 12 orders. We estimated the models separately for each group of observations and report the *Unemployment* coefficients for the discounted and undiscounted orders in Figure 3.1 (detailed results including samples sizes are reported in the Appendix).

Figure 3.1 Probability of an Order in the Quarter: Core vs. Non-Core Items



The figures report the *Unemployment* (β_1) coefficients from estimating Equation 3.1 on different groups of items. Errors bars indicate 95% confidence intervals (with standard errors clustered at the country x quarter level). Detailed results including samples sizes are reported in the Appendix.

Higher unemployment is associated with a significantly lower probability of some types of orders and a higher probability for others. For items that are ordered frequently, the probability of a discounted order increases in quarters with high unemployment, while the probability of an undiscounted order

decreases.[8] This is precisely the pattern we would anticipate if distributors react to an adverse demand shock by consolidating their purchases of core items to increase the volume discounts that they qualify for.

This outcome is most pronounced for the most frequently purchased items. From a practical perspective, consolidating orders is easier on items that are ordered more frequently. Consolidating multiple orders for the same item requires either over-ordering in anticipation of future demand, or under-ordering while waiting for future demand. If demand is infrequent the wait for future demand will be longer, and the costs associated with under-ordering or over-ordering will be higher.

In this initial analysis we investigated the probability that distributors placed an order in the quarter. In our next analysis we investigate how changes in unemployment affected the average size of the orders that distributors placed.

Average Order Size

We use three measures to characterize the average size of distributors' orders. The first two measures describe the frequency and magnitude of the volume discounts that distributors receive:

Proportion Discounted$_{dct}$	The proportion of transactions that qualified for a volume discount for distributor *d* in country *c* in quarter *t*.
Average Discount$_{dct}$	The average volume discount that transactions qualified for.

The *Proportion Discounted* and *Average Discount* are constructed using the standard price list (which is static). Therefore, variation in these measures results solely from changes in the orders, not from the manufacturer changing the discount schedule. Changes in distributors' orders that could affect the *Proportion Discounted* and *Average Discount* measures include changes in which items distributors order, and/or changes in the size of the orders that they place. We previously investigated which items distributors order and how that is affected by variation in unemployment. We can also study how the size of the orders varies. We do so using the following measure:

Order Size$_{dct}$	The size of the order relative to the median order size for item *i* distributor *d* in country *c*.[9]

We use these three measures as dependent variables in the following weighted OLS model:

[8] Notice that β_1 is identified by variation within an item. Therefore, the positive interaction for discounted orders does not simply reflect a higher probability of ordering items when *Nbr Orders* is large. Instead, a more accurate interpretation is that for items that the distributor orders frequently, the probability of an order is larger in quarters that unemployment is high compared to quarters in which it is low.

[9] The *Order Size* for transaction *z* is calculated as: $\frac{Order\ Quantity_z - Median\ Order\ Quantity_{cid}}{0.5*Order\ Quantity_z + 0.5*Median\ Order\ Quantity_{cid}}$. This calculation ensures that orders that are larger and smaller than the median order size are treated symmetrically. In the Appendix we also report findings when calculating the *Order Size* using the average order quantity as the benchmark.

$$Y_{cdt} = \boldsymbol{\beta}\,\boldsymbol{\Theta}_{cd} + \boldsymbol{\beta}\,\textbf{Quarter}_t + \beta_1\, Unemployment_{ct} + \beta_2\, Price\ Index_{cdt} + \varepsilon_{cdt} \quad (3.2)$$

The unit of analysis is a country x distributor quarter and so the fixed effects are at the country x distributor ($\boldsymbol{\Theta}_{cd}$) and quarter levels. Similar to Equation 3.1, the fixed effects capture cross-sectional (country x distributor) and temporal (quarter) effects. When estimating the probability of an order (Equation 3.1) we explicitly controlled for manufacturer price changes. For completeness, we also include the *Price Index* in Equation 3.2. The earlier evidence that distributors respond differently on frequently versus infrequently ordered items also suggests that it is important to weight the observations. We weight each observation by the distributor's revenue in that country that quarter. In the Appendix we also report findings using static weights (A*verage Quarterly Revenue* calculated across the data period). The findings are reported in Table 3.2, where the standard errors are clustered at the country x quarter level. We report alternative clusters in the Appendix.

Table 3.2. Average Order Size

	Proportion Discounted	**Average Discount**	**Order Size**
Unemployment	2.061%**	0.660%**	2.611%**
	(0.437%)	(0.134%)	(0.611%)
Price Index	3.792%	2.380%	-6.115%
	(7.519%)	(2.911%)	(8.512%)
R^2	0.8896	0.9091	0.7953

The table reports the coefficients from estimating Equation 3.2. Fixed effects were estimated but are omitted from the table. Standard errors clustered at the country x quarter level are in parentheses. The sample size in all three models is 1,060. The observations are weighted by revenue.

The findings confirm that distributors' orders qualify for larger volume discounts in quarters with higher unemployment. A 1% increase in unemployment is associated with a 2.06% increase in the proportion of orders that qualify for discounts, and a 0.66% increase in the average size of the discounts that the orders qualify for. Our earlier findings suggest that these effects occur for two reasons. For items that are ordered infrequently, distributors place fewer orders in quarters that unemployment is high (compared to quarters in which it is low). This increases the average discount because the sample includes fewer orders that do not qualify for large discounts. For the frequently purchased, items distributors consolidate their orders so that more orders qualify for volume discounts.

Both effects imply consolidation, indicating that order sizes should be larger. The *Order Size* result confirms this. A 1% increase in unemployment is associated with a 2.61% increase in the *Order Size*

(measured relative to the median order size). The implication is that volume discounts lead to distributors paying lower prices during periods of higher unemployment.

In the Appendix we report several robustness checks. First, perhaps unsurprisingly, there is a positive pair-wise correlation between revenue in the quarter and average order size in the quarter (distributors place larger orders in quarters in which total revenue is higher). As a result, the weighting factors may be amplifying the estimated effects; increases in the order size are magnified by the greater weight given to larger revenue. Although this arguably makes the weighting of the observations even more appropriate, for completeness we also estimate the findings when using static weights (*Average Quarterly Revenue* calculated across the entire data sample). The pattern of findings is unchanged.

The second robustness check uses all of the items to calculate the *Proportion Discounted* and *Average Discount* for that distributor x country x quarter, instead of just those items for which volume discounts are available (in the standard price list). For the items without a volume discount the *Proportion Discount* and *Average Discount* are always zero (by construction).[10] The pattern of findings survives and is slightly strengthened when including these items.

Finally, we re-calculate the *Order Size* measure using the average order size rather than the median order size. The findings are even stronger when using this approach.

While these findings confirm that buyer behavior contributes to lower prices when unemployment increases, they do not reveal how large this effect is relative to manufacturer price changes (seller behavior). The findings are also based upon the standard price list, which may be different from the prices that distributors actually pay. In the next analysis we address both limitations.

Change in Seller Profit Margins

In the final analysis in this section we investigate how much volume discounts contribute to price cyclicality. Our analysis of how individual distributors change their orders during quarters with higher unemployment reveals evidence that they substitute between items. Therefore, to evaluate the contribution of volume discounts to price cyclicality we need to account for this type of substitution. One potential concern is that items may offer different "value"; although two items may have the same unit price, the manufacturer may earn different profit margins on them. Therefore, in order to evaluate price variation between items we need a basis for comparing value across items. To do so we will focus on variation in profit margins instead of variation in prices.

The profit from transaction *z* can be written as: $Profit_z$ = $Profit\ Margin_z$ * $Revenue_z$ where:

$$Profit\ Margin_z = \frac{Price_z - Cost_z}{Price_z}.$$

[10] The inclusion of these items could potentially increase or decrease the estimated *Unemployment* coefficient. The outcome depends upon which quarters customers purchase these items. If they tend to purchase them in quarters that unemployment is low (high) this will tend to make the *Unemployment* coefficient more (less) negative.

We will investigate how the *Profit Margin* varies as unemployment varies, when weighting each transaction by the transaction *Revenue*. In particular, we estimate the following weighted OLS model:

$$\textit{Profit Margin}_z = \boldsymbol{\beta}\, \boldsymbol{\Theta}_{cd} + \boldsymbol{\beta}\, \textbf{Quarter}_t + \beta_1\, \textit{Unemployment}_{ct} + \varepsilon_z \qquad (3.3)$$

The unit of analysis is a transaction, and the fixed effects are at the country x distributor ($\boldsymbol{\Theta}_{cd}$) and quarter levels. Similar to the two previous sets of analysis, the fixed effects capture cross-sectional (country x distributor) effects, and quarterly (temporal) effects. We report the findings of this model in the Appendix. The findings reveal that a 1% increase in unemployment in a country x quarter is associated with a 1.29% decrease in the manufacturer's profit margin on orders in that country x quarter. We conclude that in periods in which unemployment is high the effective price paid by distributors is lower.

There are three possible explanations for this effect:

Seller Changes

1. Seller Changes: Variation in the manufacturer's prices (or costs).

Buyer Behavior

2. Substitution Between Items: Variation in which items distributors order.
3. Variation in Order Quantities: Variation in distributors' order quantities for the same item, leading to price changes along the volume discount schedule.

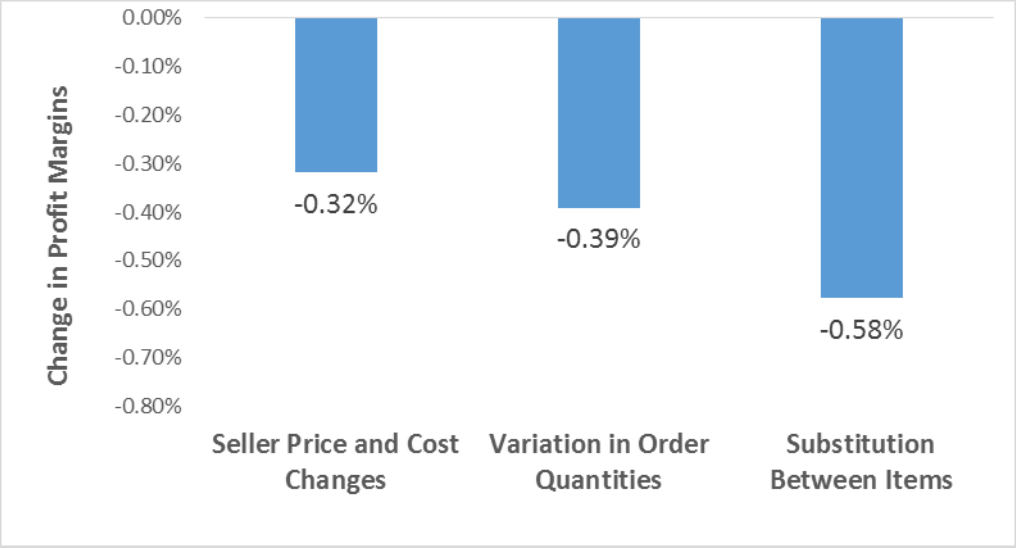
We disentangle these three effects by averaging to isolate the source of variation in the seller's (manufacturer's) profit margin. In the first step we average out variation over time and across countries at the distributor x item level, but allow variation across different purchase quantities. This removes the manufacturer sources of variation (identified as effect 1 above). In the second step we average out variation over time, countries and purchase quantities for a distributor x item. This incremental change removes variation due to substitution between order quantities (effect 2 above). The only remaining variation in profit margins is due to substitution between items (effect 3 above). A more extensive discussion of this analysis together with detailed findings are provided in the Appendix. We summarize the findings in Figure 3.2.

When unemployment increases by 1%, substitution by distributors between items contributes 0.58% of the total 1.29% reduction in profit margins. Substitution between order quantities contributes another 0.39% reduction, while changes in the manufacturer's (seller's) costs and prices contribute the remaining 0.32% decrease.

Because the manufacturers' costs do not vary with purchase quantities, the only source of variation in profit margins when quantities change is variation in prices. Therefore, changes in the manufacturer's profit margins due to variation in order quantities are fully attributable to volume discounts. In the absence of volume discounts, prices would not change when purchase quantities change. Under this

interpretation, when unemployment increases by 1%, approximately a third of the overall reduction in manufacturer profit margins is attributable to volume discounts.

Figure 3.2 Decomposing the Change in the Manufacturer's Profit Margins



The figure reports the incremental contribution of each effect to the change in the manufacturer's profit margins when unemployment increases by 1%. Details of the analysis and the findings are provided in the Appendix.

This is a conservative estimate of the total contribution of volume discounts to the reduction in the seller's profit margins. Volume discounts may also contribute to distributors' decisions to substitute between items, and the results in Figure 3.2 indicate that this substitution further contributes to the reduction in manufacturer profit margins. In particular, it is possible that at least some of the substitution between items is due to distributors consolidating their orders to benefit from larger volume discounts. However, there are other reasons that a distributor may choose to substitute between items when unemployment increases (including the possibility that demand for different items changes). As a result, we cannot attribute substitution between items solely to the availability of volume discounts.

Summary

We might have expected that when unemployment increases, demand would go down, and that this would lead to smaller orders and smaller volume discounts. Our findings reveal that the reverse is true. During quarters with higher unemployment, distributors organize their orders so that they qualify for larger volume discounts, and this reduces the effective price that distributors pay.

The changes in distributors' behavior are reflected in both the size of their orders and in the probability that they place an order. On their core items that they purchase frequently, distributors are more likely to place orders that qualify for discounts when unemployment is high and less likely to place orders that

do not qualify for discounts. Increases in unemployment are also associated with; a significant increase in average order sizes, an increase in the proportion of orders that qualify for volume discounts, and an increase in the average discount that orders qualify for.

We conclude that volume discounts contribute to cyclicality in paid prices. This contribution comes in two forms: lower prices due to higher order volumes, and substitution from high profit margin items to low profit margin items. The first effect, which can be solely attributed to volume discounts, contributes approximately a third of the decrease in manufacturer profit margins when unemployment increases. The second effect, substitution between items, is not solely attributable to volume discounts. Distributors may substitute and purchase items that earn lower profit margins for the manufacturer irrespective of volume discounts. However, consolidating orders to earn larger volume discounts provides the distributors with an additional incentive to engage in this practice.

Regressions of prices on measures of economic activity are potentially confounded if productivity is correlated with unemployment rates. However, in this section we have focused on distributors' orders rather than the manufacturer's prices. [11] It is possible that productivity changes could contribute to changes in distributors' order. This might help to explain why distributors are able to consolidate their orders in the way that we document. Although we acknowledge this possibility, we caution that a relatively elaborate productivity effect would be required to explain why distributors' orders vary with unemployment in the manner described in this section.

Distributors sometimes purchase items in one country and then sell them in another country. This might occur if they have a regional warehouse for some parts that they use as a base for shipping to multiple countries in that region (for example shipping from a Spanish warehouse to Portuguese customers). We do not believe that this provides an alternative explanation for the findings. Instead, it seems that this a source of noise in the data, making it more difficult to identify a relationship between unemployment in a country and purchases by distributors in that country (the findings occur despite not because of this noise).

Because the data involves a single manufacturer we might be concerned about the generalizability of the findings. In defense of the findings we note that we focus on purchasing decisions (rather than the manufacturer's decisions) and the data includes purchasing decision by a large number of distributors together with items from a broad range of industries. We also explicitly address concerns about generalizability in the next two sections. In Section 4 we focus on sales by the distributors to their customers and again study variation in unemployment across countries. In Section 5 we focus on retail data, and investigate how retail customers respond to higher regional unemployment in the United States.

[11] We also note that in our profit margin analysis we control for changes in manufacturer costs, which would include changes in the manufacturer's productivity. Moreover, for approximately 47% of the item x country pairs represented in the analysis, the items are not manufactured in the countries in which they are sold. Instead, the items are manufactured in one country and shipped to the other country. Variation in manufacturing productivity in one country is unlikely to explain variation in distributors' orders in another country.

4. Distributors' Sales to their Customers

In the previous section we investigated whether industrial distributors' respond to higher unemployment by placing smaller or larger orders. Surprisingly, in quarters with higher unemployment, we show that distributors place larger orders that qualify for larger volume discounts. In this section we investigate whether this finding replicates when we study onward transactions between distributors and their end customers. These end customers are almost all product manufacturers who incorporate the components into their products. There are two notable features that are common between the manufacturer sales to distributors and distributors' sales to their customers. First, the distributors are common between the datasets, so that the sellers in this data are buyers in the data studied in the previous section. Second, the items sold in these transactions are the same items that we studied in Section 3.

We begin by reporting how the probability of a discounted or undiscounted order in the quarter varies with the unemployment rate. The unit of analysis is a country *c* x item *i* x distributor *d* x end customer *e* x quarter *t* and we estimate the following weighted OLS model:

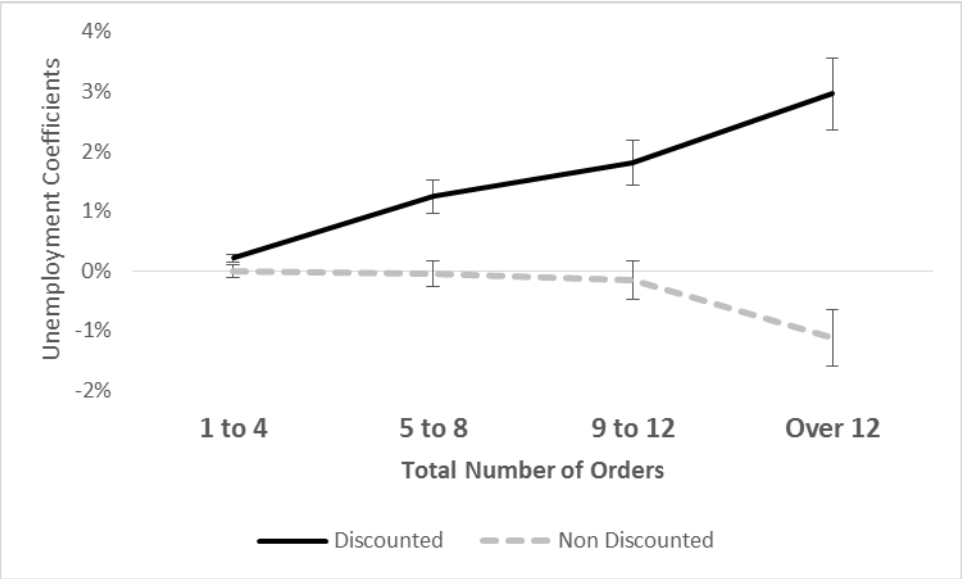
$$Y_{cidet} = \boldsymbol{\beta}\, \boldsymbol{\Theta}_{cid} + \boldsymbol{\beta}\, \mathbf{Quarter}_t + \beta_1\, Unemployment_{ct} + \beta_2\, Price\ Index_{cdt} + \varepsilon_{cidet} \qquad (4.1)$$

This equation is analogous to Equation 3.1. The $\boldsymbol{\Theta}_{cide}$ and $\mathbf{Quarter}_t$ terms identify country x item x distributor and quarter fixed effects (respectively). These fixed effects play the same role they played in the previous analysis. They ensure that the relationship between unemployment and the dependent measures cannot be explained by cross-sectional variation across countries or by common temporal shocks.

The estimation sample is again balanced: we include an observation for every quarter for every item x country x distributor x end customer that has at least one transaction. The sample includes 41,400,840 observations, representing 13 quarters (Q1 2011 through Q1 2014), 5 distributors, 515,266 customers, 31 countries (listed in the Appendix), and 32,546 items.[12] The findings are summarized in Figure 4.1, where we summarize the outcome according to how frequently customers purchase the item. Complete findings are provided in the Appendix.

[12] The six countries represented in this data that do not appear in the manufacturer sales data are Chile, Czech Republic, Estonia, Hungary, Portugal and Slovakia. These items are purchased by distributors in other countries and shipped to customers in these countries.

Figure 4.1 Probability of an Order in the Quarter: Core vs. Non-Core Items
Distributor Sales



The figures report the *Unemployment* (β_1) coefficients from estimating Equation 4.1 on different groups of items. Errors bars indicate 95% confidence intervals (with standard errors clustered at the country x quarter level). Detailed results including samples sizes are reported in the Appendix.

The findings in Figure 4.1 (distributor sales) are remarkably similar to Figure 1 (distributor purchases). Higher unemployment is associated with a higher probability of some types of orders, and a lower probability of others. In quarters in which unemployment is high there is a significant increase in the probability of orders that qualify for discounts, but a decrease (or no change) in the probability of orders that do not qualify for discounts. The difference is again most pronounced for the core items that customers order most frequently. When macroeconomic conditions are unfavorable customers appear to be more careful to organize their frequently ordered items so that they qualify for volume discounts.

We next report how the *Average Discount*, *Proportion Discounted* and *Order Size* vary with the *Unemployment Rate*.[13] The unit of analysis is a country *c* x distributor *d* x quarter *t* and we estimate the following weighted OLS model (which is analogous to Equation 4.2):

$$Y_{cdt} = \boldsymbol{\beta}\, \boldsymbol{\Theta}_{cd} + \boldsymbol{\beta}\, \textbf{Quarter}_t + \beta_1\, Unemployment_{ct} + \beta_2\, Price\ Index_{cdt} + \varepsilon_{cdt} \qquad (4.2)$$

The fixed effects are at the country x distributor ($\boldsymbol{\Theta}_{cd}$) and quarter levels. The fixed effects capture cross-sectional (country x distributor) and temporal (quarter) effects. We weight each observation by the distributor's revenue in that country that quarter. The findings are reported in Table 4.1, where the

[13] When calculating the *Order Size* in the distributor sales data we use the median order size for item *i* from distributor *d* in country *c*.

standard errors are clustered at the country x quarter level. In the Appendix we also report the findings with alternative clusters.

Table 4.1. Average Order Size: Distributor Sales

	Proportion Discounted	Average Discount	Order Size
Unemployment	1.420%**	0.192%**	6.775%**
	(0.191%)	(0.052%)	(0.843%)
Price Index	-45.438%**	-7.043%*	-214.408%**
	(7.657%)	(2.787%)	(35.288%)
R^2	0.8852	0.9773	0.3098

The table reports the coefficients from estimating Equation 4.2. Fixed effects were estimated but are omitted from the table. Standard errors clustered at the country x quarter level are in parentheses. The sample size in all three models is 1,105. The observations are weighted by revenue.

Recall that orders by the distributors qualify for larger volume discounts in quarters with higher unemployment. We see the same pattern for sales by the distributors. A 1% increase in unemployment is associated with a 1.42% increase in the proportion of orders that qualify for discount, and the average discount that an order qualifies for increases by 0.19%. The average order size is also larger in quarters with higher unemployment. Recall that the *Order Size* measures the order size relative to the median order size for that item from that distributor in that country. A 1% increase in unemployment is associated with a 6.78% increase in the *Order Size*.

In the Appendix we report the same robustness checks that we used in Section 3. First, we re-estimate the models using static weights. This yields a very similar pattern of outcomes. The next robustness check uses all of the items to calculate the *Proportion Discounted* and *Average Discount* for that distributor x country x quarter, instead of just those items for which volume discounts are available (in the scraped price list).[14] We also re-calculate the *Order Size* measure using the average order size (for item *i* from distributor *d* in country *c*) as the benchmark for comparison, instead of the median order size. The findings are robust to both of these alternative approaches.

We finish this section by investigating how the distributors' profit margins changed when unemployment increased. Unlike the manufacturer sales data, the distributor sales data does not include the distributors' cost of goods sold (the unit price they paid to the manufacturer). However, we

[14] Recall that for the observations without a volume discount the *Proportion Discount* and *Average Discount* are always zero (by construction) and that the inclusion of these observations could potentially increase or decrease the estimated *Unemployment* coefficient (see earlier discussion).

can calculate the cost of goods sold from the dataset describing the manufacturer's sales to these distributors.[15]

The distributors' purchases from the manufacturer do not coincide with their sales of the same inventory. As a result, it is not possible to accurately evaluate how the cost of goods sold in the distributors sales vary over time, across countries or across order quantities:

Time: sales inventory could come from different distributor purchases.

Countries: distributors' purchases of inventory may not occur in same the countries that they sell the inventory.

Order Quantities: the distributors often sell in different quantities than they purchase.

For this reason we cannot construct the same decomposition of the distributors' profit margins as we constructed for the manufacturer's profit margins. Instead, we calculate a static item cost for each distributor by aggregating their purchases for that item over time, countries and purchase quantities. We then use this static cost of goods sold in our profit margin calculations. This eliminates any variation in the distributors' profit margins due to variation in the cost of goods sold. However, we can still evaluate other sources of changes in the distributors' profit margins, including distributors' price changes, end customers varying their order quantities, and end customers substituting between items.

We summarize this decomposition in Figure 4.2, and provide details of the models in the Appendix. Because we use a static cost measure, the findings are not directly comparable to the findings reported in Section 3 for the manufacturer sales data. However, we see several similarities in the findings. Higher unemployment is associated with lower profit margins for the distributors. In particular, a 1% increase in unemployment is associated with a 1.71% decrease in the distributors' profit margins. Most of this effect comes from substitution between items, with variation in order quantities and seller price changes contributing only a relatively small portion of the overall effect.

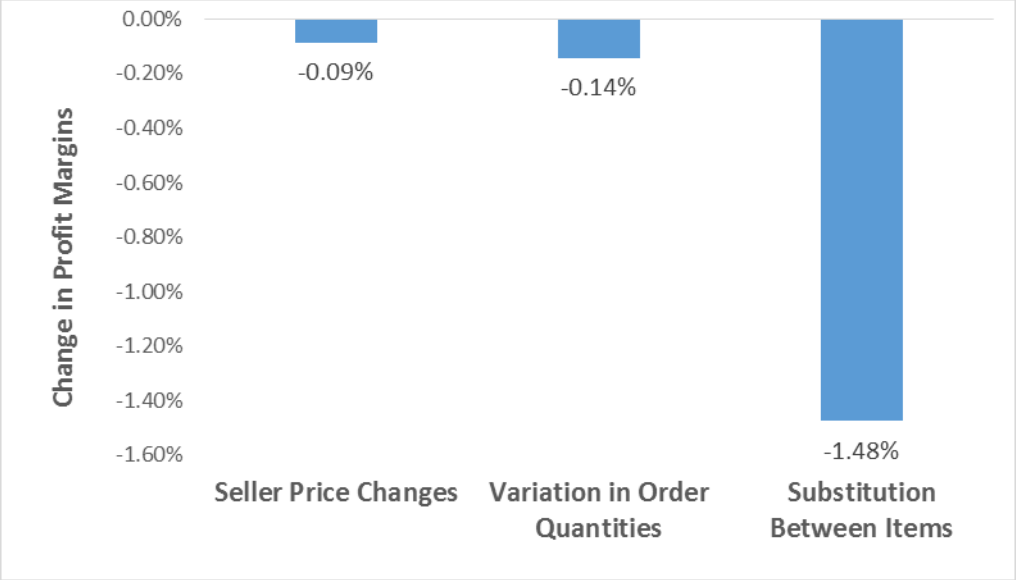
Summary

The difference between the findings in this section and the previous section is that in this section we studied sales by the distributors, while previously we studied purchases by the distributors. Despite this difference we see the same pattern of findings. In quarters with higher unemployment we might have expected that the distributors' customers would place smaller orders. We see the reverse; they place larger orders that qualify for larger volume discounts when unemployment increases. This leads to lower distributor profit margins and contributes to cyclicality in paid prices. As with the manufacturer sales, the reduction in profit margins largely result from customers substituting to lower margin items in periods in which unemployment is high.

[15] This matching process results in some missing observations because distributors sometimes assign different item numbers than the manufacturer's item numbers. This means that it is not possible to accurately match the distributors' sales data with the data describing their purchases from the manufacturer. To validate the process of matching item numbers the distributors' sales and purchases, we exclude transactions where the cost of goods sold in the purchasing dataset exceeds the sales price in the sales dataset. This is an indication that the matching between the sales and purchasing datasets may be subject to error (and affects approximately 8.4% of the transactions).

These findings help to confirm that the findings in the previous section generalize when we change the identity of the industrial buyers and sellers. We further investigate generalizability in the next section, where we focus on retail data, and investigate how retail consumers respond to higher regional unemployment in the United States.

Figure 4.2 Decomposing the Change in Distributors' Profit Margins



The figure reports the incremental contribution of each effect to the change in the distributors' profit margins when unemployment increases by 1%. The changes in the profit margins are calculated using static cost measures. Details of the analysis and the findings are provided in the Appendix.

5. Replicating the Findings Using Retail Transactions

In the previous two sections we investigated whether industrial purchasers respond to higher unemployment by placing smaller or larger orders. Surprisingly, in quarters with higher unemployment, we show that they place larger orders that qualify for larger volume discounts. In this section we investigate whether this finding replicates in retail markets.

There are at least three important differences between the retail setting and the industrial markets studied in sections 3 and 4. First, we do not have transaction level data. Instead, we will analyze aggregate sales at the item x store x week level. Second, while we again use a combination of temporal and cross-sectional variation in unemployment, the cross-sectional variation is regional variation in unemployment within the United States, instead of variation across countries. Third, only a small portion of the items sold by this retailer offer quantity discounts. We will focus on the 536 items that are sold in both small and large package sizes. Although they encompass a wide range of categories, they

represent less than 3% of the products that the retailer sells. As a result, it is not meaningful to measure the overall impact of volume discounts on this retailer's profit margin. The limited range of products that have volume discounts means that it is also not meaningful to investigate switching between items. Instead, we focus on estimating how changes in unemployment affect the choices of small versus large package sizes.

Recall from our discussion of the data (Section 2) that the data describes aggregate sales at the SKU level in 102 retail stores for 195 weeks. From this weekly data we construct two measures that are analogous to those used in the previous section:

Proportion Discounted$_{ist}$	The number of large sizes of item *i* sold in store *s* in week *t* divided by the total number of items (of both sizes) sold in that store x week.
Average Discount$_{ist}$	The average discount received for purchases of item *i* in store *s* in week *t*.

Calculation of the *Average Discount* is detailed in the Appendix. Our focus is on the customer response. Therefore, to ensure that variation in these measures is not explained by variation in the prices themselves, throughout this section we use static prices and costs (the price and cost for each package size in each store averaged across the entire data period). These measures are used as dependent variables in the following weighted OLS model (analogous to Equations 3.2 and 4.2 in the industrial data):

$$Y_{ist} = \boldsymbol{\beta}\, \boldsymbol{\Theta}_{is} + \boldsymbol{\beta}\, \mathbf{Week}_t + \beta_1\, Unemployment_{st} + \beta_1\, Price\ Index_{ist} + \varepsilon_{ist} \qquad (5.1)$$

The unit of analysis is an item x store x week. The fixed effects again control for temporal effects and cross-sectional variation (at the item x store level).[16] The estimation sample is necessarily restricted to store x weeks in which at least one size of the item had sales (the dependent variables are otherwise undefined). Like our earlier analysis, the *Price Index* is calculated using deviations from the average price across the data period.

We initially focus on item x store combinations that offer volume discounts (the unit price is lower on the large package size). However, we will also check robustness by re-estimating the findings using all of the observations, including items with quantity premiums (higher unit prices on the large package sizes). We weight the observations by revenue calculated at the item x store x week level, and report findings in Table 5.1. The standard errors are clustered by CBSA x month, which is the level of variation in the *Unemployment* variable. In the Appendix, we also report the findings when clustering the standard errors by store, month, or store x month.

[16] Recall that in this analysis our focus is on switching between packaging sizes of the same item, rather than switching between items.

Table 5.1. Retail Data Results

	Proportion Discounted	Average Discount	Profit Margin
Unemployment	0.199%*	0.056%**	-0.024%**
	(0.081%)	(0.020%)	(0.006%)
Price Index	1.209%**	27.726%**	
	(0.326%)	(0.082%)	
R^2	0.6851	0.7242	0.9689

The table reports the coefficients from estimating Equation 5.1 (*Proportion Discounted* and *Average Discount*) and Equation 5.2 (*Profit Margin*). Fixed effects were estimated but are omitted from the table. Standard errors clustered at the CBSA x month level are in parentheses. The sample size in all three models is 2,110,952. The observations are weighted by revenue calculated at the item x store x week level.

The evidence that during periods of higher unemployment customers consolidate their orders to purchase in larger quantities extends beyond industrial markets to also include retail markets. A 1% increase in unemployment is associated with a 0.20% increase in the proportion of sales that are discounted and a 0.056% increase in the average discount. This replicates our earlier findings: customers appear to organize their purchases so that more of them qualify for volume discounts when economic conditions are unfavorable.

We repeated the same robustness checks that we conducted for the industrial components analysis. These results are provided in the Appendix. First, we repeated the analysis when including the item x store combinations that did not offer volume discounts (items with quantity premiums). As we might expect, inclusion of these observations resulted in slightly smaller effects.[17] Second, we re-estimated the findings when using static weights. The magnitude and significance of the findings are strengthened when using static weights.

We can also investigate the resulting impact on the retailer's profit margins. We used an equation analogous to Equations 3.3 and 4.3:

$$Y_{ist} = \boldsymbol{\beta}\, \boldsymbol{\Theta}_{is} + \boldsymbol{\beta}\, \mathbf{Week}_t + \beta_1\, Unemployment_{st} + \varepsilon_{ist} \qquad (5.2)$$

[17] We caution that there are too few of these items to draw reliable conclusions about the consumer response to items that offer quantity premiums (higher unit prices on larger package sizes). We also note that it is possible that some consumers do not calculate unit prices and so do not recognize when a larger size does not offer a volume discount (recall that unit prices are generally not printed on this retailer's shelf stickers)

Calculation of the *Profit Margin* is detailed in the Appendix and the findings are also reported in Table 5.1. The *Proportion Discounted* and *Average Discount* results indicated that retail customers respond to adverse economic conditions by organizing their purchases so that more of them qualify for volume discounts. This is reflected in the retailer's profit margins. A 1% increase in unemployment is associated with a 0.24% decrease in the retailer's profit margins. We calculated the retailer's profit margin using static prices and costs (see the discussion in the Appendix) and so this effect is fully attributable to variation in which items retail consumers purchase, and the package sizes that they select.

In our analysis of the industrial data we compared the response to changes in unemployment for items that distributors ordered frequently and those that they ordered infrequently. Unfortunately, the absence of transaction level data means that this comparison is not possible with our retail data. However, we can evaluate how these aggregate findings vary across items. In particular, we might expect that the response would be larger on products that tend to be purchased by more price sensitive customers. The marketing literature recognizes that private label products generally attract more price sensitive customers than national brand products.[18] This would suggest that we should see larger effects for private label products than for products that have national brands.

Private Label vs. National Brand Items

The estimation sample includes 345 items that have national brands and 140 items that have the retailer's own store brand (there are a small number of items for which brand information is not available). We re-estimate Equation 5.1 separately for each sample of items, and summarize the resulting *Unemployment* coefficients in Table 5.2.

The findings reveal clear differences. For private label items a 1% increase in unemployment is associated with a 0.369% increase in the proportion of large sizes purchased. For national brand items we do not observe a statistically significant change. Similarly, the *Average Discount* increases by 0.094% for private label items, but only 0.039% for national brand items.

These findings might suggest that increases in unemployment would lead to larger reductions in the retailer's *Profit Margin* for private label items than for national brand items. Surprisingly, the reduction in profit margins is similar on national brand and private label items. Further investigation reveals that for some national brand items there is a lot of substitution from small to large sizes during periods of high unemployment. Unfortunately for the retailer, this substitution tends to occur on the items for which the retailer's unit profit margin is a lot lower on the large sizes (compared to the small sizes).[19]

[18] For example, Hansen, Singh, and Chintagunta (2006) report that household tendency to buy store brands is correlated across categories, and is also positively correlated with price sensitivity. Hoch (1996) reports that private labels are more prevalent (i.e. have higher shares) in stores with more price-sensitive consumers. See also Hoch, Kim, Montgomery, and Rossi (1995).

[19] Further investigation reveals that higher unemployment is associated with more substitution to large sizes on national brand items for which the large size has a big size advantage (many more units than the small size) and a little price disadvantage (the large package price is closer to the price of the small size). This is consistent with customers substituting during periods of higher unemployment to take advantage of larger available volume discounts. However, we caution that the size advantage and price disadvantage are endogenous. As a result, it is difficult to know whether these interactions reflect customers taking advantage of larger available discounts versus the factors that explain the variation in the size of the volume discounts.

Table 5.2. Retail Data: Private Label vs. National Brand Comparison

	Private Label	National Brand	Difference
Proportion Discounted	0.369%**	0.123%	0.246%**
	(0.107%)	(0.083%)	(0.090%)
Average Discount	0.094%**	0.039%†	0.056%†
	(0.033%)	(0.019%)	(0.029%)
Profit Margin	-0.019%**	-0.024%**	0.005%
	(0.007%)	(0.007%)	(0.010%)
Sample Size	768,255	1,297,162	2,065,417

The table reports the (β_1) coefficients from estimating Equation 5.1 (*Proportion Discounted* and *Average Discount*) and Equation 5.2 (*Profit Margin*). Fixed effects were estimated but are omitted from the table. Standard errors clustered at the CBSA x month level are in parentheses. The observations are weighted by revenue calculated at the item x store x week level. Missing observations reflect items for which no brand information is available.

Like our analysis of the industrial components, our analysis of the retail data is limited to a single seller. This may raise concerns that the findings will not generalize to other retailers. However, this limitation overlooks the purpose of this retail analysis, which was to demonstrate generalizability. In particular, it is reassuring that we were able to replicate evidence that customers organize their purchases so that they receive larger volume discounts in periods in which unemployment is high. Notably, this replication occurs using very different products (consumer packaged goods instead of industrial components), and a different source of variation in unemployment (regional variation within the United States rather than across countries). We also note that while we use data from a single retailer, the consumer behavior that we document reflects the purchasing decisions of many consumers.

A recent paper has raised another potential limitation for papers that use regional variation to investigate aggregate effects. Beraja, Hurst and Ospina (2015) argue that different shocks could drive cross-sectional variation than time-series variation, and this might lead to different outcomes. However, for at least three reasons this limitation is much less relevant to this paper than to other papers that rely upon regional variation. First, we do not claim that the variation in volume discounts contributes to important cyclicality in retail prices. Products with volume discounts represent only a small fraction of the products that this retailer sells, and so it is not plausible to claim that the aggregate effects in this retail setting lead to meaningful cyclicality. Instead, we simply conclude that when unemployment increases, retail consumers exhibit behaviors similar to industrial distributors: they organize their purchases in order to qualify for more volume discounts. Second, Beraja, Hurst and Ospina (2015) focus on sources of price variation introduced by sellers. Instead, we focus on how buyer decisions contribute to changes in the prices that they pay. Third, our analysis of industrial prices in Sections 3 and 4 exploits variation in unemployment measured at the national level (not the regional level).

Summary

In this section we investigated how retail consumers use volume discounts in response to changes in unemployment. When unemployment is high consumers purchase a higher proportion of large package sizes, leading to larger volume discounts. This contributes to a reduction in retailer profit margins, and replicates the pattern of findings that we reported for the industrial distributors in the previous section.

We showed that this effect is stronger on private label items than on national brands. This is consistent with evidence in the marketing literature that private label products tend to be purchased by consumers who are more price sensitive.

6. Conclusions

We have studied how volume discounts contribute to cyclicality by measuring how the prices that industrial distributors and retail consumers pay vary as unemployment varies. A straw man argument might predict that when unemployment increases then demand will contract, leading to smaller orders that qualify for smaller volume discounts. We show that the opposite happens. In periods with higher unemployment industrial purchasers consolidate their purchases into larger orders that qualify for larger volume discounts. Similarly, retail consumers are more likely to purchase larger package sizes when unemployment is high. In both settings this results in lower effective prices and lower profit margins for the sellers.

The use of multiple datasets allows us to replicate the key finding that higher unemployment leads to purchasers organizing their purchases so that they qualify for larger volume discounts. This replication occurs in very different product markets, and exploits different sources of variation in unemployment. The industrial data exploits variation in unemployment across countries and over time, while the retail analysis uses variation in regional unemployment within the United States. This increases confidence that our findings represent a general customer response to changes in economic conditions when volume discounts are available.

7. References

Ailawadi, Kusum L., Bari A. Harlam, Jacques César, and David Trounce (2007), "Quantifying and Improving Promotion Effectiveness at CVS," *Marketing Science*, 26(4), 566-575.

Anderson, Eric T., Nir Jaimovich, and Duncan Simester (2015), "Price Stickiness: Empirical Evidence of the Menu Cost Channel," *Review of Economics and Statistics*, 97(4), 813-826.

Anderson, Eric, Benjamin A. Malin, Emi Nakamura, Duncan Simester and Jón Steinsson (2016) "Informational Rigidities and the Stickiness of Temporary Sales," working paper, Columbia University.

Beraja, Martin, Erik Hurst, and Juan Ospina (2015), "The Aggregate Implications of Regional Business Cycles," working paper, University of Chicago.

Bils, M. J. and P. Klenow (2004), “Some Evidence on the Importance of Sticky Prices,” *Journal of Political Economy*, 112(5), 947-985.

Carlton, Dennis (1986), “The Rigidity of Prices,” *American Economic Review*, 76(4), 637-656.

Chevalier, J. A. and A. K. Kashyap (2014), “Best Prices: Price Discrimination and Consumer Substitution,” NBER Working Paper No. 20768.

Coibion, O., Y. Gorodnichenko and G.H. Hong (2015), “The Cyclicality of Sales, Regular, and Effective Prices: Business Cycle and Policy Implications,” *American Economic Review*, 105(3), 993-1029.

Eichenbaum, M, N. Jaimovich and S. Rebelo (2011), “Reference Prices and Nominal Rigidities,” *American Economic Review*, 101(1), 234-262.

Hansen, Karsten, Vishal Singh, and Pradeep Chintagunta (2006), “Understanding Store-Brand Purchase Behavior Across Categories,” *Marketing Science*, 25(1), 75-90.

Hoch, Stephen J., Byung-Do Kim, Alan L. Montgomery and Peter E. Rossi (1995), “Determinants of Store Level Price Elasticity,” *Journal of Marketing Research*, 32(1), 17-29.

Hoch, Stephen J. (1996), “How Should National Brands Think About Private Labels,” *Sloan Management Review*, 37(2), 89-102.

Ito, Koichiro (2013), “How do Consumers Respond to Nonlinear Pricing? Evidence from Household Water Demand,” working paper, Stanford University.

McShane, Blakely B., Chaoqun Chen, Eric T. Anderson, and Duncan I. Simester (2016), “Decision Stages and Asymmetries in Regular Retail Pass-Through,” *Marketing Science*, forthcoming.

Maćkowiak, Bartosz and Wiederholt, Mirko (2009), “Optimal Sticky Prices under Rational Inattention,” ECB Working Paper No. 1009.

Nakamura, E., and J. Steinsson (2008), “Five Facts about Prices: A Reevaluation of Menu Cost Models,” *Quarterly Journal of Economics*, 123(4), 1415-1464.

Stigler, George and James Kindahl (1970), *The Behavior of Industrial Prices*, NBER General Series, No. 90, New York: Columbia University Press,

Zbaracki, Mark J., Mark Ritson, Daniel Levy, Shantanu Dutta, Mark Bergen (2006), “Managerial and Customer Costs of Price Adjustment: Direct Evidence from Industrial Markets,” *Review of Economics and Statistics*, 86(2), 514-533.

Appendix

List of Findings in the Appendix

Section 2

Summary Statistics

Section 3 Manufacturer Sales to Distributors

Countries Represented in the Analysis

Alternative Clusters of Standard Errors

Probability of an Order: Frequently vs. Infrequently Purchased Items

Robustness Checks

Changes in the Manufacturer's Profit Margins

Section 4 Distributor Sales to End Customers

Countries Represented in the Analysis

Probability of an Order

Robustness Checks

Changes in the Distributors' Profit Margins

Section 5 Retail Transactions

Definition of Variables

Alternative Clusters of Standard Errors

Robustness Checks

Summary Statistics

		Mean	Std Deviation	Sample Size
Section 3 Manufacturer Sales				
Prob. Of An Order	Any Discounted Order	17.21%	37.74%	2,341,980
	Any Undiscounted Order	8.76%	28.27%	2,341,980
	Unemployment	7.22	1.86	2,341,980
	Price Index	-0.0029	0.0297	2,341,980
Average Order Size	Proportion Discounted	59.84%	21.85%	1,060
	Average Discount	19.70%	8.70%	1,060
	Order Size	-0.0150	0.1919	1,060
	Unemployment	7.13	3.29	1,060
	Price Index	-0.0018	0.0376	1,060
Profit Margin	Unemployment	7.09	2.10	974,117
Section 4 Distributor Sales				
Prob. Of An Order	Any Discounted Order	3.91%	19.38%	41,400,840
	Any Undiscounted Order	7.31%	26.04%	41,400,840
	Unemployment	8.16	2.92	41,400,840
	Price Index	-0.0011	0.02138	41,400,840
Average Order Size	Proportion Discounted	38.93%	12.53%	1,105
	Average Discount	10.78%	7.68%	1,105
	Order Size	-0.0056	0.1800	1,105
	Unemployment	8.09	4.16	1,105
	Price Index	-0.0018	0.0634	1,105
Profit Margin	Unemployment	8.23	3.12	5,796,927
Section 5 Retail Data	Proportion Discounted	41.03%	45.19%	2,110,952
	Average Discount	12.60%	17.04%	2,110,952
	Unemployment	5.81	2.19	2,110,952
	Price Index	-0.05	0.28	2,110,952

The table reports summary statistics for the variables used in the analysis. The *Profit Margin* statistics are omitted as this information is confidential.

Countries Represented in the Analysis: Manufacturer Sales

Asia Pacific	Australia New Zealand South Korea
EMEA (Europe, Middle East, Africa)	Austria Belgium Denmark Finland France Ireland Israel Italy Netherlands Norway Poland Slovenia South Africa Spain Sweden Switzerland Turkey United Kingdom
Japan	Japan
Americas	Brazil Canada Mexico United States

Alternative Clusters of Standard Errors: Manufacturer Sales

		Clustering by Country	Clustering by Quarter	Clustering by Country x Qtr
Prob. Of an Order	Any Discounted Order	-0.132%	-0.132%	-0.132%
		(0.617%)	(0.574%)	(0.653%)
	Any Undiscounted Order	-0.886%**	-0.886%**	-0.886%**
		(0.278%)	(0.169%)	(0.174%)
Order Size	Proportion Discounted	2.061%*	2.061%**	2.061%**
		(0.753%)	(0.250%)	(0.437%)
	Average Discount	0.660%**	0.660%**	0.660%**
		(0.224%)	(0.109%)	(0.134%)
	Order Size	2.611%*	2.611%**	2.611%**
		(1.091%)	(0.600%)	(0.611%)

The table reports the *Unemployment* (β_1) coefficients from estimating Equations 3.1 and 3.2 using alternative standard errors. The sample size in the *Probability of an Order* models is 2,341,980 and the sample size in the *Order Size* models is 1,060.

Frequently vs. Infrequently Purchased Items: Manufacturer Sales

	Items with 1 to 4 orders	Items with 5 to 8 orders	Items with 9 to 12 orders	Items with more than 12 orders
Any Discounted Order				
Unemployment	-1.120%* (0.595%)	1.231%† (0.673%)	2.236%** (0.818%)	3.130%** 0.842%)
Price Index	-4.521% (3.671%)	-24.212%* (9.951%)	-29.806%** (8.186%)	-45.126%** (8.990%)
R^2	0.1104	0.2190	0.3243	0.4663
Any Undiscounted Order				
Unemployment	-1.130%** (0.154%)	-1.376%** (0.358%)	-1.223%** (0.429%)	-2.688%** (0.832%)
Price Index	-2.401% (2.651%)	-10.742% (7.959%)	-20.205%* (9.286%)	-20.702% (15.252%)
R^2	0.1291	0.3021	0.4234	0.5333
Sample Size	2,078,880	335,080	123,810	130,480

The table report the coefficients from estimating Equation 3.1 on different groups of items. Standard errors clustered at the country x quarter level are reported in parentheses.

Robustness Checks: Manufacturer Sales

	Proportion Discounted	Average Discount	Order Size
Static Weights (*Average Quarterly Revenue*)	2.177%** (0.408%)	0.708%** (0.128%)	2.789%** (0.579%)
No Volume Discount Restriction	2.345%** (0.577%)	0.761%** (0.177%)	
Order Size Calculated Using Average Order Size			3.853%** (0.840%)

The table reports the (β_1) coefficients from estimating Equation 3.2. Fixed effects were estimated but are omitted from the table. Standard errors clustered at the country x quarter level are in parentheses. The sample size in all of the models is 1,060. The observations are weighted by revenue in all of the models (except the Static Weights models).

Changes in the Manufacturer's Profit Margins

In Model 1 in the table below we report the coefficient of interest when estimating Equation 3.3. We omit the fixed effects and report standard errors clustered at the country x quarter levels.

	Model 1	Model 2	Model 3
Unemployment	-1.290%**	-0.969%**	-0.578%**
	(0.231%)	(0.192%)	(0.152%)
R^2	0.1258	0.1289	0.1344

The table reports the *Unemployment* coefficient (β_1) when estimating Equation 3.3 on each dependent variable. Standard errors clustered at the country x quarter level are reported in parentheses. The sample size in all three models is 974,117. The observations are weighted by each transaction's revenue.

In the other two models in this table we disentangle how these three effects impact the cyclicality in profit margins. In particular, we re-calculate the dependent variable (*Profit Margin*) to remove different sources of variation in prices and costs:

Model 2: We calculate the *Profit Margin* using the prices and costs averaged at the distributor x item x order quantity (aggregating across time and countries).[20]

Model 3: We calculate the *Profit Margin* using the prices and costs averaged at the distributor x item level (aggregating across time, countries and purchase quantities).

Using the average price and cost calculated at the distributor x item level (Model 3) removes any price variation due to both: (a) manufacturer changes, and (b) substitution between order quantities. With this measure we isolate just the third effect identified in the main text: substitution between items.

Using the average price and cost calculated at the distributor x item x *order quantity* level (Model 2) preserves variation due to customers ordering different quantities. Intuitively, this treats a product as an item x order quantity, and so preserves variation due to substitution either between items or between order quantities for the same item. However, variation due to manufacturer cost and price changes is averaged out.

[20] Averaging prices and the manufacturer's costs across countries (at a distributor x item level) requires that we use a common currency. We convert all prices and costs to USD and use static exchange rates to avoid introducing variation through the exchange rates. Further investigation reveals that allowing the exchange rates to vary by quarter has little impact on the findings. This is because the exchange rates only affect the averaging of prices where a distributor purchases the same item in multiple currencies.

We can use these results to decompose the sources of cyclicality in profit margins. We summarize this decomposition in the table below. In the base model (Model 1), the *Unemployment* coefficient is identified by all three effects. In Model 2 the *Unemployment* coefficient is no longer identified by changes in the manufacturer's costs and prices. In Model 3 the *Unemployment* coefficient is no longer identified by either manufacturer cost and price changes or variation in order quantities. Therefore, the differences in the *Unemployment* coefficients between the models provide an estimate of the incremental effect contributed by each component.

Contribution to Cyclicality in Profit Margins	**Comparison of *Unemployment* Coefficients**	**Contribution**
Total Change From All 3 Effects	Model 1	-1.288%
Seller Cost and Price Changes	Model 1 – Model 2	-0.319%
Variation in Order Quantities	Model 2 – Model 3	-0.391%
Substitution Between Items	Model 3	-0.578%

We offer two additional comments on this decomposition. First, a strength of the decomposition is that these three effects fully explain the overall effect (the decomposition is exhaustive). This can be seen by asking what happens if we were to calculate the *Profit Margin* using prices and costs averaged at the distributor x country level (averaging across items). Equation 3.3 would be over-identified as all of the remaining variation in the dependent variable would be fully explained by the fixed effects.

Second, we acknowledge that this estimate of the contribution of the two buyer behavior effects (substitution between items and order quantities) is arguably inflated in at least one respect. To the extent that there are interactions between the buyer and seller decisions, then these are attributed to the buyer not the seller (the decomposition attributes residual variation to the buyer decisions). This increases the estimated role of buyer behavior. As justification for this attribution we note that the price variation would not have occurred without the buyer decision, but could have occurred without the manufacturer price change.

Countries Represented in the Analysis: Distributor Sales

Asia Pacific	Australia New Zealand South Korea
EMEA (Europe, Middle East, Africa)	Austria Belgium Czech Republic Denmark Estonia Finland France Ireland Israel Italy Netherlands Norway Poland Portugal Slovakia Slovenia South Africa Spain Sweden Switzerland Turkey United Kingdom
Japan	Japan
Americas	Brazil Canada Chile Mexico United States

Probability of an Order: Distributor Sales

	Items with 1 to 4 orders	Items with 5 to 8 orders	Items with 9 to 12 orders	Items with more than 12 orders	All Items
Any Discounted Order	0.213%**	1.241%**	1.803%**	2.964%**	0.243%**
	(0.032%)	(0.143%)	(0.196%)	0.306%)	(0.037%)
Any Undiscounted Order	-0.009%	-0.055%	-0.158%	-1.124%**	-0.005%
	(0.054%)	(0.108%)	(0.163%)	(0.240%)	(0.056%)
Sample Size	39,069,446	1,488,487	416,546	426,361	41,400,840

The table report the *Unemployment* (β_1) coefficients from estimating Equation 4.1 on different groups of items. Standard errors clustered at the country x quarter level are reported in parentheses.

Robustness Checks: Distributor Sales

	Proportion Discounted	Average Discount	Order Size
Static Weights	1.434%**	0.195%**	6.893%**
(*Average Quarterly Revenue*)	(0.188%)	(0.052%)	(0.830%)
No Volume Discount Restriction	1.421%**	0.191%**	
	(0.188%)	(0.0517%)	
Order Size Calculated Using Average Order Size			6.287%**
			(0.770%)

The table reports the *Unemployment* (β_1) coefficients from estimating Equation 4.2. Fixed effects were estimated but are omitted from the table. Standard errors clustered at the country x quarter level are in parentheses. The sample size in all of the models is 1,105. The observations are weighted by revenue in all of the models (except the Static Weights models).

Alternative Clusters of Standard Errors: Distributor Sales

		Clustering by Country	Clustering by Quarter	Clustering by Country x Qtr
Order Size	Proportion Discounted	1.420%**	1.420%**	1.420%**
		(0.369%)	(0.214%)	(0.188%)
	Average Discount	0.192%*	0.192%*	0.192%**
		(0.076%)	(0.074%)	(0.052%)
	Order Size	6.775%**	6.775%**	6.775%**
		(1.830%)	(0.835%)	(0.843%)

The table reports the *Unemployment* (β_1) coefficients from estimating Equation 4.2 using alternative standard errors.

Changes in the Distributors' Profit Margins

In the table below we use the distributors' sales data to re-estimate the same three models that we used to decompose the manufacturer's profit margin (see earlier discussion). In particular, we estimate a model analogous to Equation 3.3. The unit of analysis is a transaction, and the fixed effects are at the country x distributor ($\boldsymbol{\Theta}_{cd}$) and quarter levels, which capture the cross-sectional and temporal effects. Unlike the analysis in Section 3, where the *Profit Margin* is calculated using costs identified at the transaction level, in this analysis item costs are averaged at the distributor level (averaging across time, countries and distributor purchase quantities). We also report findings for two variants of this model:

Model 2: We calculate the *Profit Margin* using prices averaged at the end customer x distributor x item x order quantity (aggregating across time and countries).

Model 3: We calculate the *Profit Margin* using the prices averaged at the end customer x distributor x item level (aggregating across time, countries and purchase quantities).

The findings for all three models are reported in the table below, where we omit the fixed effects and report standard errors clustered at the country x quarter levels. The unit of analysis is a transaction and the observations are weighted by each transaction's revenue. Because we use static cost measures when calculating the distributors' profit margins, the results of this analysis are not directly comparable with the findings reported for the manufacturer in Section 3.

	Model 1	**Model 2**	**Model 3**
Distributor Profit Margins	-1.707%** (0.151%)	-1.620%** (0.144%)	-1.475%** (0.136%)
R^2	0.4444	0.4479	0.4565

The table reports the *Unemployment* coefficient (β_1) when estimating Equation 3.3 on each dependent variable. Standard errors clustered at the country x quarter level are reported in parentheses. The sample sizes are 5,796,927. The observations are weighted by each transaction's revenue.

Contribution to Cyclicality in Profit Margins	**Comparison of *Unemployment* Coefficients**	**Distributor Profit Margins**
Total Change From All 3 Effects	Model 1	-1.707%
Seller Price Changes	Model 1 – Model 2	-0.087%
Variation in Order Quantities	Model 2 – Model 3	-0.145%
Substitution Between Items	Model 3	-1.475%

Definition of Variables: Retail Data

Average Discount

We calculate the *Average Discount* as the percentage difference in how much customers would have spent if they had purchased the same number of total units at the small size unit price. The *Average Discount* is calculated as:

$$\frac{Revenue\ at\ Small\ Size\ Price_{is} - Actual\ Revenue_{ist}}{Revenue\ at\ Small\ Size\ Price_{is}}$$

where:

$$Revenue\ at\ Small\ Size\ Price_{is} = Unit\ Price\ of\ Small\ Size_{is} * Total\ Units\ Purchased_{ist}.$$

Profit Margin

The *Profit Margin*$_{ist}$ is defined as the percentage profit margin calculated using static prices and costs. It is calculated as:

(*Proportion Large*$_{ist}$ * *Large Profit Margin*$_{is}$) + (*Proportion Small*$_{ist}$ * *Small Profit Margin*$_{is}$)

The proportions are calculated at the unit level. For example, in a store x week that has sales of 10 packages of the (small) 10 unit size and 15 packages of the (large) 20 unit size, the *Proportion Large* would be 75% (300 units out of 400 total) and the proportion small is 25%. The profit margins for each size are calculated using the average prices and costs:

$$Small\ Profit\ Margin_{\ is} = \frac{\overline{Small\ Price}_{is} - \overline{Small\ Cost}_{is}}{\overline{Small\ Price}_{is}}$$

$$Large\ Profit\ Margin_{\ is} = \frac{\overline{Large\ Price}_{is} - \overline{Large\ Cost}_{is}}{\overline{Large\ Price}_{is}}$$

In these expressions $\overline{Large\ Price}_{is}$ and $\overline{Large\ Cost}_{is}$ refer to the average unit price and unit cost for the large package size calculated across the entire data period (and similarly for the small package sizes).

Alternative Clusters of Standard Errors: Retail Data

In the table below we report the key results using three alternative clusters: the country, the quarter, and the country x quarter. We focus on the coefficient of interest.

	Clustering by Store	Clustering by Month	Clustering by Store x Month	Clustering by CBSA x Month
Proportion Discounted	0.199% (0.122%)	0.199%* (0.096%)	0.199%** (0.060%)	0.199%* (0.081%)
Average Discount	0.056%† (0.029%)	0.056%* (0.025%)	0.056%** (0.016%)	0.056%* (0.020%)
Profit Margin	-0.024%** (0.008%)	-0.024%** (0.005%)	-0.024%** (0.005%)	-0.024%** (0.005%)

The table reports the *Unemployment* (β_1) coefficients from estimating Equation 5.1. The sample size is 2,110,952.

Robustness Checks: Retail Data

	Proportion Discounted	Average Discount	Profit Margin
Static Weights	0.300%** (0.087%)	0.096%** (0.023%)	-0.025%** (0.007%)
All Observations	0.176%* (0.078%)	0.044%* (0.019%)	-0.021%** (0.005%)

The table reports the (β_1) coefficients from estimating Equation 5.1(*Proportion Discounted* and *Average Discount*) and Equation 5.2 (*Profit Margin*). Fixed effects were estimated but are omitted from the table. Standard errors clustered at the CBSA x month level are in parentheses. The sample size in the static weights models are 2,110,952. The sample size in the models with all observations (including store x item combinations without volume discounts) are 2,284,475.