NBER WORKING PAPER SERIES

TRADING DOWN AND THE BUSINESS CYCLE

Nir Jaimovich
Sergio Rebelo
Arlene Wong

Working Paper 21539
http://www.nber.org/papers/w21539

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
September 2015

This research was conducted with restricted access to the Bureau of Labor Statistics (BLS) data. The views expressed here are those of the authors and do not necessarily reflect the views of the BLS, the Federal Reserve Bank of Minneapolis, the Federal Reserve System, or the National Bureau of Economic Research. We thank our project coordinator at the BLS, Ryan Ogden, for help with the data and Yuriy Gorodnichenko and Michael Weber for making their data and code available to us. We also thank Julie Hoffmeyer and Joe Derochowski from the NPD Group for sharing their data with us. We are grateful for comments from David Berger, Guido Menzio, Emi Nakamura, Aviv Nevo, Tommaso Oliviero, Raphael Schoenle, Vincent Sterk, Henry Siu, and Jón Steinsson.

Trading Down and the Business Cycle
Nir Jaimovich, Sergio Rebelo, and Arlene Wong
NBER Working Paper No. 21539
September 2015, Revised January 2017
JEL No. E1,E2,E3

ABSTRACT

We document two facts. First, during the Great Recession, consumers traded down in the quality of the goods and services they consume. Second, the production of low-quality goods is less labor intensive than that of high-quality goods. Hence, when households traded down, labor demand fell, increasing the severity of the recession. We find that the trading-down phenomenon accounts for a substantial fraction of the fall in U.S. employment in the recent recession. We then study a general equilibrium model that embeds quality choice and is consistent with our empirical findings.

Nir Jaimovich
Department of Finance and Business Economics
Marshall Business School
University of Southern California
Los Angeles, CA 90089-0804
and NBER
nir.jaimovich@marshall.usc.edu

Sergio Rebelo
Northwestern University
Kellogg School of Management
Department of Finance
Leverone Hall
Evanston, IL 60208-2001
and NBER
s-rebelo@northwestern.edu

Arlene Wong
Federal Reserve Bank of Minneapolis
Research Division, 90 Hennepin Ave
Minneapolis, MN 55401
arlenewong@princeton.edu

1 Introduction

One of the classic research areas in macroeconomics is the study of how households make consumption choices and how these choices impact the economy. There is a large empirical literature on this topic going back at least to the work of Burns and Mitchell (1946). This literature has gotten renewed attention after the Great Recession.[1]

In this paper, we contribute to this line of research as follows. First, we show that during the Great Recession, consumers traded down in the quality of the goods and services they consume. Second, we show that the production of low-quality goods is less labor intensive than that of high-quality goods. Hence, when households traded down, labor demand fell, increasing the severity of the Great Recession.

To quantify the implications of "trading down" for employment during the Great Recession, we combine various data sources to construct a data set with firm-level measures of product quality, labor intensity, and market share. For most of our analysis, we use prices as a proxy for quality. Our assumption is that, if consumers are willing to pay more for an item, they perceive it to be of higher quality. We obtain price measures from two sources: data scraped from the Yelp! website and the confidential micro data set used to construct the Producer Price Index (PPI). We merge these data with Compustat data to measure labor intensity and market share for each firm in our sample. We estimate that around half of the decline in employment in the Great Recession is accounted for by consumers trading down in the quality of the goods and services they purchased.

To study the effects of trading down from a theoretical perspective, we embed quality choice into a general equilibrium model with endogenous determination of employment where heterogeneous households who choose both the quantity and quality of the goods they consume. In the model, households face a natural trade-off: consuming higher

[1]Recent contributions to this literature include Aguiar, Hurst, and Karabarbounis (2013), Kaplan and Menzio (2015), Nevo and Wong (2015), and Coibion, Gorodnichenko, and Hee (2015).

quality goods yields higher utility but higher quality goods are more expensive.

We first discuss some analytical results on the general equilibrium effect of quality choice on employment. Then, we calibrate our model to study its quantitative implications for the importance of trading down during the Great Recession.

Our paper is organized as follows. In Section 2, we describe our data and present our empirical results. We present the model in Section 3 and study its quantitative properties in Section 4. Section 5 concludes.

2 Empirical findings

In this section, we analyze the impact of trading down on labor demand. By trading down we mean shifts in the composition of consumption across firms *within* narrowly defined sectors towards lower quality goods. We relate the composition shifts to the labor intensity distribution within sectors to examine the impact of trading down on employment. Our within-sector trading-down channel differs from other channels of demand that can affect employment, such as shifts in spending *across* sectors (e.g. from restaurants to the grocery sector)[2], or declines in total spending.

Our empirical approach to study the effects of trading down is as follows. We denote by M the number of sectors in the economy. Total aggregate employment across these sectors is:

$$N_t = \sum_{m=1}^{M} N_{m,t}, \tag{1}$$

where $N_{m,t}$ denotes employment at time t in sector m.

In each sector, goods can belong to one of J levels of quality. The market share of quality tier j in sector m ($S_{j,m,t}$) is the ratio of sales of goods in quality tier j ($Y_{j,m,t}$) to total sales in sector m ($Y_{m,t}$):

$$S_{j,m,t} = Y_{j,m,t}/Y_{m,t}.$$

[2]See discussion in Section 2.5

The measure of labor intensity ($LI_{j,m,t}$) that we construct with our data is the ratio of employees ($N_{j,m,t}$) to sales ($Y_{j,m,t}$):

$$LI_{j,m,t} = N_{j,m,t}/Y_{j,m,t}.$$

Using this notation, we can write total employment in sector m in period, t as:

$$N_{m,t} = Y_{m,t} \sum_{j=1}^{J} S_{j,m,t} LI_{j,m,t} \,. \tag{2}$$

Combining equations (1) and (2), we can write aggregate employment as:

$$N_t = Y_t \sum_{m=1}^{M} \left(\frac{Y_{m,t}}{Y_t} \sum_{j=1}^{J} S_{j,m,t} LI_{j,m,t} \right). \tag{3}$$

Using equation (3), we can write the log-percentage change in employment, $\log(N_{t+1}/N_t)$ as:

$$\begin{aligned} \log(N_{t+1}/N_t) \quad &= \quad \log\left(Y_{t+1}/Y_t\right) + \log\left(\sum_{m=1}^{M} \sum_{j=1}^{J} \frac{Y_{m,t+1}}{Y_{t+1}} S_{j,m,t+1} LI_{j,m,t+1} \right) \\ &\quad - \log\left(\sum_{m=1}^{M} \sum_{j=1}^{J} \frac{Y_{m,t}}{Y_t} S_{j,m,t} LI_{j,m,t} \right). \end{aligned} \tag{4}$$

In order to quantify the role played by quality trading down in employment, we calculate a counterfactual value for employment in period $t+1$:

$$N_{t+1}^{CF} = Y_{t+1} \sum_{m=1}^{M} \left(\frac{Y_{m,t+1}}{Y_{t+1}} \sum_{j=1}^{J} S_{j,m,t} LI_{j,m,t+1} \right). \tag{5}$$

This counterfactual value corresponds to the level employment without quality trade down–for each sector; i.e. the market shares of each quality tier in period $t+1$ are the same as in period t. From this we can calculate what, in the absence of trading down, the change in employment would have been:

$$
\begin{aligned}
\log(N_{t+1}^{CF}/N_t) &= \log\left(Y_{t+1}/Y_t\right) + \log\left(\sum_{m=1}^{M}\sum_{j=1}^{J}\frac{Y_{m,t+1}}{Y_{t+1}} S_{j,m,t} LI_{j,m,t+1}\right) \\
&\quad - \log\left(\sum_{m=1}^{M}\sum_{j=1}^{J}\frac{Y_{m,t}}{Y_t} S_{j,m,t} LI_{j,m,t}\right).
\end{aligned}
\tag{6}
$$

The difference between the actual change in employment and the counterfactual change in employment if no trading down occurred, gives us an estimate of the importance of the trading down phenomenon for employment. We proceed with computing this estimate using our data.

2.1 Empirical measures

We start by using a data set that merges data from Yelp!, the Census of Retail Trade, and Compustat. We then extend our analysis to the manufacturing sector by using the micro data gathered by the BLS to construct the PPI. Finally, we consider several other data sets.

2.2 Results obtained with Yelp! and Census of Retail Trade data

In this section, we discuss the results we obtain using data from Yelp! and from the Census of Retail Trade. The combined data set covers five North American Industry Classification System (NAICS) sectors: apparel, grocery stores, restaurants, home furnishing, and general merchandise. These sectors represent 17 percent of private non-farm employment.

We focus on two main time periods. Our recession period covers 2007-2012 period, while our pre-recession period covers 2004-2007.[3] Since income inequality has been

[3]Even though the NBER determined that the recession ended in June 2009, average and median household income continued to fall until 2012. In addition, employment recovered very slowly: in December 2012 employment was still 3 percent below its December 2007 level. In Appendix E, we report results when we use 2007-2009 as the recession period.

rising, trading down was occurring before the Great Recession. The data for the period 2004-2007 allows us to control for trends in trading down so that we can isolate the cyclical component of trading down that is associated with the Great Recession.

Yelp!

For sectors other than General Merchandise, we collect information on prices by scraping data from Yelp!, a website where consumers share reviews about different goods and services.[4] For each store and location pair, Yelp! asks its users to classify the price of the goods and services they purchased into one of four categories: $ (low), $$ (middle), $$$ (high), and $$$$ (very high). Since there are few observation in the very-high category, we merge the last two categories into a single high-price category.

We construct the Yelp! data set as follows (see Appendix A for more details). We first associate each firm (for example, Cost Plus, Inc.) with its brand names and retail chains (for example, Cost Plus owns the retail chain World Market). We find the Yelp! profile for each retail chain and brand in the 18 largest U.S. cities and collect the first match (for example, the first match for World Market in Chicago is the store on 1623 N. Sheffield Av.). We then compute the average price category across the first match for each of the 18 cities (to compute this average, we assign 1 to category low, 2 to middle and 3 to high).[5]

U.S. Census of Retail Trade

For General merchandise, the U.S. Census of Retail Trade splits firms into three price tiers that correspond to three different levels of quality: non-discount stores (high quality), discount department stores (middle quality), other general merchandise stores,

[4]Yelp! users also rate the quality of the goods and services they consume. These ratings are not useful for our purpose because they are not an absolute measure of quality. Instead, they measure the quality of an item relative to the price paid for that item. For example, a fast-food restaurant that receives five stars might be worse than a high-priced restaurant that receives three stars.

[5]The dispersion in price categories across cities is relatively small; it is rare for firms to be included in different price categories in different cities.

including family dollar stores (low quality). For each of these three tiers, the Census provides information about employment and sales. We use this information to construct labor intensity measures and market shares.

Compustat

We merge the price information for each firm in our Yelp! data set with data from Compustat on the number of employees, sales, operating expenses, and cost of goods sold. Our measure of labor intensity is the ratio of employees to sales.[6]

Findings

Tables 1 and 2 present our estimates of labor intensity and market share by quality tier. We use these estimates as inputs into equations (4) and (6). Labor intensity increases with quality in all five sectors. For example, in 2007, the number of employees per million dollar of sales is 6.33, 9.23 and 10.86 for low-, middle-, and high-quality apparel stores, respectively (Table 1). Overall, the middle quality uses 46 percent more workers per million dollar of sales than the low-quality producers. Similarly, the high-quality producers use 18 percent more workers per million dollar of sales than the middle-quality producers. Other things equal, a shift of one million dollar of sales from middle to a low quality stores reduces employment by roughly three jobs.

Between 2007 and 2012, firms that produce middle- and high-quality items lost market share relative to firms that produce low-quality items. In 2007, the low-, middle- and high-price categories account for 42, 52 and 6 percent of sales, respectively. In contrast by 2012, high-quality producers lost about 0.5 percentage points in market

[6]For the sample of firms that report the labor share in cost (less than 1/4 of the firms included in Compustat data), the correlation between labor share and the labor intensity measure of employees/sales is 0.94. The correlation between the labor share in cost and employees/gross margin is 0.97. As a robustness check, we also use the ratio of employees to gross margin as a measure of labor intensity (Appendix D). Gross margin, which is sales minus cost of goods sold, is a measure that is close to value added. Value added is equal to the gross margin minus energy and services purchased. We cannot compute value added because Compustat does not report data on energy and services purchased. The correlation between employees/gross margin and employees/sales is 0.72.

share, middle-quality producers lost 6.5 percentage points, and low-quality producers gained 7 percentage points. This pattern is present is all the sectors we consider with one exception: the market share of high-quality grocery stores increased.[7]

With the information in Tables 1 and 2, we can then implement the empirical approach in equations 4 and 6. Overall employment in the sectors included in our data set fell by 3.39 percent. Using equation (6) we find that in the absence of trading down employment would have fallen by only 0.39 percent. This result implies that trading down accounts for $(3.39 - 0.39)/3.39 = 88$ percent of the fall in employment.[8]

Table 1: Market shares and Labor Intensity: 2007

Industry	$m Sales	Labor intensity			Market share		
		Low	Middle	High	Low	Middle	High
Home furnishing and appliances	654,535	4.40	4.97	7.10	1%	95%	5%
Grocery stores	547,837	3.37	4.68	7.58	39%	59%	2%
Food services and drinking places	444,551	15.63	24.02	22.43	52%	41%	7%
Clothing stores	221,205	7.55	9.43	16.49	11%	78%	11%
General merchandise stores	578,582	3.72	6.92	7.19	64%	23%	13%
Total	2,446,710	6.33	9.23	10.86	35%	58%	7%

Note: This table depicts the 2007 total sales, labor intensity and market share for different retail sectors. The last row is a sales-weighted measure. Sales are from the US Census of Retail Trade, and the labor intensities and market shares are from Compustat. Labor intensity is defined as the number of employees per million dollars of sales and the market share are the share of sales for each price tier within each sector. Price tiers are denoted by low, middle and high, and are based on Yelp! classifications of prices $, $$, and $$$, respectively. See text for more information.

[7]This exception is driven by an outlier: WholeFoods, a high-quality supermarket that gained market share despite the recession.

[8]Our calculations are based on Census estimates of sector expenditures. We are implicitly assuming that the market shares and labor intensities in our data are representative of each sector in a whole.

Table 2: Market shares and Labor Intensity: 2012

Industry	$m Sales	Labor intensity			Market share		
	Census	Low	Middle	High	Low	Middle	High
Home furnishing and appliances	609,323	3.49	4.92	5.93	1%	94%	5%
Grocery stores	631,486	1.92	4.15	6.06	43%	53%	5%
Food services and drinking places	524,892	13.43	19.49	22.40	61%	33%	6%
Clothing stores	241,386	6.50	9.16	15.09	15%	77%	7%
General merchandise stores	649,754	3.72	6.92	7.19	72%	18%	10%
Total	2,656,841	5.41	8.49	10.36	42%	52%	6%

Note: This table depicts the 2012 total sales, labor intensity and market share for different retail sectors. The last row is a sales-weighted measure. Sales are from the US Census of Retail Trade, and the labor intensities and market shares are from Compustat. Labor intensity is defined as the number of employees per million dollars of sales and the market share are the share of sales for each price tier within each sector. Price tiers are denoted by low, middle and high, and are based on Yelp! classifications of prices $, $$, and $$$, respectively. See text for more information.

Trend versus Cycle

Trading down was occurring before the recession, so there are both trend and cycle components of trading down. In order to disentangle these two components, we proceed as follows. We compute market shares by quality tier for each sector for the period 2004-2007. We use the change in these shares over this period to linearly extrapolate what the market shares of different quality tiers would have been in 2012. Using these market shares, we construct the 2012 employment implied by the extrapolated market shares

$$N_{2012}^{CF2} = Y_{2012} \sum_{m=1}^{5} \left(\frac{Y_{m,2012}}{Y_{2012}} \sum_{j=1}^{3} \left[S_{j,m,2007} + \frac{S_{j,m,2007} - S_{j,m,2004}}{(2007-2004)} \times (2012-2007) \right] [LI_{j,m,2012}] \right). \tag{7}$$

This counterfactual measure of employment is 2.38 percent lower than the level of employment in 2007. We conclude that $(3.39 - 2.38)/3.39 = 30$ percent of the fall in employment is due to trend factors. Recall that trading down accounted for a 3 percentage points fall in employment, so the part of trading down associated with cyclical factors is: $[3 - (3.39 - 2.06)] = 1.67$. In other words, $1.67/3.39 = 50$ percent of the fall in employment is due to trading down associated with cyclical factors. The remaining 11 percent of the fall in employment are unrelated to the trading down phenomenon.

2.3 The PPI data

In order to extend the analysis to the manufacturing sector, we use the confidential micro data collected by the Bureau of Labor Statistics (BLS) to construct the PPI.[9] As with the Yelp! data, we merge the PPI data with Compustat to obtain price, labor intensity, and market share for each firm. This combined data set has $62,000$ monthly observations for the period 2007-2012. Overall, the sectors covered by the merged PPI and Compustat data account for 15 percent of private non-farm employment. We refer the reader to Appendix B for more details on the construction of this data set.

We focus on the 2-digit NAICS manufacturing sectors 31, 32, and 33 because in these sectors we are able to merge the PPI and Compustat data for more than 10 firms per sector and span a range of quality tiers.

In order to construct an indicator of quality for each firm, we need the level of the unit price per item rather than the inflation rate per item. The PPI provides information on the unit of measure for each item which we use to ensure that prices in our sample refer to the same unit of measurement (e.g. pounds). Unfortunately, there is a large number of observations on the unit of measure missing before 2007. This limitation restricts our ability to account for "pre-Great Recession" trend.

In order to construct a quality measure per each firm we proceed as follows. For each product k that establishment e sells in year t, we calculate its price, p_{ket}, relative to the median price in the industry for product k in year t, $\bar{p}_{kt}$:[10]

[9] Examples of other papers that use these data include Nakamura and Steinsson (2008), Gilchrist et al (2014), Gorodnichenko and Weber (2014), and Weber (2015).

[10] Our analysis is based on products defined at a six-digit industry code and then further disaggregated by product type. Therefore, although the results are presented at a 2-digit level, all relative

$$R_{ket} = p_{ket}/\bar{p}_{kt}.$$

For single-product establishments, we use this relative price as the measure of the quality of the product produced by establishment e. For multi-product establishments, we compute the establishment's relative price as a weighted average of the relative price of different products, weighted by shipment revenue in the base year (w_{ke}):[11]

$$R_{et} = \sum_{k \in \Omega} w_{ke} R_{ket}.$$

where Ω denotes the set of all products in the PPI data set that we examined.

To make our results comparable with those obtained with Yelp! data, we proceed as follows. Once we rank establishments by their relative price, we assign the top 7 percent to the high-quality category, the middle 58 percent to the middle-quality category, and the bottom 35 percent to the low-quality category. Recall that this is the distribution of firms by quality tier that characterize the firms included in the Yelp! data set.

We aggregate the establishment quality tier assignment to firm level by taking a shipment-value weighted average of the quality tier and rounding to the closest quality tier. Finally, we merge the firm-level quality tier assignment from the PPI with the Compustat sample of firms.[12] This merged data set allows us to compute labor intensity

prices are defined at a narrow 6-digit level for comparability. For details of this disaggregation, see Table 11 of the BLS PPI Detailed Report. The variable $\bar{p}_{kt}$ is a shipment-value weighted average within the product category. For reporting purposes, we aggregate the results to the two-digit level. The aggregation is based on shipment revenue.

[11]This approach for constructing firm-level price indices is similar to that used by Gorodnichenko and Weber (2014), and Gilchrist et al (2014). However, we compute relative prices using a much finer product definition than these authors. We refer the reader to Section II in Gorodnichenko and Weber (2014) for a discussion of how the BLS samples products and firms.

[12]The aggregation of establishments up to firm level uses the matching done by Gorodnichenko and Weber (2014), who shared their code with us. In their work, they manually matched the names of establishments to the name of the firm. They also searched for names of subsidiaries and checked for any name changes of firms within the Compustat data set. See Gorodnichenko and Weber (2014) for more detail. A similar exercise of matching establishments to firms is used in Gilchrist et al (2014).

by quality tier.[13]

Tables 3 and 4 shows that our two key facts hold in the PPI data. First, low-quality firms gained market share between 2007 and 2012 at the cost of middle and high-quality firms. Second, quality is correlated with labor intensity. High-quality producers have higher labor intensity than middle-quality producers and middle-quality producers have higher labor intensity than low-quality producers.[14]

We now use the PPI data to implement our empirical approach. Overall employment in the sectors included in our data fell by approximately 8.6 percent. The counterfactual fall in employment that would have occurred without trading down is 3.9 percent. Hence, trading down accounts for 54 percent of the fall in employment.

Table 3: PPI Sectors 2007

Industry	$m Expenditure in 2007	Labor Intensity			Market share		
		Low	Middle	High	Low	Middle	High
31	811,751	0.74	3.41	n.a.	23%	77%	n.a.
32	1,434,885	2.73	2.99	4.62	26%	45%	29%
33	2,457,336	2.04	2.60	4.05	31%	63%	6%
Total	4,703,972	2.03	2.86	3.53	28%	58%	14%

Note: This table depicts the 2007 total sales, labor intensity and market share for different manufacturing sectors. The last row is a sales-weighted measure. Sales are from Census, and the labor intensities and market shares are from Compustat. Labor intensity is defined as the number of employees per million dollars of sales and the market share are the share of sales for each price tier within each sector. Price tiers are denoted by low, middle and high, and are based on firm-level producer price data. See text for more information.

[13] We use the entire sample of establishments within the PPI to rank the establishments, not just those that we are able to match with Compustat.

[14] We do not have any firms within the high-quality tier (defined based the PPI data) that could be merged with the Compustat data.

Table 4: PPI Sectors 2012

Industry	$m Expenditure in 2012	Labor Intensity			Market share		
		Low	Middle	High	Low	Middle	High
31	956,083	0.40	3.41	n.a.	34%	66%	n.a.
32	1,461,253	2.69	2.85	4.59	27%	47%	26%
33	2,494,959	1.40	2.41	3.32	38%	57%	5%
Total	4,912,295	1.59	2.74	3.06	33%	54%	13%

Note: This table depicts the 2012 total sales, labor intensity and market share for different manufacturing sectors. The last row is a sales-weighted measure. Sales are from Census, and the labor intensities and market shares are from Compustat. Labor intensity is defined as the number of employees per million dollars of sales and the market share are the share of sales for each price tier within each sector. Price tiers are denoted by low, middle and high, and are based on firm-level producer price data. See text for more information.

In sum, our results using the PPI data are consistent with those obtained with Yelp! and Census of Retail Trade data. Higher priced stores, which are generally more labor intensive, lost market share during the recent recession. This loss of market share accounts for about half of the overall decline in employment.

2.4 NPD data

In this subsection, we discuss results obtained using data on the evolution of market shares in restaurants of different quality levels. This data set collected by the NPD Group (a marketing consulting firm) includes restaurant traffic (number of meals served) and consumer spending in restaurants broken into four categories of service: quick-service restaurants, midscale restaurants, casual dining, and fine dining/upscale hotel. These categories are designed to represent different levels of quality.

These data can shed light on the appropriateness of our assumption that the price of a good or service is a good proxy for its quality. If we sort firms using the average price of a meal as a proxy for quality, we obtain a sorting by quality tiers similar to

NPD's.[15] The average price of dinner (lunch) is \$6.5 (\$5.8) in quick-service restaurants, \$11.2 (\$9.2) in midscale restaurants, and \$14.9 (\$11.7) in casual dining.[16]

We find clear evidence of trading down in the NPD data. Consider first the number of meals served. Table 5 shows that the percentage of meals served by quick-service restaurants increased from 76.1 percent in 2007 to 78.2 percent in 2012. At the same time, the fraction of meals served declined in all the other segments: midscale, casual and fine-dining.[17] Table 6 reports results for market share. We see that over the period 2007-2012 the market share of quick-service restaurants rose from 57.7 percent to 60 percent. At the same time, the market share declined in all the other segments.[18]

Unfortunately, we cannot do our accounting calculations with these data because we do not have the breakdown of employment across the different segments used by NPD.

2.5 Substitution across categories

In our analysis, we focus on the implications of trading down for employment. We also studied the employment implications of substitution across categories, for example from luxuries to necessities. Our analysis is based on CEX and NIPA PCE data (see Appendix C for more detail on the calculations).

It is well known that different categories of expenditure have different income elasticities. As a result they differ in their cyclical properties. For example, expenditure

[15]There is a literature on the role of search frictions in generating price dispersion. While search frictions are clearly important, the price differences across categories in our dataset are clearly too large to be accounted for by these frictions alone. Aguiar and Hurst (2006) estimate that doubling of shopping frequency lowers the price paid for a given good by 7 to 10 percent. The price differences across different categories in our data are almost an order of magnitude larger than these estimates.

[16]These price data were collected in March 2013. We do not have average meal prices for fine-dining restaurants.

[17]There is also some evidence in the NDP data that consumers traded down in terms of the meal they choose to eat at restaurants, eating out at breakfast and lunch instead of at dinner.

[18]Tables 5 and 6 show that after the worst of the recession was over in 2010, fine dining started to recover. But overall, the fraction of meals served and market share of fine dining are still lower in 2012 than in 2007.

Table 5: Percentages of Restaurant Traffic by Year and Quality Segment

Quality segment	2007	2008	2009	2010	2011	2012
Quick service restaurants	76.1	76.4	76.8	77.2	77.8	78.2
Midscale	11.4	11.1	11.0	10.7	10.3	10.0
Casual dining	11.2	11.1	11.1	10.9	10.7	10.4
Fine dining/upscale hotel	1.4	1.4	1.2	1.2	1.3	1.3

Note: The data is from NPD Group.

Table 6: Restaurant Market Share by Year and Quality Segment

Quality segment	2007	2008	2009	2010	2011	2012
Quick service restaurants	57.7	58.0	58.7	59.0	59.4	60.0
Midscale	15.4	15.2	15.1	14.8	14.5	14.1
Casual dining	21.5	21.4	21.4	21.3	20.9	20.5
Fine dining/upscale hotel	5.5	5.3	4.9	5.0	5.2	5.4

Note: The data is from NPD Group.

on Food away from Home falls during recessions by much more than expenditures on Personal Care.

We find that substitution across categories has a negligible effect on employment. This result is driven by the low correlation between income elasticities of different categories and labor intensity is quite low. For example, both Food away from home and Vehicle purchases fall during recessions. But Food away from Home has high labor intensity, while Vehicle Purchases has low labor intensity. We summarize our results in Appendix C.

3 A model

In this section, we interpret our results using a model where households choose both the quantity and quality of the goods they consume. The goal of the model is to show how the choice of the quality of the goods consumed affects employment in general equilibrium.

In the model households face a natural trade-off: consuming higher quality goods yields higher utility but costs more. In order to be consistent with our empirical findings, the model has to satisfy five properties. First, there must be goods of different quality consumed in equilibrium. Second, employment must be endogenously determined. Third, prices must increase with quality. Fourth, labor intensity must increase with quality. Finally, it is natural to require that richer households consume higher quality goods.

We consider below a simple general equilibrium model with heterogeneous agents that satisfies these properties. We first derive some analytical results that highlight the impact of quality choice on employment. We then turn to a calibrated version of the model where we quantify the impact of quality trade down on aggregate employment.

3.1 Households

We assume that individuals differ in their ability. They are endowed with different levels of efficiency units of labor, ϵ, distributed according to the cumulative distribution function $\Gamma(\epsilon)$.

Utility A consumer with ability ϵ derives utility U_ϵ from both the quantity (C_ϵ) and quality (q_ϵ) of the goods consumed:

$$U_\epsilon = U(C_\epsilon, q_\epsilon). \tag{8}$$

The consumer maximizes the above utility subject to the following general static budget constraint:

$$P_\epsilon C_\epsilon = Y_\epsilon, \tag{9}$$

where P_ϵ is the price of one unit of consumption of quality q_ϵ and Y_ϵ is the income of an individual with ability ϵ. As usual, we assume that utility increases with the quantity consumed:

$$U_1(C_\epsilon, q_\epsilon) > 0, \tag{10}$$

We also assume that the marginal utility of quality is positive:

$$U_2(C_\epsilon, q_\epsilon) > 0. \tag{11}$$

The ratio of the first-order condition of the consumer problem is given by:

$$\frac{U_2(C_\epsilon, q_\epsilon)}{U_1(C_\epsilon, q_\epsilon)} = \frac{P'_\epsilon}{P_\epsilon} C_\epsilon, \tag{12}$$

where P'_ϵ, denotes the derivative of the price function with respect to quality (we return to this expression below).

Our maintained assumption is that quality is a normal good, so that higher income consumers choose goods of higher quality. While this condition seems natural, it imposes restrictions on the form of the utility function. The first-order condition above implies that if $U(C_\epsilon, q_\epsilon)$ is homogeneous in C_ϵ, then quality is independent of income. So, in order for quality to be a normal good, U must be non-homothetic in consumption.

With this requirement in mind, we assume that the momentary utility of an individual with ability ϵ takes the form:

$$U_{\epsilon,t} = q_{\epsilon,t} \log(C_{\epsilon,t}). \tag{13}$$

An advantage of this functional form is that it nests the usual logarithmic utility as a special case. This property simplifies the comparison of versions of the model with and without quality choice.

In what follows, we concentrate on the short-run behavior of the economy. To simplify, we assume that there is a fixed amount of a productive durable good which can be interpreted as land.[19]

The budget constraint of an individual with ability ϵ is given by:

$$P_{\epsilon,t}C_{\epsilon,t} = [I_{\epsilon,t} \times W_t\epsilon + (1 - I_{\epsilon,t}) \times V] + k_{\epsilon,t} - \tau_t, \tag{14}$$

where $P_{\epsilon,t}$ denotes the price of a good of quality $q_{\epsilon,t}$ and I_t is an indicator function that equals one if the individual chooses to work in period t and zero otherwise. The variable k_t denotes the individual holdings of the productive durable asset. Without loss of generality, we choose this asset to be the numeraire. The variable τ_t denotes the level of lump-sum taxes.

The first-order condition for $C_{\epsilon,t}$ is:

$$\frac{q_{\epsilon,t}}{C_{\epsilon,t}} = \lambda_{\epsilon,t}P_{\epsilon,t}, \tag{15}$$

where $\lambda_{\epsilon,t}$ denotes the Lagrange multiplier associated with the budget constraint.

[19] Solving a version of the model with aggregate uncertainty and capital accumulation requires an algorithm that can handle a large number of heterogeneous agents and a large number of sectors (one for each quality produced). To the best of our knowledge, this problem has not been solved in the literature. The key difficulty is that solving for the competitive equilibrium requires making sure that the sum of capital and labor demand across sectors equals the supply.

The first-order conditions for $q_{\epsilon,t}$, the quality that is optimal for an individual with skill ϵ, is:

$$\log(C_{\epsilon,t}) = \lambda_{\epsilon,t} P'_{\epsilon,t} C_{\epsilon,t}, \tag{16}$$

Combining equations (15) and (16), we obtain:

$$C_{\epsilon,t} = \exp\left(\frac{q_{\epsilon,t} P'_{\epsilon,t}}{P_{\epsilon,t}}\right). \tag{17}$$

Labor Supply In order to study the effect of quality choice on the labor market, employment must be endogenous. We assume that individuals make a discrete choice of whether or not to work. Individuals who do not work receive a transfer V from the government that is independent of ϵ. An employed worker with ability ϵ earns ϵW_t, where W_t denotes the wage per efficiency unit of labor.

This implies that the decision to work is governed by a cut-off rule. We denote by ϵ_t^* the ability of an agent who is indifferent between working and not working:

$$\epsilon_t^* W_t = V. \tag{18}$$

At each point in time, individuals with ability $\epsilon < \epsilon_t^*$ choose not to work and receive a transfer V from the government. Those with ability $\epsilon \geqslant \epsilon_t^*$ choose to work. The fraction of individuals out of the labor force is $\Gamma(\epsilon_t^*)$.

3.1.1 Government

The government finances its transfers to those out of the labor force by levying a lump-sum tax, τ_t. The budget constraint of the government is:

$$V\Gamma(\epsilon_t^*) = \tau_t. \tag{19}$$

3.2 Production

We assume that goods of all quality levels are produced by perfectly competitive producers. To produce $Y_{\epsilon,t}$ units of a consumption good with quality $q_{\epsilon,t}$, producers combine labor and land according to the following CES production function:

$$Y_{\epsilon,t} = A_t \left[\alpha \left(\frac{H_{\epsilon,t}}{q_{\epsilon,t}} \right)^{\rho} + (1-\alpha) \left(K_{\epsilon,t} \right)^{\rho} \right]^{\frac{1}{\rho}}, \tag{20}$$

where A_t denotes the aggregate level of total factor productivity (TFP). The variables $H_{\epsilon,t}$ and $K_{\epsilon,t}$ denote the labor and land used by the producer, respectively. We assume that $\rho < 0$ so there is less substitution between the two factors of production that in a Cobb-Douglas production function. As we show below, this assumption is necessary so that, as in the data, higher quality goods are more labor intensive.

The producer's problem is:

$$\max P_{\epsilon,t} Y_{\epsilon,t} - W_t H_{\epsilon,t} - K_{\epsilon,t}. \tag{21}$$

This problem implies that the price schedule, $P_{\epsilon,t}$, is given by:

$$P_{\epsilon,t} = \frac{1}{A_t} \left[\alpha^{\frac{1}{1-\rho}} (q_{\epsilon,t} W_t)^{\frac{\rho}{\rho-1}} + (1-\alpha)^{\frac{1}{1-\rho}} \right]^{\frac{\rho-1}{\rho}}. \tag{22}$$

The elasticity of the price with respect to quality is:

$$\frac{P'_{\epsilon,t}(q_{\epsilon,t}) q_{\epsilon,t}}{P_{\epsilon,t}} = \frac{1}{1 + \left(\frac{1-\alpha}{\alpha}\right)^{\frac{1}{1-\rho}} \left(\frac{1}{q_{\epsilon,t} W_t} \right)^{\frac{\rho}{\rho-1}}}. \tag{23}$$

Given our assumption that $\rho < 0$, both the price and the price elasticity are increasing in quality.

Combining the optimal demand for labor by the firm and the price schedule (equation (22)) we obtain the following expression of the labor intensity measure we use in our empirical work:

$$\frac{H_{\epsilon,t}}{P_{\epsilon,t} Y_{\epsilon,t}} = \frac{\left[\frac{(1-\alpha)}{\alpha} W_t \right]^{\frac{1}{\rho-1}}}{\left(\frac{\alpha}{1-\alpha}\right)^{\frac{1}{1-\rho}} W_t^{\frac{\rho}{\rho-1}} + q_{\epsilon,t}^{\frac{\rho}{1-\rho}}}.$$

For $\rho < 0$ this measure of labor intensity is increasing in quality.

Equation (17) implies that both quality and quantity consumed increase with income.[20]

[20]To see this, replace the expression from equation ((17)) into the budget constraint, obtaining

3.3 Equilibrium

The equilibrium definition is standard. Households maximize utility taking the wage rate per efficiency unit and the price-quality schedule as given. The household optimization conditions are summarized by equations (15) and (16), together with the budget constraint (14) and cut-off condition (18). Firms maximize profits and their first-order condition are summarized by the solution to the problem in equation (21).

The labor market clears, so total demand for labor equals total supply:

$$\int_{\epsilon_{min}}^{\epsilon^{max}} H_{\epsilon,t} d\epsilon = \int_{\epsilon_t^*}^{\epsilon^{max}} \epsilon \Gamma'(\epsilon) d(\epsilon). \tag{24}$$

The land market clears, so total demand for land equals total supply:

$$\int_{\epsilon_{min}}^{\epsilon^{max}} K_{\epsilon,t} d\epsilon = \int_{\epsilon_{min}}^{\epsilon^{max}} k(\epsilon)\Gamma'(\epsilon) d(\epsilon) = K. \tag{25}$$

Finally, the budget constraint of the government (equation (19)) must hold. Since the rental price of the fixed factor is the numeraire, real output, Y_t, is given by:[21]

$$Y_t = \int_{\epsilon_{min}}^{\epsilon_{max}} P_{\epsilon,t} Y_{\epsilon,t} d(\epsilon). \tag{26}$$

3.4 The importance of quality choice

We can show analytically that the presence of quality choice amplifies the response of the real wage and employment to a productivity shock. To derive these analytical results we assume that land holdings are distributed in proportion to individual ability, so more productive individuals hold an higher share of the productive asset:

$$k_\epsilon = \frac{\epsilon}{\int_{\epsilon_{min}}^{\epsilon^{max}} \epsilon \Gamma'(\epsilon) d\epsilon} K = \frac{\epsilon}{E(\epsilon)} K.$$

$P_{\epsilon,t} C_{\epsilon,t} = P_{\epsilon,t} \exp\left(\frac{q_{\epsilon,t} P'_{\epsilon,t}}{P_{\epsilon,t}}\right)$ which is increasing with quality. Since quality increases with income it follows from equation ((17)) that consumption is an increasing function of income as well.

[21] This expression assumes that real output is computed using hedonic adjustments: when the price of consumption rises, the statistical authorities recognize that this rise is solely due to an increase in the quality of the goods consumed.

This assumption simplifies the analysis, since an individual's income depends only on ϵ, but it is not essential to our results.

In this case, the budget constraint for a worker with ability ϵ is:

$$P_{\epsilon,t}C_{\epsilon,t} = \epsilon\left[W_t + \frac{K}{E(\epsilon)}\right] - \tau_t, \tag{27}$$

while budget constraint of a non-working individual with ability ϵ is:

$$P_{\epsilon,t}C_{\epsilon,t} = \left[V + \epsilon\frac{K}{E(\epsilon)}\right] - \tau_t. \tag{28}$$

The market-clearing condition for goods consumed by individuals with ability ϵ is:

$$Y_{\epsilon,t} = A_t\left[\alpha\left(\frac{H_{\epsilon,t}}{q_{\epsilon,t}}\right)^{\rho} + (1-\alpha)\left(K_{\epsilon,t}\right)^{\rho}\right]^{\frac{1}{\rho}} = C_{\epsilon,t}\Gamma'(\epsilon). \tag{29}$$

Finally, the firms' first-order condition implies:

$$\frac{H_{\epsilon,t}}{K_{\epsilon,t}} = \left[W_t\frac{(1-\alpha)}{\alpha}\right]^{\frac{1}{\rho-1}}\left(q_{\epsilon,t}\right)^{\frac{\rho}{\rho-1}}. \tag{30}$$

3.4.1 Local analysis

We use a circumflex to denote logarithmic deviations of a variable from its steady state value. A variable without a time subscript denotes its steady state value.

The percentage deviation of the price of the quality consumed by an individual with ability ϵ is:

$$\hat{P}_{\epsilon} = \frac{\left(\frac{\alpha}{1-\alpha}\right)^{\frac{1}{1-\rho}}\left(Wq_{\epsilon}\right)^{\frac{\rho}{\rho-1}}}{\left(\frac{\alpha}{1-\alpha}\right)^{\frac{1}{1-\rho}}\left(Wq_{\epsilon}\right)^{\frac{\rho}{\rho-1}} + 1}\left(\hat{W} + \hat{q}_{\epsilon}\right) - \hat{A}. \tag{31}$$

Given equations ((29))-((30)), the deviations of consumption and output from their steady state values are given by:

$$\hat{C}_{\epsilon} = \hat{Y}_{\epsilon} = \hat{K}_{\epsilon} + \hat{A} + \frac{\left[\left(\frac{\alpha}{1-\alpha}\right)^{\frac{1}{1-\rho}}\left(Wq_{\epsilon}\right)^{\frac{\rho}{\rho-1}}\right]}{\left[\left(\frac{\alpha}{1-\alpha}\right)^{\frac{1}{1-\rho}}\left(Wq_{\epsilon}\right)^{\frac{\rho}{\rho-1}} + 1\right]}\frac{1}{\rho-1}\left(\hat{W} + \hat{q}_{\epsilon}\right). \tag{32}$$

The linearized budget constraint for workers is:

$$\hat{P}_\epsilon + \hat{C}_\epsilon = \frac{\epsilon W}{P_\epsilon C_\epsilon}\hat{W} - \frac{\tau}{P_\epsilon C_\epsilon}\hat{\tau}. \tag{33}$$

The linearized budget constraint for a non-working individuals is:

$$\hat{P}_\epsilon + \hat{C}_\epsilon = -\frac{\tau}{P_\epsilon C_\epsilon}\hat{\tau}. \tag{34}$$

Here $\hat{\tau}$ denotes the percentage deviations of the tax rate implied by the government budget constraint. Finally, the government's budget constraint implies that the lump sum tax satisfies:

$$\hat{\tau} = \left[\frac{\Gamma'(\epsilon^*)\epsilon^*}{\Gamma(\epsilon^*)}\right]\hat{\epsilon^*}. \tag{35}$$

3.4.2 Model without quality choice

We now compare the response to a negative TFP shock in economies with and without quality choice. Consider first an economy without quality choice. To make this economy as comparable as possible to our model with quality choice we assume that the single consumption good is produced with the same CES production function described in equation (29):

$$Y_t = A_t \left[\alpha \left(H_t\right)^\rho + (1-\alpha)\left(K\right)^\rho\right]^{\frac{1}{\rho}} = C_t. \tag{36}$$

Total hours worked in the economy are given by:

$$H_t = \int_{\epsilon_t^*}^{\epsilon^{max}} \epsilon\Gamma'(\epsilon)d(\epsilon),$$

where ϵ_t^* is determined by the cut-off condition:

$$\epsilon_t^* W_t = V.$$

The firms' problem implies that the optimal ratio of production factors is:

$$\frac{H_t}{K} = \left[W_t\frac{(1-\alpha)}{\alpha}\right]^{\frac{1}{\rho-1}}. \tag{37}$$

Substituting the expression for efficiency hours and the cut-off condition, we obtain the following equation for the cut-off value,

$$(\epsilon_t^*)^{\frac{1}{\rho-1}} \int_{\epsilon_t^*}^{\epsilon^{max}} \epsilon\Gamma'(\epsilon)d(\epsilon) = K\left[V\frac{(1-\alpha)}{\alpha}\right]^{\frac{1}{\rho-1}}. \tag{38}$$

The right-hand side of this equation is constant while the left-hand side is monotonically decreasing in ϵ^*.[22] Hence, there is only one possible intersection. Since, ϵ_t^* is constant, the real wage is also constant and given by:

$$W = \frac{V}{\epsilon^*}.$$

In sum, in an economy without quality choice a TFP shock does not affect the real wage or the level of employment.

3.4.3 Model with quality choice

In what follows we show that in the model with quality choice employment and the real wage must fall in response to a negative TFP shock. Since the model has a continuum of consumption sectors indexed by ϵ, we limit the analysis to a symmetric equilibrium in which the changes in all sectors are of the same sign.

Ruling out a constant real wage rate Consider first the case where the wage does change. The cut-off condition implies that ϵ^* is constant, i.e. $\hat{\epsilon^*} = 0$. Equation (35) implies then that the tax rate is also constant. Equations (33) and (34) imply that the household income is constant and thus: $\hat{P}_\epsilon + \hat{C}_\epsilon = 0$.

It is useful to combine equations (31), (32) and (33) to rewrite a linearized version of the budget constraint of a worker with ability ϵ as:

$$\left(\frac{\rho}{\rho-1}\right)\frac{\left[\left(\frac{\alpha}{1-\alpha}\right)^{\frac{1}{1-\rho}}(Wq_\epsilon)^{\frac{\rho}{\rho-1}}\right]}{\left[\left(\frac{\alpha}{1-\alpha}\right)^{\frac{1}{1-\rho}}(Wq_\epsilon)^{\frac{\rho}{\rho-1}}+1\right]}\left(\hat{W}+\hat{q}_\epsilon\right)+\hat{K}_\epsilon = \hat{P}_\epsilon+\hat{C}_\epsilon = \frac{\epsilon W}{P_\epsilon C_\epsilon}\hat{W}-\frac{\tau}{P_\epsilon C_\epsilon}\hat{\tau}. \tag{39}$$

[22]To see this result, apply Leibnitz's rule to obtain the derivative of the left-hand side with respect to ϵ^*: $\frac{1}{\rho-1}(\epsilon^*)^{\frac{2-\rho}{\rho-1}}\int_{\epsilon^*}^{\epsilon^{max}}\epsilon\Gamma'(\epsilon)d\epsilon - (\epsilon^*)^{\frac{\rho}{\rho-1}}\Gamma'(\epsilon^*)$.

Similarly, we can rewrite the budget constraint of a non-working individual with ability ϵ as:

$$\left(\frac{\rho}{\rho-1}\right)\frac{\left[\left(\frac{\alpha}{1-\alpha}\right)^{\frac{1}{1-\rho}}\left(Wq_\epsilon\right)^{\frac{\rho}{\rho-1}}\right]}{\left[\left(\frac{\alpha}{1-\alpha}\right)^{\frac{1}{1-\rho}}\left(Wq_\epsilon\right)^{\frac{\rho}{\rho-1}}+1\right]}\left(\hat{W}+\hat{q}_\epsilon\right)+\hat{K}_\epsilon=-\frac{\tau}{P_\epsilon C_\epsilon}\hat{\tau}. \tag{40}$$

Suppose that quality increases. Since the wage and the tax rate do not change, the right-hand side of equations (39)-(40) equals zero. So, an increase in quality must be accompanied by a fall in the land employed in the production of each quality level. This property is inconsistent with the market-clearing condition. A similar argument implies that quality cannot decrease.

Finally, to rule out the case where both the wage and quality do not change, we note that the household optimization (equation (17)) implies that:

$$\hat{C}_\epsilon=\frac{\left(\frac{1-\alpha}{\alpha}\right)^{\frac{1}{1-\rho}}\left(\frac{1}{Wq_\epsilon}\right)^{\frac{\rho}{\rho-1}}}{\left[1+\left(\frac{1-\alpha}{\alpha}\right)^{\frac{1}{1-\rho}}\left(\frac{1}{Wq_\epsilon}\right)^{\frac{\rho}{\rho-1}}\right]^2}\frac{\rho}{\rho-1}\left(\hat{W}+\hat{q}_\epsilon\right). \tag{41}$$

Equation (41) implies that if the wage and quality are constant then consumption is also constant. However, since we argued above that $\hat{P}_\epsilon+\hat{C}_\epsilon=0$, prices are also constant. This property contradicts equation (31) which implies that if q and W do not change, then prices must increase given the fall in TFP. Hence, the wage cannot remain constant in response to a TFP shock.

Ruling out an increase in the real wage Suppose that the real wage goes up. Then the cut-off condition implies that $\hat{\epsilon}_t^*<0$. Equation (35) implies that there is a fall in the level of lump sum taxes–since more individuals are working, the value of τ required to fund the payments to the non-working falls. It follows from equations (39)-(40) that, for each consumer, $\hat{P}_\epsilon+\hat{C}_\epsilon>0$.

Consider first the case where quality increases. Then from equation (41) it follows that consumption increases for every consumer. However, equation (32) implies that the land hired by producers increases for all values of ϵ. This property is inconsistent with clearing the land market. A similar argument rules out the possibility that quality stays constant.

To rule out the case where the real wage increases and quality falls, we proceed as follows. First, consider the case where quality falls but this fall is not sufficient for consumption to fall. Then from equation (41) it follows that in each market $\hat{W} + \hat{q}_\epsilon > 0$. This condition together with equation (32) implies that land use increases for all producers, violating the market clearing condition.

Second, consider the case where quality falls so as to offset the increase in the real wage. In that case, equation (41) implies that consumption does not change. Again, this property together with equation (32) implies that the land use has to increase for all ϵ, violating the market clearing condition.

Finally, consider the case where quality falls by more than the real wage. Since equations (39)-(40) imply that the overall income of each agent increases then since $\hat{W} + \hat{q}_\epsilon < 0$ then land use has to increase for every ϵ, violating the market clearing condition.

Wages must fall By a process of elimination we established that real wages fall in response to a negative TFP shock. This fall leads to an increase in the cut-off value and a fall in employment. These properties stand in contrast to those of the model without quality choice where employment and the real wage remain constant in response to a TFP shock.

4 Numerical Analysis

In this section, we quantify the impact of quality choice on the labor market. We begin by describing the calibration of the model economy and the solution algorithm used to compute the equilibrium.

4.1 Calibration and Solution Algorithm

Without loss of generality, we normalize both the level of TFP and the supply of land to one ($A = 1$, $K = 1$). To match the fat right tail of the U.S. income distribution, we assume that $\Gamma(\epsilon)$ is a Pareto distribution with shape parameter κ and scale parameter ϵ_{min}. We normalize $\epsilon_{min} = 1$ and, following Diamond and Saez (2011), we choose $\kappa = 1.5$. We calibrate ϵ_{max} to match the income ratio of the top 5 percent to the second income quintile between 2010-2014. This empirical ratio equals 4.88.[23] We then discretize the support of the distribution, $[\epsilon_{min}\ , \epsilon_{max}]$, with half a million grid points, resulting in a difference between two adjacent points equal to 0.000058.

Solution Algorithm The remaining three parameters are ρ, α, and V. As we discuss below, we iterate until in equilibrium we match the empirical values of the ratio of: (i) the labor intensity in the "middle quality" market to the labor intensity in the "low quality" market, (ii) the ratio of the labor intensity in the "high quality" market to the labor intensity in the "middle quality" market, and (iii) a value of ϵ^* such that the model is consistent with the 63 percent employment rate of the 16+ population in 2007, i.e. $\Gamma(\epsilon^*) = 0.37$.[24]

Given an initial guess for these two parameters, we guess a value for the real wage, W. Given that we target $\Gamma(\epsilon^*) = 0.37$ we compute the value of ϵ^*. This value, together with the cut-off rule gives us the value of V and thus the value of τ. With these values we obtain the income for each consumer so we can solve for their optimal quality and quantity of consumption.

Using equations (29)-(30) we recover the land and hours worked employed by the producers of the quality associated with each value of ϵ. We check whether the market clearing condition for hours worked and land (equations (24) and (25)) hold. We iterate

[23]See the data in http://www.taxpolicycenter.org/statistics/historical-income-distribution-all-households

[24]We match the 2007 labor intensity weighted average estimates reported in Table 1, $\log(9.23/6.33) = 0.38$ and $\log(10.86/9.23) = 0.16$.

on the wage rate until these market-clearing conditions hold together with the cut-off rule and the household budget constraint.

Once we clear the two factor markets, we calculate the two moments discussed above. Specifically, we first find the quality threshold, q_{low}, below which we have 35 percent of total consumption, i.e.

$$\sum_{j=q_{min}}^{q_{low}} \frac{P_j Y_j}{\sum_{z=q_{min}}^{q_{max}} P_z Y_z} = 0.35.$$

Similarly, we find the quality level q_{middle} such that consumption in the middle quality segment accounts for 58 percent of total consumption, i.e.,

$$\sum_{j=q_{low}}^{q_{middle}} \frac{P_j Y_j}{\sum_{z=q_{min}}^{q_{max}} P_z Y_z} = 0.58.$$

The remaining 7 percent of consumption correspond to the high-quality segment. These values are the market shares we estimated using our data.[25] As in our empirical analysis, we calculate the revenue-weighted average of the labor intensity within each of the three quality categories, i.e.

$$LI_{low} = \sum_{j=q_{min}}^{q_{low}} \left(\frac{H_j}{P_j Y_j}\right)\left(\frac{P_j Y_j}{\sum_{z=q_{min}}^{q_{low}} P_z Y_z}\right),$$

$$LI_{middle} = \sum_{j=q_{low}}^{q_{middle}} \left(\frac{H_j}{P_j Y_j}\right)\left(\frac{P_j Y_j}{\sum_{z=q_{low}}^{q_{middle}} P_z Y_z}\right),$$

$$LI_{high} = \sum_{j=q_{middle}}^{q_{max}} \left(\frac{H_j}{P_j Y_j}\right)\left(\frac{P_j Y_j}{\sum_{z=q_{middle}}^{q_{max}} P_z Y_z}\right).$$

We iterate on ρ and α and repeat all the steps discussed above until we match the following two empirical moments

$$\log\left(\frac{LI_{middle}}{LI_{low}}\right) = 0.38,$$

[25] Again, as with the ratios of the labor intensities we match the 2007 weighted average numbers in Table 1.

$$\log\left(\frac{LI_{high}}{LI_{middle}}\right) = 0.16.$$

Our algorithm yields the following values: $\alpha = 0.69$ and $\rho = -0.7$.

4.2 Results

Figure 1 shows that wealthier individuals consume more quantity and quality and that higher levels of quality have higher prices and are produced with higher labor intensities. The top left panel in Figure 1 depicts the optimal quality choice by income level (ϵ). The top right panel depicts the price of the consumption good chosen by an individual with ability ϵ. The left bottom panel depicts the quantity of the good that consumed by this individual. Finally, in the right bottom panel we depict the the labor intensity of the good that the individual consumes.

Response to a TFP shock Recall that in the model without quality choice, a change in TFP has no impact on the level of employment or the real wage rate. Now consider the response of the economy with quality choice to a decline in TFP. For every value of the shock we consider below, we use the model's predictions for market shares and labor intensities and we compute the changes in employment that are due to trading down by following the same empirical approach used in Section 2, summarized by equations (4) and (6).

Our findings are as follows. Consider first the case that A falls by 1 percent to $A = 0.99$. Overall, the employment rate falls from 63 percent to 62.5 percent. We find that absent trading down the employment rate would have fallen to 62.77 percent. Thus, trading down accounts for around half of the observed changes in employment. The property that half of the change in employment is explained by trading down is robust to different shock magnitudes.[26]

[26] For example, for the case where A falls by 2 percent to $A = 0.98$ we find overall employment to fall from 63 percent to 62 percent where absent trading down employment would have fallen to 62.5 percent. Similarly, for bigger shocks where A for example falls to 0.94 we find that overall employment

Figure 1: Optimal Choices and Implications

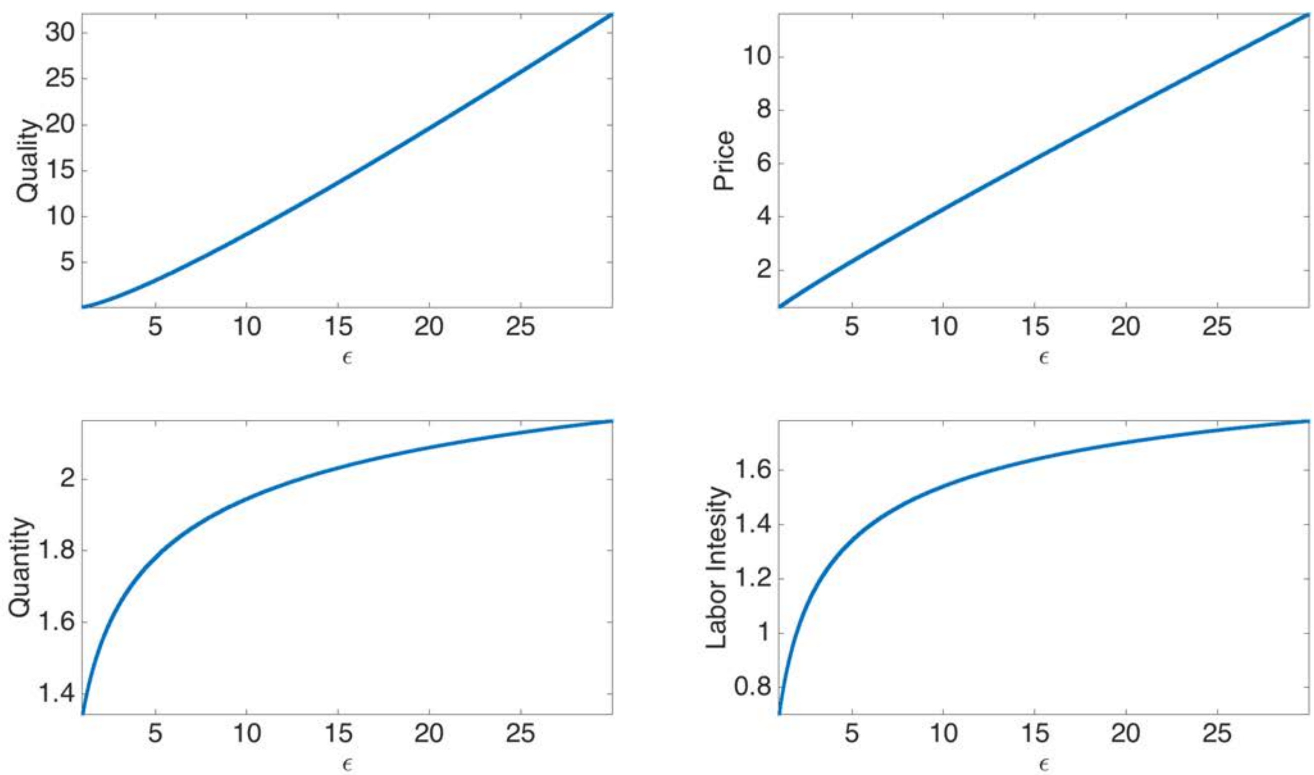

Recall that we found in the Yelp! data that around half of the observed fall in employment was due to cyclical trading down. Similarly, in the PPI data we found trading down to also account for approximately half of the fall in employment. Thus, our model, is roughly consistent with these estimates even though they were not targeted in the calibration.

to fall from 63 percent to 60 percent where absent trading down employment would have fallen to 61.6 percent.

5 Conclusion

In this paper, we show that during the Great Recession consumers traded down in the quality of the goods and services they consumed. We also show that lower quality products are generally less labor intensive, so trading down reduces the demand for labor. Our calculations suggest that trading down accounts for about half of the decline in employment during the Great Recession.

We introduce quality choice into a general equilibrium model where consumers optimally chose the quality and quantity of consumption. We show that in equilibrium a decline in TFP is associated trading down in quality and declines in employment that are broadly consistent with our empirical findings.

6 References

Aguiar, Mark and Erik Hurst "Lifecycle Production and Prices," *The American Economic Review*, 97(5): 1533-59, 2007.

Aguiar, Mark, Erik Hurst and Loukas Karabarbounis "Time use during the Great Recession," *The American Economic Review*: 1664-1694, 2013.

Burns, Arthur and Wesley Mitchell *Measuring Business Cycles*, New York, National Bureau of Economic Research, 1946.

Burstein, Ariel, Eichenbaum, Martin, and Sergio Rebelo, 2005. Large devaluations and the real exchange rate. Journal of Political Economy 113 (4), 742–784.

Chen, Natalie and Luciana Juvenal "Quality and the Great Trade Collapse," manuscript, International Monetary Fund, 2016.

Coibion, Olivier, Yuriy Gorodnichenko and Gee Hee Hong "The Cyclicality of Sales, Regular and Effective Prices: Business Cycle and Policy Implications" *The American Economic Review,* 105, 993-1029, 2015.

Diamond, Peter, and Emmanuel Saez, 2011. "The Case for a Progressive Tax: from Basic Research to Policy Recommendations" *The Journal of Economic Perspectives*, 165-190.

Gilchrist, Simon, Jae Sim, Egon Zakrajsek, and Raphael Schoenle "Inflation dynamics during the financial crisis," manuscript, 2015.

Gorodnichenko, Yuriy and Michael Weber "Are sticky prices costly? Evidence from the stock market," manuscript, University of California, Berkeley, 2015.

Kaplan, Greg and Guido Menzio "Shopping externalities and self-fulfilling unemployment fluctuations," *Journal of Political Economy*, 124(3), 2016.

Nevo, Aviv and Arlene Wong "The elasticity of substitution between time and market goods: Evidence from the Great Recession," manuscript, Northwestern University, 2015.

Weber, Michael "Nominal rigidities and asset pricing," manuscript, University of Chicago, 2015.

Appendix

In this appendix, we provide some additional details about the construction of the Yelp! and PPI data set and discuss the computation of the indirect effects of trading down on employment. We also analyze the relation between quality and income and quality and employment implied by our model. In addition, we discuss a version of the model that is consistent with balanced growth.

A Yelp! data

We scraped data for Yelp in April 2014. For firms that own more than one brand, we compute the average price category for each brand and then compute the average price category for the firm, weighting each brand by their sales volume. One concern about this procedure is that we might be averaging high-quality and low-quality brands. In practice, this situation is rare: 73 percent of the firms in our sample have a single brand. For multi-brand firms, 54 percent have all their brand in the same price category. For example, the firm Yum! Brands owns three brands (Taco Bell, KFC, and Pizza Hut), but they are all in the same price category (low price). For robustness, we redid our analysis including only firms that either have a single brand or have all their brands in the same price category. We obtain results that are very similar to those we obtain for the whole universe of firms.

In merging the data with Compustat we note that for companies with operations outside of the U.S., we use the information on sales by business region to compute U.S. sales. We also use the break down of employment by business region to compute labor intensity in the U.S. We exclude from our sample manufacturing firms for which this breakdown is not available. For retail firms, foreign operations are generally small, so we include companies with foreign operations in our sample. As we robustness check, we redo our analysis excluding these companies. The results are similar to those we obtain for the full sample.

Table 7 presents some description of the data used to analyze quality shifts in expenditure in five retail sectors. It describes the data source (column I), the number of firms covered in the sample in 2007 (II), the average annual firm sales revenue (III), and the percent of the overall sector sales that our sample covers (IV).

Table 7: Data Sample Description

Sectors	Data Source	Number of Firms	2007 Annual Sales of Average Firm ($m)	% of U.S. Sector
	(I)	(II)	(III)	(IV)
Apparel	Compustat and company reports	54	1,648	41%
Grocery stores	Compustat and company reports	9	34,348	56%
Restaurants	Compustat and company reports	74	1,012	19%
Home furnishing	Compustat	41	4,750	39%
General Merchandise	U.S. Census	n.a.	n.a.	100%

Note: This table describes for each sector the data source used (I), the number of firms within the sample (II), and the average annual sales of each firm (III). (IV) reports the share of the sales of the entire sector that our data set covers.

B PPI

Using the PPI data presents two challenges. First, in the PPI, firms in the same industry report prices that correspond to different units of measurement, e.g. some firms report price per pound, others price per dozen. To circumvent this problem, we first convert prices into a common metric whenever possible (for example, converting ounces to into pounds). We then compute the modal unit of measurement for each 6-digit NAICS category and restrict the sample to the firms that report prices for this model unit.

This filtering procedure preserves 2/3 of the original data, which is comprised of 16,491 establishments out of a sample of about 25,000 establishment surveyed by the PPI.[27]

Second, some of the firms included in the PPI data offshore their production, so their reported employment does not generally include production workers. It includes primarily head-office workers and sales force in the U.S. Using information in the firms' annual reports, we exclude firms that have most of their production offshore. The resulting data set preserves over half of the merged PPI/Compustat data.

C Substitution across categories

Our main analysis focuses on trading-down behavior within categories. We also examined substitution across consumption categories and the effect on employment. We use data from the CEX Survey and NIPA personal consumption expenditures (PCE). We consider 31 different consumption categories (see Figure 2).

We examine the effect of across-category substitution on employment in two steps. First, we construct the labor intensity for each consumption category. We match CEX consumption categories with the NIPA PCE commodity definitions.[28] This allows us to the Input/Output commodity-level data to construct labor intensity measures for each consumption category. Second, we compute the change in budget share for each consumption category over 2007-2012 for the average household. To isolate out the cyclical component of the budget reallocation across consumption categories, we estimate elasticities of the category budget shares to total household expenditure. We then multiply the elasticities by the actual change in household expenditure to obtain the change in budget allocation for each category.

We derive shifts in expenditure associated with the recessionary drop in household

[27] Note some establishments are excluded because we only include items that are recorded at the modal unit of measure within the 6-digit product category.

[28] Further details of this matching process is available upon request.

income by estimating the following Engel curve elasticities:

$$w^k_{ht} = \alpha^k + \beta^k \ln(X_{ht}) + \sum_j \gamma_j \ln(P_{jt}) + \theta^k_{ht} \cdot Z_{ht} + \epsilon^k_{ht}, \tag{42}$$

where w^k_{ht} is the budget share allocated to category k by household h at time t; X_{ht} is total household expenditure; and P_{jt} is the price index of each expenditure category j at time t. The variable Z_{ht} is a vector of household demographics variables, including the age and square of the age of the head of household, dummies based on the number of earners (<2,2+), and household size (<2,3-4,5+). The error term is denoted by ϵ^k_{ht}. We estimate equation (42) using household sample weights given in the CEX data based on the 1980-2012 waves of the CE Surveys. The coefficient β^k gives the fraction change in budget share allocation to expenditure category k, given a 100 percentage point change in total household expenditure.[29]

[29]There potential issues that arise in estimating Equation (42). For instance, mismeasurement of individual goods may be cumulated into total expenditure, which would bias the estimated coefficients. Aguiar and Bils (2013) for a more detailed discussion of these measurement issues associated with using the CE Survey to estimate elasticities. Therefore for robustness, we also use the standard approach of instrumenting total expenditure with total income reported by the household. The estimated elasticities yield similar results to our base estimation without instrumenting.

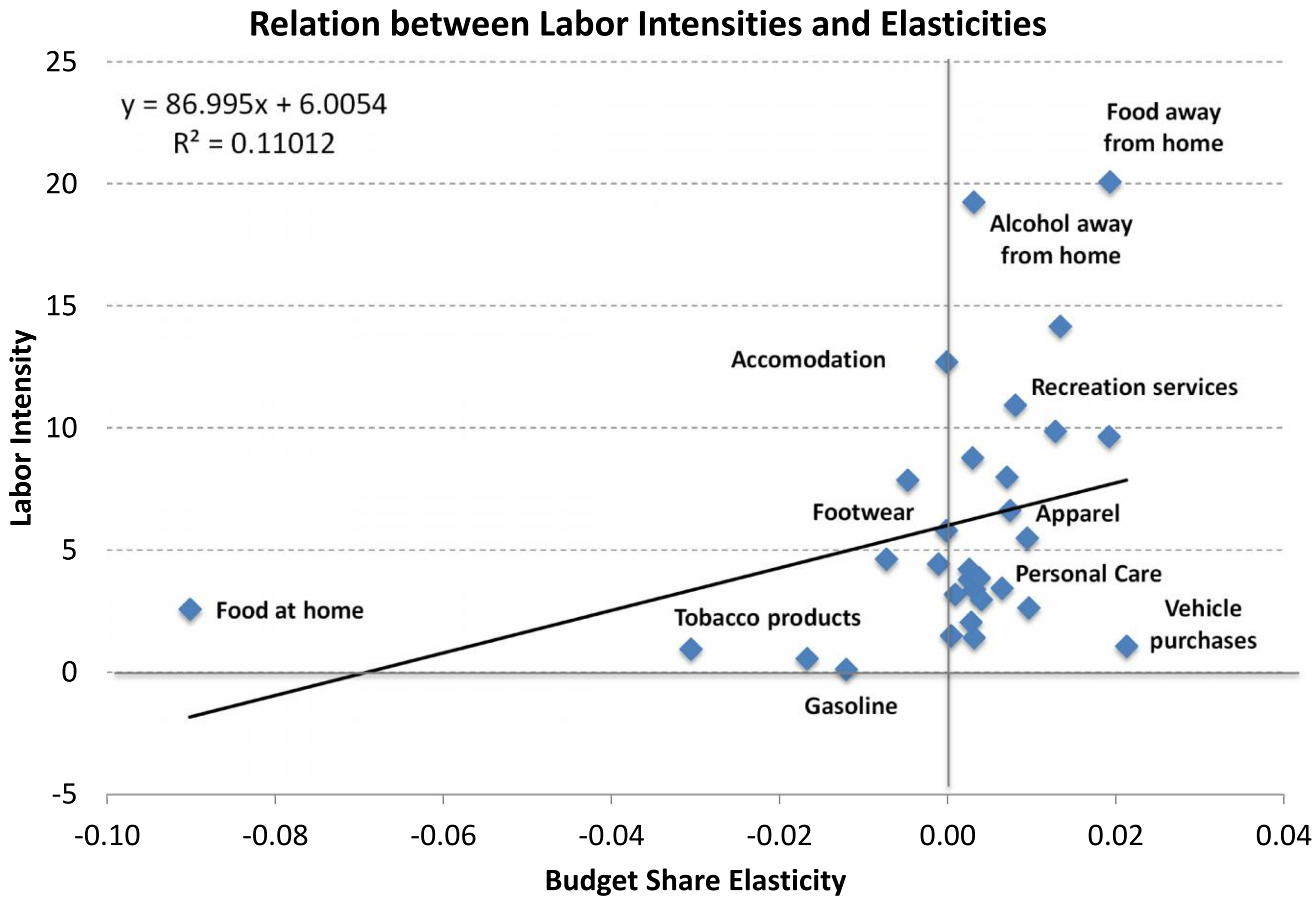

Figure 2

Figure 2 shows there is a low positive relation between a category expenditure elasticity and its labor intensity measure. To examine the effect of across-category substitution on employment, we perform a similar exercise as our trading-down calculation. We compute the change in the number of employed workers between 2007 and 2012 due to changes in the shares of the expenditures categories, holding fixed the measured labor intensity.[30] This exercise yields a negligible effect of the substitution across categories on aggregate employment, in contrast to our findings of large effects of quality trading-down within categories.

D Results based on alternative labor intensities

As a robustness check, Tables 8 and Tables 9 reports our calculation using our second measure of labor intensity, employment/gross margin. For the period 2007-12, the change in employment accounted for by trading down represents 37 percent of the fall in employment. For the period 2007-09, this fraction represents 28 percent of the fall in employment.

Table 8: Market shares and Labor Intensity: 2007

Industry	$m Sales - COGS	Labor intensity			Market share		
		Low	Middle	High	Low	Middle	High
Home furnishing and appliances	248,751	14.05	15.27	27.90	1%	95%	4%
Grocery stores	353,472	16.23	18.06	21.22	16%	80%	4%
Food services and drinking places	270,390	68.08	131.75	93.70	58%	36%	6%
Clothing stores	119,579	27.34	24.34	29.59	8%	76%	16%
General merchandise stores	425,163	16.23	24.26	19.49	56%	25%	19%
Total	1,417,355	26.68	41.65	36.41	33%	58%	10%

Note: This table depicts the 2007 sales less cost of goods sold, labor intensity and market share for different retail sectors. The last row is a value-added-weighted measure. Sales and cost of goods sold are from the US Census of Retail Trade, and the labor intensities and market shares are from Compustat. Labor intensity is defined the ratio the number of employees (thousands) per sales less cost of goods sold (millions), and the market share are the share of sales less cost of goods sold for each price tier within each sector. Price tiers are denoted by low, middle and high, and are based on Yelp! classifications of prices $, $$, and $$$, respectively. See text for more information.

[30] We also did the computation using actual observed change in expenditure allocations in the NIPA PCE data and found similar results.

Table 9: Market shares and Labor Intensity: 2012

Industry	$m Sales - COGS	Labor intensity			Market share		
		Low	Middle	High	Low	Middle	High
Home furnishing and appliances	212,650	8.25	14.68	15.57	1%	93%	5%
Grocery stores	410,500	14.62	16.61	16.33	21%	72%	7%
Food services and drinking places	313,782	53.53	85.08	125.80	70%	26%	4%
Clothing stores	129,444	21.20	21.62	26.68	11%	79%	10%
General merchandise stores	473,441	16.23	24.26	19.49	65%	21%	14%
Total	1,539,817	22.72	33.07	40.37	41%	51%	8%

Note: This table depicts the 2012 sales less cost of goods sold, labor intensity and market share for different retail sectors. The last row is a value-added-weighted measure. Sales and cost of goods sold are from the US Census of Retail Trade, and the labor intensities and market shares are from Compustat. Labor intensity is defined the ratio the number of employees (thousands) per sales less cost of goods sold (millions), and the market share are the share of sales less cost of goods sold for each price tier within each sector. Price tiers are denoted by low, middle and high, and are based on Yelp! classifications of prices $, $$, and $$$, respectively. See text for more information.

E Results based on alternative definition of the recession

As a robustness check, Table 10 and 11 reports our calculation using for the period 2007-09 for the retail sector and manufacturing sectors, respectively. For the period 2007-09, the change in employment accounted for by trading down represents 20 percent and 16 percent of the fall in employment for the retail sector and manufacturing sectors, respectively.

Table 10: Market shares and Labor Intensity: 2009

Industry	$m Sales in 2009	Labor intensity			Market share		
		Low	Middle	High	Low	Middle	High
Grocery stores	558,337	n.a.	4.40	4.83	n.a.	95%	5%
Food services and drinking places	690,225	1.83	4.01	5.91	44%	51%	5%
Clothing stores	621,902	12.34	21.67	21.97	43%	47%	10%
General merchandise stores	255,052	6.69	8.40	9.27	19%	71%	10%
Total	673,729	3.61	5.74	7.24	76%	16%	9%

Note: This table depicts the 2009 total sales, labor intensity and market share for different retail sectors. The last row is a sales-weighted measure. Sales are from the US Census of Retail Trade, and the labor intensities and market shares are from Compustat. Labor intensity is defined as the number of employees per million dollars of sales and the market share are the share of sales for each price tier within each sector. Price tiers are denoted by low, middle and high, and are based on Yelp! classifications of prices $, $$, and $$$, respectively. See text for more information.

Table 11: Market shares and Labor Intensity: 2009

Industry	$m Sales in 2009	Labor Intensity			Market share		
		Low	Middle	High	Low	Middle	High
31	824,269	0.48	3.27	n.a.	34%	66%	n.a.
32	1,201,015	3.10	3.36	4.74	27%	47%	26%
33	1,902,289	1.98	2.96	3.86	38%	57%	5%
Total	3,927,573	2.01	3.15	3.32	33%	54%	13%

Note: This table depicts the 2009 total sales, labor intensity and market share for different manufacturing sectors. The last row is a sales-weighted measure. Sales are from Census, and the labor intensities and market shares are from Compustat. Labor intensity is defined as the number of employees per million dollars of sales and the market share are the share of sales for each price tier within each sector. Price tiers are denoted by low, middle and high, and are based on firm-level producer price data. See text for more information.