

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. 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 Jon Steinsson. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2015 by Nir Jaimovich, Sergio Rebelo, and Arlene Wong. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

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

ABSTRACT

We document two facts. First, during recessions consumers trade 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. So, when households trade down, labor demand falls, increasing the severity of recessions. We find that the trading-down phenomenon accounts for a substantial fraction of the fall in U.S. employment in the recent recession. We study two business cycle models that embed quality choice and find that the presence of quality choice magnifies the response of these economies to real and monetary shocks.

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

Arlene Wong
Department of Economics
Northwestern University
2001 Sheridan Road
Evanston, IL 60208-2600
arlene.wong@u.northwestern.edu

Sergio Rebelo
Northwestern University
Kellogg School of Management
Department of Finance
Leverone Hall
Evanston, IL 60208-2001
and NBER
s-rebelo@northwestern.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).¹

In this paper, we contribute to this literature by documenting two facts. First, during recessions consumers trade down in the quality of the goods and services they consume. Second, the production of low-quality goods is generally less labor intensive than that of high-quality goods. So, when households trade down, labor demand falls, increasing the severity of recessions.

To quantify the implications of “trading down” for employment, 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 also use data from the Census of Retail Trade.

We estimate that 22 to 36 percent of the decline in employment in the 2007-2012 period is accounted for by consumers trading down in the quality of the goods and services they purchased.²

To study the effects of trading down from a theoretical perspective, we embed quality choice into two otherwise standard models: a flexible-price model and a Calvo-style

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

²For data availability reasons, our analysis focuses on the period 2007-2012. But we also provide some evidence of trading down for the 1990-1992 and the 2001 recessions.

sticky-price model. We find that the presence of quality choice magnifies the response of these economies to real and monetary shocks, generating larger booms and deeper recessions. This amplification results from stronger shifts in both labor demand and labor supply. Consider the case of an expansion. In standard business-cycle models, the response of workers to an increase in the real wage is muted by the presence of decreasing marginal utility of consumption. As workers who supply more labor use the additional income to raise their consumption, their marginal utility of consumption declines. The possibility of consuming higher quality goods reduces this fall, resulting in a larger increase in the labor supply. At the same time, the production of higher quality goods requires more labor, generating a larger increase in the labor demand than in a model without quality choice.

The quality-augmented model has two other interesting properties. First, it can generate comovement between employment in the consumption and investment sectors, a property that is generally difficult to obtain (see Christiano and Fitzgerald (1998) for a discussion). Second, the model produces an endogenous, countercyclical labor wedge. As Shimer (2009) discusses, this type of wedge is necessary in order to reconcile business-cycle models with the empirical behavior of hours worked.

Our paper is organized as follows. In Section 2, we describe our data and present our empirical results. The flexible and sticky price models are presented in Sections 3 and 4, respectively. Section 5 concludes.

2 Empirical findings

In this section we present our empirical findings. We start by studying a data set that combines data from Yelp!, the Census of Retail Trade, and Compustat. We complement this analysis using two additional data sets. The first, is data on restaurant expenditures collected by the Bureau of Economic Analysis (BEA) as part of the National Income and Product Accounts (NIPA) and by the BLS as part of the Current Establishment

Survey (CES). The second, is data on traffic and revenue for restaurants by quality segment collected by the NPD Group, a marketing consulting firm. We then extend our analysis to all the sectors in the economy using the micro data gathered by the BLS to construct the PPI.

We focus on the 2007-2012 period because, 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. For robustness, we also report results for the 2007-2009 period.

2.1 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 six North American Industry Classification System (NAICS) sectors: accommodation, apparel, grocery stores, restaurants, home furnishing, and general merchandise. These sectors represent 17 percent of private non-farm employment.

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.³ 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.

³Yelp! 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.

We provide below a brief discussion of the construction of the Yelp! data set. In Appendix A, we discuss this construction in more detail. 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).⁴

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, 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. The primary labor intensity measure we use is the ratio of employees to sales. The choice of this measure is dictated by data availability. Less than 1/4 of the firms included in Compustat data report the share of labor in total cost, which is a natural measure of labor intensity. In the sample of firms that report the labor share in cost, the correlation between labor share and employees/sales is 0.94.

As a robustness check, we also use the ratio of employees to gross margin. Gross margin, which is sales minus cost of goods sold, is a measure that is close to value

⁴The dispersion in price categories across cities is relatively small; it is rare for firms to be included in different price categories in different cities.

added.⁵ The correlation between employees/gross margin and employees/sales is 0.72. The correlation between the labor share in cost and employees/gross margin is 0.97.

Findings

The distribution of firms by price category is as follows. The low-, middle- and high-price categories account for 25, 57, and 18 percent of the firms, respectively. The distribution of sales across price categories is similar. The low-, middle- and high-price categories account for 35, 53, and 12 percent of sales, respectively.⁶

Table 1 documents our first fact: between 2007 and 2012, firms that produce middle- and high-quality items lost market share relative to firms that produce low-quality items. This pattern emerges for the six sectors we consider with one exception: the market share of high-quality grocery stores increased. This exception is driven by an outlier: WholeFoods, a high-quality supermarket, gained market share despite the recession.

Table 2 documents our second fact: both our measures of labor intensity are increasing in quality.⁷ For example, the number of employees per million dollar of sales is 15.1, 9.2, and 6.5, for high-, middle- and low-quality apparel stores, respectively. So, other things equal, a shift of one million dollar of sales from a middle-quality to a low-quality apparel store destroys roughly three jobs.

As Table 3 shows, we obtain broadly similar results using our second measure of labor intensity, employment/gross margin.

⁵Value 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.

⁶Here are some examples of firms categorized into different quality bins using Yelp! data. In the Accommodation sector: Choice (low), IHG, (middle), Starwood (high). In the Home Furnishing sector: Lumber Liquidators (low), Lowe's, (middle), Williams-Sonoma (high). In the Grocery Store sector: Sam's Club (low), Safeway, (middle), WholeFoods (high). In the Restaurant sector: McDonald's (low), Cheesecake Factory, (middle), Del Frisco (high). In the Apparel Stores sector: Ross (low), Gap, (middle), Abercrombie Fitch (high). In the General Merchandise Store sector: Dollar stores (low), Discount, (middle), Non-discount (high).

⁷This table is based on 2012 labor intensity measures. We obtain similar results using 2007 labor intensity measures.

Table 1: Market Share Changes by Quality Segments

Sector	Low	Middle	High
Accommodation	0.5%	0.4%	-0.9%
Home furnishing stores	0.1%	-0.5%	0.4%
Grocery stores	3.9%	-6.4%	2.6%
Restaurants	8.6%	-7.9%	-0.6%
Apparel stores	4.3%	-0.6%	-3.7%
General merchandise stores	8.5%	-4.8%	-3.8%
Total	4.8%	-4.1%	-0.7%

Note: This table reports the changes in market share for each quality tier (low, middle, and high) within each retail sector over 2007 to 2012. Market shares are based on the sales revenue of firms. Quality tiers are based on the U.S. Census of Retail Trade categorization for General merchandise stores, and on the Yelp! price categories for firms in the other sectors. See text for more details.

Table 2: Employees per Million Dollar Sales by Quality Segments

Sector	Low	Middle	High
Accommodation	6.3	8.6	21.6
Home furnishing stores	3.5	4.9	5.9
Grocery stores	1.9	4.2	6.1
Restaurants	13.4	19.5	22.4
Apparel stores	6.5	9.2	15.1
General merchandise stores	3.7	6.9	7.2
Total	5.5	8.5	11.1

Note: This table reports the labor intensity in 2012 for each quality tier (low, middle, and high) for each retail sector. Labor intensity is defined as the number of employees per million dollar of sales. Quality tiers are based on the U.S. Census of Retail Trade categorization for General merchandise stores and the Yelp! price categories for firms in the other sectors. See text for more details.

Table 3: Employees per Million Dollars of Gross Margin, by Quality Segments

Sector	Low	Middle	High
Accommodation	6.3	8.6	21.6
Home furnishing stores	8.3	14.7	15.6
Grocery stores	14.6	16.6	16.3
Restaurants	53.5	85.1	125.8
Apparel stores	21.2	21.6	26.7
General merchandise stores	16.2	24.3	19.5
Total	21.5	31.3	39.0

Note: This table reports the labor intensity in 2012 for each quality tier (low, middle, and high) within each retail sector. Labor intensity is defined as the number of employees per million dollar of (sales-cost of goods sold). Quality tiers are based on the U.S. Census of Retail Trade categorization for General merchandise stores and the Yelp! price categories for firms in the other sectors. See text for more details.

We quantify the effects of trading down on employment by using a simple accounting method. For each sector j , we define the market share of each of the three quality tiers (low, middle or high, denoted by i) in a given year t as:

$$\text{market share}_{t,i,j} = \frac{\text{sales}_{t,i,j}}{\text{total sales}_{t,j}}. \quad (1)$$

For each sector j , we write employment in 2012 using the following identity:

$$N_{2012,j} = \text{total sales}_{2012,j} \sum_i \text{market share}_{2012,i,j} \left(\frac{N_{i,j}}{\text{sales}_{i,j}} \right)_{2012}, \quad (2)$$

where $N_{i,j}$ is the number of workers employed by firm in quality tier i and sector j .

We then compute, N_{2012}^* , the employment that would have resulted in 2012 if the market shares of different firms were the same as in 2007:

$$N_{2012,j}^* = \text{total sales}_{2012,j} \sum_i \text{market share}_{2007,i,j} \left(\frac{N_{i,j}}{\text{sales}_{i,j}} \right)_{2012}. \quad (3)$$

Finally, we compute the change in employment accounted for by changes in market share as:

$$N_{2012,j} - N_{2012,j}^* = \text{total sales}_{2012,j} \sum_i (\text{market share}_{2012,i,j} - \text{market share}_{2007,i,j}) \left(\frac{N_{i,j}}{\text{sales}_{i,j}} \right)_{2012}. \quad (4)$$

Table 4 reports the results of this accounting exercise.⁸ Between 2007 and 2012, high-quality producers lost 1 percent in market share, middle-quality producers lost 4 percent, and low-quality producers gained 5 percent. Overall employment in the sectors included in our data fell by 7.9 percent between 2007 and 2012. The change in employment accounted for by trading down is -2.6 percent, which represents 33 percent of the fall in employment. When we consider the period 2007-2009, we find

⁸We compute the change in employment accounted for by trading down as the sum of two components: the direct effect of quality shifts on employment within the same sector and the indirect effect on employment in other sectors that provide intermediate inputs to the sector that experienced a change in market share. These indirect effects are quite small, so our results are similar whether we include them or not. We discuss how we compute these indirect effects in Appendix C.

that employment in the sectors included in our data declined by 4.5 percent. The change in employment accounted for by trading down is -1.4 percent, which represents 31 percent of the fall in employment.⁹

Table 4: Employment Effect of Quality Trade-Down

Quality	Δ Market share	Emp/\$m sales	Implied Δ Emp	Actual Δ Emp
Low	5 ppt	5.6	978,000	2,105,000
Middle	-4 ppt	9.1	-1,351,000	-3,635,000
High	-1 ppt	11.1	-210,000	-214,000
Total Change			-582,000	-1,744,000
Total Percentage Change			-2.6%	-7.9%
Share accounted for by trading down:				33%

Note: This table reports the employment effect of quality trade-down for the sectors that we examined. Column 2 reports the change in market share over 2007-2012. Column 3 reports the labor intensity (number of employees per million dollars of sales) in 2012. Column 4 reports the change in employment implied by the shift in market share. This change is computed as the product of columns 2-3 times the sales in 2012. Column 5 reports the actual change in employment.

⁹Using employment/sales and total sales measured in 2007 instead of in 2012 in equation (4) we obtain very similar results: 32 percent of the change in employment is accounted for by trading down.

As a robustness check, Table 5 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 42 percent of the fall in employment. For the period 2007-09, this fraction represents 32 percent of the fall in employment.

Table 5: Employment Effect of Quality Trade Down

Quality (q)	Δ Market share	Labor intensity	Implied Δ Emp	Actual Δ Emp
Low	6ppt	21.5	1,266,854	2,105,000
Middle	-5ppt	31.3	-1,539,960	-3,635,000
High	-1ppt	39.0	-453,121	-214,000
Total Change			-726,227	-1,744,000
Total Percentage Change			-3.3%	-7.9%
Share accounted for by trading down:				42%

Note: This table reports the employment effect of quality trade-down for the sectors that we examined. Column 2 reports the change in market share over 2007-2012 for each quality tier. Column 3 reports the labor intensity (the number of employees per million dollars of gross margin) in 2012. Column 4 reports the change in employment implied by the shift in market share in column 2. This change is computed as the product of columns 2 and 3 times the value of gross margin in 2012 for each quality tier. Column 5 reports the actual change in employment.

Table 6 provides another way of seeing the effects of trading down on employment. The table shows that during the recession, low-quality producers generally expanded employment while middle- and high-quality producers contracted employment.¹⁰

Table 6: Employment Changes by Sector and Quality Segment

Sector	Low	Middle	High	Total
Accommodation	-23	-7	-118	-149
Home furnishing stores	7	-947	-74	-1,014
Grocery stores	99	-291	40	-152
Restaurants	1,613	-1,882	101	-167
Apparel stores	1	-231	-92	-322
General merchandise stores	408	-276	-72	61
Total	2,105	-3,635	-214	-1,744

Note: This table reports the actual change in employment over 2007-2012 by quality tier within each sector. The employment numbers for the General Merchandise sector are from the U.S. Census of Retail Trade. For all of the other sectors, the employment numbers are computed using the actual firm-level employment data from Compustat.

¹⁰There are some exceptions to this pattern. Employment increased in high-quality supermarkets, reflecting the expansion of Wholefoods. Employment also expanded in high-quality restaurants as a result of trading up by some consumers.

2.2 The NIPA and NPD data sets

To complement our analysis we study two additional data sets: NIPA data for restaurant expenditures and data from the NPD Group, a marketing consulting firm, on the evolution of market shares in restaurants of different quality levels.

The BEA disaggregates NIPA quarterly measures of real expenditures on “Food away from home” into two categories: expenditures on limited-service (lower quality) restaurants and on full-service (higher quality) restaurants.

While overall expenditures on food away from home fell by 7.4 percent between 2007 and 2012, real expenditure on limited-service restaurants fell by only 4.7 percent. In contrast, real expenditure on full-service restaurants fell by 10.2 percent. Overall, limited-service restaurants gained 1.5 percentage points of market share between 2007 and 2009 and 1.3 percentage points between 2007 and 2012. These shifts in expenditure provide clear evidence of trading down.

The Current Establishment Statistics conducted by the BLS reports employment data for limited-service and full-service restaurants. Using these data, we find that the employment per sales in 2012 is 21.4 for full-service restaurants and 15.4 for limited-service restaurants.

The change in employment implied by changes in market share is -0.5 percent for 2007-2009, and -0.4 percent over 2007-2012. The effect of trading down in the 2007-2012 (2007-2009) period represents 36 percent (17 percent) of the total change in employment in the food away from home sector.¹¹

Our second source of restaurant data is the NPD Group. This data set has 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.

¹¹This change was -2.7 percent for the period 2007-2009 and -1.2 percent for the period 2007-2012.

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. 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.¹²

We find clear evidence of trading down in the NPD data. Consider first the number of meals served. Table 7 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.¹³ Table 8 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.¹⁴

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.

¹²These price data were collected in March 2013. We do not have average meal prices for fine-dining restaurants.

¹³There 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.

¹⁴Tables 7 and 8 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 7: 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.

2.3 The PPI data

In order to extend the analysis to other sectors of the economy, we use the confidential micro data collected by the Bureau of Labor Statistics (BLS) to construct the PPI.¹⁵ The PPI data set measures producers' prices for manufacturing, services, and all the other sectors of the economy.

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 22 percent of private non-farm employment. We refer the reader to Appendix B for more details on the construction of this data set.

¹⁵Examples of other papers that use these data include Nakamura and Steinsson (2008), Gilchrist et al (2014), Gorodnichenko and Weber (2014), and Weber (2015).

Table 8: 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.

We focus on the 2-digit NAICS manufacturing sectors 31, 32, and 33, and the retail trade sector 44, 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 extend the analysis on the PPI back to prior recessions, and so we focus our analysis on the recent recession.

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} :¹⁶

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

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}):¹⁷

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

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 15 percent to the high-quality category, the middle 55 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.¹⁸ This merged data set allows us to compute labor intensity

¹⁶Our analysis is based on products defined at a six-digit level. The variable \bar{p}_{kt} is a shipment-value weighted average within the six-digit level. For reporting purposes, we aggregate the results to the two-digit level. The aggregation is based on shipment revenue.

¹⁷This approach for constructing firm-level price indices is similar to that used by Gorodnichenko and Weber (2014), and Gilchrist et al (2014). We refer the reader to Section II in Gorodnichenko and Weber (2014) for a discussion of how the BLS samples products and firms.

¹⁸The 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.¹⁹

Tables 9-11 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.

Table 9: Market Share Changes by Quality

NAICS Sector	Low	Middle	High
31 Manufacturing: Food, textiles, etc.	11.3%	-11.3%	n.a.
32 Manufacturing: Wood, chemical, etc.	0.9%	1.9%	-2.8%
33 Manufacturing: Computers, equip., etc.	6.8%	-6.5%	-0.3%
44 Retail trade	3.2%	-2.8%	-0.4%
Total	4.5%	-3.6%	-0.9%

Note: This table reports the changes in market share for each quality tier (low, middle, and high) within each sector over 2007 to 2012. Market shares are based on sales revenue. Quality tiers are determined using the PPI dataset.

¹⁹We use the entire sample of establishments within the PPI to rank the establishments, not just those that we are able to match with Compustat.

Table 10: Employment per Million Dollars of Sales

NAICS Sector	Low	Middle	High
31 Manufacturing: Food, textiles, etc.	0.4	3.4	n.a.
32 Manufacturing: Wood, chemical, etc.	2.7	2.9	4.6
33 Manufacturing: Computers, equip., etc.	1.4	2.4	3.3
44 Retail trade	2.9	5.0	12.8
Total	2.2	3.5	6.6

Note: This table reports the labor intensity for each quality tier (low, middle, and high) within each sector in 2012. Labor intensity is defined as the number of employees per million dollars of sales. Quality tiers are determined using the PPI dataset.

Table 11: Employment per Million Dollars of Gross Margin

NAICS Sector	Low	Middle	High
31 Manufacturing: Food, textiles, etc.	5.8	6.9	n.a.
32 Manufacturing: Wood, chemical, etc.	5.7	8.4	14.0
33 Manufacturing: Computers, equip., etc.	4.3	10.6	16.3
44 Retail trade	12.8	14.6	21.0
Total	7.7	11.1	15.6

Note: This table reports the labor intensity for each quality tier (low, middle, and high) within each sector in 2012. Labor intensity is defined as the number of employees per million dollars of gross margin. Quality tiers are determined using the PPI dataset.

We report in Table 12 the results obtained by performing our accounting exercise with PPI data. We find that trading down accounts for 22 percent of the jobs lost between 2007 and 2012 and 16 percent of the the jobs lost between 2007 and 2009.²⁰

In sum, our results using the PPI data are broadly similar to those obtained with Yelp! and Census of Retail Trade data. Higher quality goods, which are generally more labor intensive, lost market share during the recent recession. This loss of market share accounts for about a quarter of the overall decline in employment in the sectors covered by PPI data.

²⁰When we use employment/gross margin as our measure of labor intensity, we find that trading down accounts for 23 percent of the jobs lost between 2007 and 2012 and 8 percent of the the jobs lost between 2007 and 2009.

Table 12: Employment Effect of Quality Trade Down

Quality	Δ Market share	Labor intensity	Implied Δ Emp
Low	5ppt	2.2	679,998
Middle	-4ppt	3.5	-1,216,488
High	-1ppt	6.6	-329,477
Total Change			-865,968
Total Percentage Change			-3.6%
Actual Total Change			-2,909,607
Actual Percentage Change			-16.2%
Share accounted for by trading down:			22%

Note: This table reports the employment effect of quality trade-down. Column 2 reports the change in market share over 2007-2012 for each quality tier. Column 3 reports the labor intensity (the number of employees per million dollars of sales) in 2012. Column 4 (first segment of the table) reports the change in employment implied by the shift in market share in column 2. This change is the product of columns 2 and 3 times the value of sales in 2012 for each quality tier. See text for more details.

2.4 Extending the analysis to the rest of the economy

So far, we focused on sectors for which we have price measures. In order to extend the analysis beyond these sectors we proceed by using labor intensity, instead of price, as a measure of quality. The rationale for this procedure is that labor intensity is positively correlated with quality. We use the same equations as before (equations (3) and (4)) but the subscript i refers to an individual firm instead of to a quality tier.

Within each 3-digit NAICS sector, we compute the market share for each firm within the sector in 2007. We then compute a counter-factual employment estimate for each firm by multiplying the firm's 2012 employment numbers with their 2007 market share. Summing up over all firms within the sector gives us the counter-factual employment that would have occurred if there had been no changes in market shares within the sector. We then sum over all sectors in the economy to get an aggregate counter-factual employment number, which we can compare to actual employment.

Doing our accounting exercise using employment/sales as our measure of labor intensity, we find that trading down accounts for 28 percent of the jobs lost between 2007 and 2012 and 23 percent of the the jobs lost between 2007 and 2009.²¹

2.5 Summary

Table 13 provides a summary of our results. This table suggests that employment effects of trading down are quantitatively large. During the 2007-2012 period, trading down accounts for 22 to 36 percent of the jobs, depending on the measure of quality used and the data set.

²¹Our results are robust to excluding sectors in which there is not much scope for the consumer to trade down: Utilities, Warehousing and Storage, Waste Management and Remediation Services, and Water Transportation.

Table 13: Summary of Employment Effect of Quality Trade Down

Data set	Quality measure	2007-12	2007-09
Yelp! and Census of Retail Trade	Price	33	31
PPI	Price	22	16
Compustat (all sectors)	Labor intensity	28	23
NIPA	Store category (Limited versus full service)	36	17

Note: This table summarizes the estimated effect of trading-down on employment over 2007-2012 and 2007-2009, across the different data sets and quality measures.

2.6 Analyzing other recessions

A natural question is whether the trading down that occurred in the 2007-2012 period also occurred in other recessions. In particular, it is important to know whether trading down occurs only in large recessions. Unfortunately, we cannot use the PPI data set to study prior recessions. Even though the data goes back to 1997, there is limited information about the units in which prices are measured prior to 2007. However we can use data from Yelp!, the Census of Retail Trade and NIPA to shed light on other recessions.

Using Yelp! and Census of Retail Trade data to study the 2000 recession

Yelp! did not exist before 2004, so we cannot use it to classify all firms within the Compustat sample into quality categories. We can, however, do the analysis with firms that are included in our Yelp! data set and existed in prior recessions by assuming that firms remained in the same price/quality category over time. We do this analysis for the 2001 recession.²²

The three sectors for which the number of firms from the Yelp! data set that exists in 2000 is relatively large are: grocery stores, apparel stores, and general merchandise stores. We find that trading down accounts for 53 percent of the decline in employment in the 2000-01 period. In comparison, in the 2007-2012 period trading down accounts for 62 percent of the decline in employment experienced by these three sectors.

Using NIPA and CES data to study the 2001 and the 1990-1991 recession

The CES data on employment in limited- and full-service restaurants starts in the 1st quarter of 1990. We can combine these data with the BEA data on expenditure on limited- and full-service restaurants to analyze the 2001 recession and the 1990-1991

²²We do not analyze the 1990-92 recession for two reasons. First, the number of firms in the Yelp! data that existed in the 1990-92 recession is relatively small. Second, the assumption that the price/quality of the firms in our sample remains constant over time is less tenable.

recession. In both cases, we measure the labor intensity of limited- and full-service restaurants in the quarter in which the recession ends.

Between the beginning of the recession in the 1st quarter of 2001 and the end of the recession in the 4th quarter of 2001, limited-service restaurants gained 0.4 percentage points of market share. The change in employment implied by this change in market share is -0.1 percent, which represents 8.9 percent of the total change (-1.2 percent) in employment in the food away from home sector. The relatively small fraction of the decline in employment accounted for by trading down is not surprising, given that the 2001 recession was short and shallow. In addition, the restaurant sector is less cyclical than the average sector in our sample.

Between the beginning of the recession in the 3rd quarter of 1990 and the end of the recession in the 1st quarter of 1991, limited-service restaurants gained 0.5 percentage points of market share. The change in employment implied by this change in market share is -0.2 percent. This change represents 19 percent of the total change (-0.8 percent) in employment in the food away from home sector.

3 A flexible-price model

In this section we interpret our results using a flexible-price model where households choose the quality of the goods they consume. In this model, households face a natural trade-off: consuming higher quality goods yields higher utility but these higher quality goods are more expensive.

We first consider a static version of the model. This simplified version allows us to highlight the key mechanisms in the model and to derive some analytical results. We and then turn to a dynamic, stochastic version of the model and assess its predictions vis-a-vis the U.S. data.

3.1 A static model

The consumption good can be produced with different levels of quality, q . The model has a representative household so, in equilibrium, all the households choose the same level of quality and only this quality is produced.

Households

Households derive utility from both the quantity, C , and the quality, q , of consumption and disutility from work, N :

$$U = U(C, q, N). \quad (5)$$

The household's budget constraint is:

$$P(q)C = wN + RK, \quad (6)$$

where $P(q)$ is the price of one unit of consumption of quality q , w is the wage rate, and R is the rental rate of capital. The stock of capital, K , is constant.

As usual, we assume that

$$U_1(C, q, N) > 0, \quad (7)$$

$$U_3(C, q, N) < 0. \quad (8)$$

We assume that the marginal utility of quality is positive:

$$U_2(C, q, N) > 0. \quad (9)$$

The first-order condition for the household problem can be written as:

$$\frac{U_2(C, q, N)}{U_1(C, q, N)} = \frac{P'(q)C}{P(q)}, \quad (10)$$

$$\frac{U_3(C, q, N)}{U_1(C, q, N)} = -\frac{w}{P(q)}. \quad (11)$$

We want quality to be 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. Equations (10) and (11) imply that if U is homogeneous in C , quality is independent of income. So, in order for quality to be a normal good, U must be non-homothetic in C .

With this requirement in mind, we assume that the utility function takes the form:²³

$$U = \frac{q^{1-\theta}}{1-\theta} \frac{C^{1-\sigma} - 1}{1-\sigma} - \phi \frac{N^{1+\nu}}{1+\nu}. \quad (12)$$

We show in appendix D that this utility function implies that quality is a normal good.

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

With this utility function, the first order conditions (10) and (11) can be rewritten as:

$$\frac{1 - C^{\sigma-1}}{1 - \sigma} = \frac{1}{1 - \theta} q \frac{p'(q)}{p(q)}, \quad (13)$$

$$\phi N^\nu = \frac{q^{1-\theta}}{1 - \theta} C^{-\sigma} \frac{w}{P(q)}. \quad (14)$$

Production

We assume that producers are perfectly competitive. To produce C units of a consumption good with quality q , they combine labor and capital according to the following CES production function:

$$C = A \left[\alpha \left(\frac{N}{q} \right)^\rho + (1 - \alpha) (K)^\rho \right]^{\frac{1}{\rho}}, \quad (15)$$

where A denotes the level of total factor productivity.

²³In order for utility to increase with quality ($U_q > 0$) we need: $(C^{1-\sigma} - 1)/(1 - \sigma) > 0$. For this condition to hold, it is sufficient that $C > 1$. We assume that income is high enough that this condition always holds. In our stochastic simulations we verify that this condition holds at each point in time.

The producer's problem is:

$$\max P(q)C - wN - rK. \quad (16)$$

We assume that $\rho < 0$, so there is less substitution between capital and labor than in a Cobb-Douglas production function. This assumption is necessary so that, as in the data, higher quality goods are more labor intensive (see the Appendix B for details).

The price schedule, $P(q)$, implied by the firm's first-order condition is:

$$P(q) = \frac{1}{A} \left[\alpha^{\frac{1}{1-\rho}} (qw)^{\frac{\rho}{\rho-1}} + (1-\alpha)^{\frac{1}{1-\rho}} r^{\frac{\rho}{\rho-1}} \right]^{\frac{\rho-1}{\rho}}.$$

The elasticity of $P(q)$ is:

$$\frac{P'(q)q}{P(q)} = \frac{1}{1 + \left(\frac{1-\alpha}{\alpha}\right)^{\frac{1}{1-\rho}} \left(\frac{R}{qw}\right)^{\frac{\rho}{\rho-1}}}.$$

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

Comparative statics

We now consider the effect of an increase in total factor productivity, A . Consider first a version of the model where quality is fixed, so that $q = 1$ and the price of the consumer good is normalized to one. For simplicity and consistent with our calibration below, we assume that $\sigma = 1$ (i.e. utility is logarithmic in consumption). The first-order conditions for the household problem imply:

$$\phi N^\nu = \frac{w}{C}. \quad (17)$$

It is easy to see that since both w and C are proportional to A , changes in A have no effect on the labor supply. The income and substitution effects of changes in wages are exactly offsetting, so N is constant.

In contrast, in the model with quality choice an increase (decrease) in A leads to an increase (decrease) in N .²⁴ The intuition for this result is as follows. When A rises, the consumer has more income and so he consumes more. Equation (13), together with the fact that $P'(q) > 0$, implies that it is optimal for households to consume both more quantity and higher quality. The rise in quality shifts up the marginal utility of consumption schedule, which is given by $q^{1-\theta}C_t^{-\sigma}/(1-\theta)$, leading to a rise in N (see equation (14)). Moreover, the increase in the demand for labor quality, increases the demand for labor since the production function implies that labor intensity increases with quality.

In summary, when A rises both the labor supply and the labor demand expand. Households are willing to work more because they can buy higher quality goods. Firms demand more labor in order to produce higher quality goods.

3.2 A dynamic model

We now consider a dynamic, stochastic version of the model. Since our evidence about trading down is only for consumption goods, we assume that the investment good is produced with a fixed level of quality.

The household's problem is:

$$\max U = E_0 \sum_{t=0}^{\infty} \beta^t \left[\frac{q_t^{1-\theta}}{1-\theta} \frac{C_t^{1-\sigma} - 1}{1-\sigma} - \phi \frac{N_t^{1+\nu}}{1+\nu} \right], \quad (18)$$

s.t.

$$P(q_t)C_t + P_t^I I_t = w_t N_t + r_t K_t, \quad (19)$$

$$K_{t+1} = I_t + (1 - \delta)K_t, \quad (20)$$

where I_t and P_t^I denote the level and the price of investment, respectively. E_0 is the conditional expectation operator.

²⁴The proof, which involves tedious algebra, is available from the authors upon request.

The first-order conditions for the household are two equations associated with the static model (equations (13) and (14)), together with the following additional condition:

$$\lambda_t = E_t \beta \lambda_{t+1} (1 + r_{t+1}). \quad (21)$$

Here, λ_t is the Lagrange multiplier associated with the household's budget constraint (equation (19)).

Consumption is produced by a continuum with unit measure of competitive firms according to the following production function:

$$C_t = A_t \left[\alpha \left(\frac{N_t^C}{q_t} \right)^\rho + (1 - \alpha) (K_t^C)^\rho \right]^{\frac{1}{\rho}}, \quad (22)$$

where N_t^C and K_t^C denote labor and capital employed in the consumption sector, respectively.²⁵ The consumption firms' problem is:

$$\max P(q_t) C_t - w_t N_t^C - R_t K_t^C. \quad (23)$$

Investment is produced by a continuum of measure one of competitive firms according to:

$$I_t = A_t \left[\alpha (N_t^I)^\rho + (1 - \alpha) (K_t^I)^\rho \right]^{\frac{1}{\rho}}, \quad (24)$$

where N_t^I and K_t^I denote labor and capital employed in the investment sector, respectively. This production function is that same used in the consumption sector but there is no choice of quality.

The investment firms' problem is:

$$\max P_t^I I_t - w_t N_t^I - R_t K_t^I. \quad (25)$$

The equilibrium conditions for capital and labor are:

$$K_t^C + K_t^I = K_t, \quad (26)$$

$$N_t^C + N_t^I = N_t. \quad (27)$$

²⁵We abstract from technical progress in both the consumption and investment sectors. See Appendix F for a version of the model with labor-augmenting technical progress that is consistent with balanced growth.

We choose the investment good as numeraire ($P_t^I = 1$). Real output (Y_t) in the economy is given by:

$$Y_t = P(q_t)C_t + I_t. \quad (28)$$

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.

We solve the model numerically, using the parameters described in Table 14, by linearizing the equilibrium conditions. We choose a high elasticity of labor supply so that the performance of the version of the model without quality choice is as good as possible. To simplify, we consider the case of $\sigma = 1$, so momentary utility is logarithmic in consumption.

The only new parameter in Table 14 is θ . To choose the value of θ , we compute the change in employment accounted for by changes in quality. We follow a procedure similar to the one we used in our empirical work. We define N_t^* as the employment that would occur if the quality and labor intensity are constant and equal to their steady-state values:

$$N_t^* = P(q_t)C_t \frac{N}{P(q)C},$$

where variables without a time subscript denote steady state values. The fraction of the change in labor accounted for by changes in quality is:

$$\psi_t = \frac{N_t^* - N}{N_t - N}.$$

We choose θ so that the average value of ψ_t computed using the first 25 quarters of the impulse response to a shock to A_t is 25 percent, a value consistent with our empirical results.

Figure 1 shows the impulse response functions of labor and output for two versions of the model: with and without quality choice. We see that the model with quality choice produces much more amplification for the reasons discussed in Section 3.1.

Table 14: Calibration

Parameter	Moment/Description	Value
β	Discount rate	0.985
ν	Inverse of Frisch elasticity	0.001
ϕ	Match N^{SS}	5.31
θ	Elasticity of utility to quality	0.5
δ	Depreciation rate	0.025
α	Production function share	0.5
ρ	EOS between K and N: $\frac{1}{1-\rho}$	-1
κ	AR(1) coefficient of TFP	0.95

Figure 1: Impulse Response Functions for a Rise in A

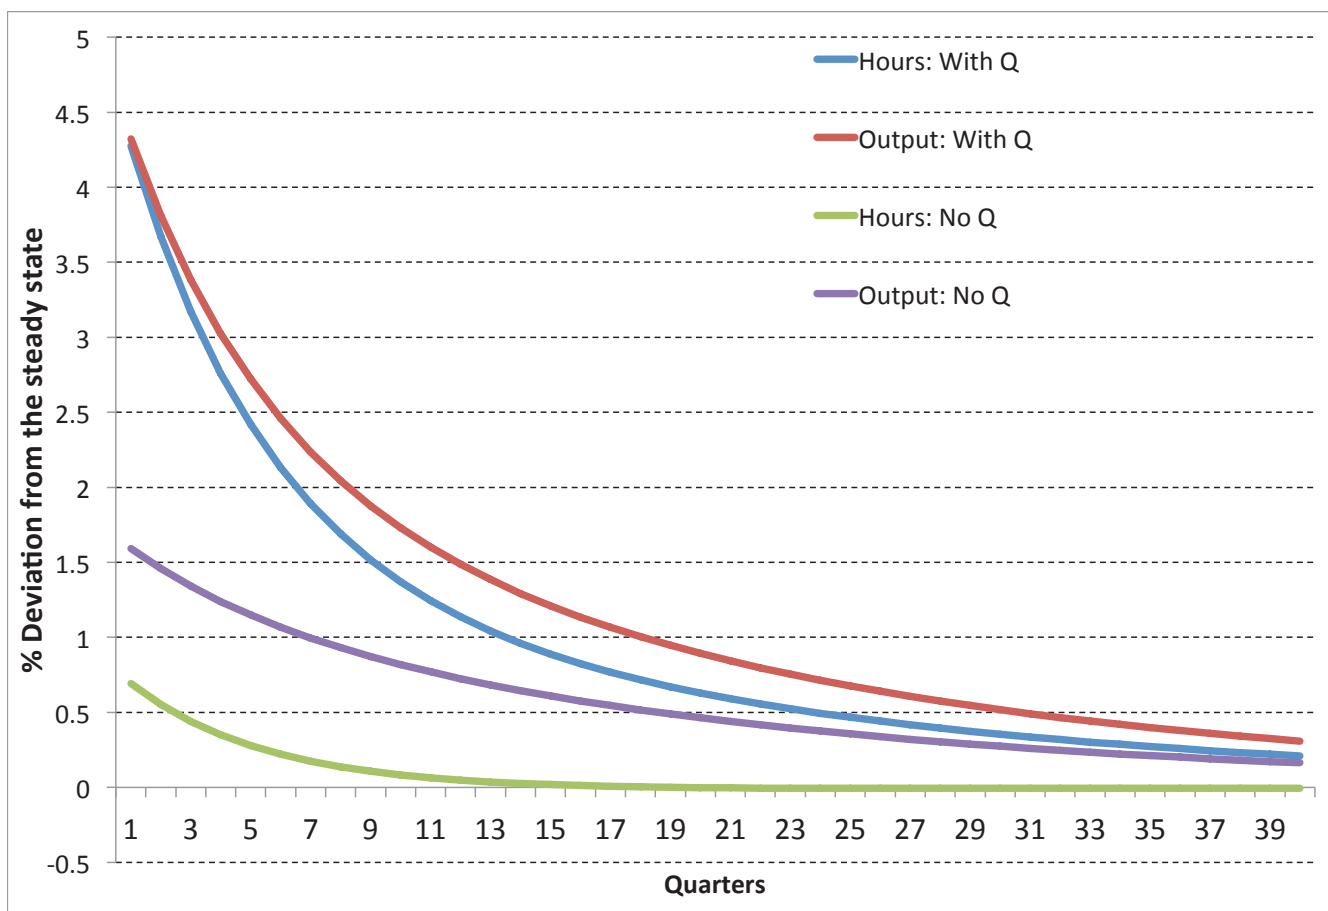


Table 15 compares the cyclical properties of quarterly U.S. data with two models driven by shocks to A_t , with and without quality choice. This table reports the variance and correlation with output for five variables: consumption, investment, total hours worked, as well as hours in the consumption and investment sectors. We see that the model with quality choice provides much more amplification of shocks to A_t than the standard model. In fact, the model generates a relative variation of hours and output that is very close to the one observed in the U.S. data.

Table 15: Second Moments

Variable	Data		Model: Quality		Model: No Quality	
	$\frac{\sigma_X}{\sigma_{GDP}}$	$Cor^{X,GDP}$	$\frac{\sigma_X}{\sigma_{GDP}}$	$Cor^{X,GDP}$	$\frac{\sigma_X}{\sigma_{GDP}}$	$Cor^{X,GDP}$
Total Hours	1.1	0.78	0.98	1.00	0.42	0.95
Hours in C	0.80	0.48	0.37	0.66	0.09	-0.45
Hours in I	2.48	0.86	3.29	0.95	5.42	0.90
Consumption	0.80	0.85	0.49	0.86	0.61	0.94
Investment	3.16	0.87	3.47	0.96	5.92	0.93
Labor Wedge	1.1	-0.69	0.37	-0.86	NA	NA

Comovement

Table 15 shows another interesting difference between the two models. As emphasized by Christiano and Fitzgerald (1998), hours worked in the consumption sector are procyclical in the data but countercyclical in the standard model.²⁶ The model with quality choice generates procyclical hours worked in the consumption sector.

To understand this property, it is useful to write the first-order condition for labor

²⁶To classify the sectors into “consumption” and “investment” we follow standard practice. We use the BEA’s 2002 benchmark I/O “use tables”. to compute the share of sectoral output that is used for private consumption vs. private investment. We assign a sector to the consumption (investment) sector if most of its final output is used for consumption (investment). For the hours/sectors data we use the Current Employment Statistics 1964:Q1 - 2012:Q4.

choice for the standard model assuming that the production function is Cobb-Douglas:

$$\phi (N_t^C + N_t^I)^\nu = \frac{\alpha}{N_t^C}.$$

It is clear that N_t^I and N_t^C are negatively correlated, so that N_t^I and N_t^C cannot be both positively correlated with aggregate output. Using a CES production function changes the form of the first-order condition but does not help generate comovement.

Consider the first-order condition for labor choice in the model with quality choice:

$$\phi (N_t^C + N_t^I)^\nu = \frac{q_t^{1-\theta-\rho}}{1-\theta} \frac{\alpha}{(N_t^C)^{1-\rho}} \left(\frac{A_t}{C_t} \right)^\rho.$$

This equation shows that N_t^C and N_t^I can be positively correlated because, when these variables rise, quality increases. The intuition for why comovement can occur is that the rise in quality increases the demand for labor in the consumption sector contributing to comovement between N_t^C and N_t^I .

An endogenous labor wedge

Shimer (2009) modifies the standard Euler equation for labor to allow for a “labor wedge,” τ_t , that acts like a tax on the labor supply:

$$\phi N_t^\nu = (1 - \tau_t) \frac{1}{C_t} w_t. \tag{29}$$

Shimer computes the labor wedge, using empirical measures of N_t , C_t , and w_t . He finds that τ_t is volatile and counter-cyclical: workers behave as if they face higher taxes on labor income in recessions than in expansions. Comparing equations (30) and the resulting experience in our model

$$\phi N_t^\nu = \frac{q_t^{1-\theta}}{1-\theta} \frac{1}{C_t} \frac{w_t}{P(q_t)}. \tag{30}$$

Since quality choice is procyclical the model can generate an endogenous counter-cyclical labor wedge. Table 15 shows that, for our benchmark parameter values, the model does indeed generate a counter-cyclical labor wedge.

Summary

To summarize, we find that the introduction of a quality choice into an otherwise standard model amplifies the response to real shocks, giving rise to higher fluctuations in hours worked. This property enables the model to match the overall relative variability of hours to output that is observed in U.S. data. Moreover, the model can also account for the sectoral comovement in hours worked.

4 A sticky-price model

In this section, we show that the same mechanism that amplifies real shocks also amplifies nominal shocks. We do so by embedding quality choice in an model with Calvo (1983)-style sticky prices. To highlight the role of quality choice in a parsimonious way, we abstract from capital accumulation.

4.1 The household problem

The representative household maximizes expected life-time utility defined in equation (18). The two constraints on the household problem are:

$$P(q_t)C_t + B_{t+1} = B_t(1 + R_t) + w_tN_t, \quad (31)$$

and

$$E_0 \lim_{t \rightarrow \infty} B_{t+1} / [(1 + r_0)(1 + r_1) \dots (1 + r_t)] \geq 0.$$

Here, B_{t+1} the number of one-period nominal bonds purchased at time t , and R_t is the one period nominal interest rate.

The first-order conditions for the household are two equations associated with the static model (Equations (13) and (14)), together with the following additional condition:

$$\lambda_t = E_t \beta \lambda_{t+1} (1 + R_{t+1}), \quad (32)$$

where λ_t is the Lagrange multiplier associated with the budget constraint (31).

Final good firms

The final good is produced by competitive firms using a continuum of intermediate goods, $Y_t^i(q_t)$:

$$Y(q_t) = \left(\int_0^1 [Y_t^i(q_t)]^{\frac{\varepsilon-1}{\varepsilon}} di \right)^{\frac{\varepsilon}{\varepsilon-1}}, \quad \varepsilon > 1. \quad (33)$$

We assume that producing a final good of quality q_t requires that all intermediate inputs have quality q_t .

The problem of firms in the final-goods sector is:

$$\max P(q_t)Y(q_t) - \int_0^1 P_t^i(q_t) Y_t^i(q_t) di,$$

where $P_t^i(q_t)$ is the price of intermediate good i . The first-order conditions of the firms' problem imply:

$$P_t^i(q_t) = P(q_t) \left[\frac{Y(q_t)}{Y_t^i(q_t)} \right]^{\frac{1}{\varepsilon}}, \quad (34)$$

where P_t is the price of the homogeneous final good. Using the first-order conditions of the firms' problem we can express this price as:

$$P(q_t) = \left(\int_0^1 P_t^i(q_t)^{1-\varepsilon} di \right)^{\frac{1}{1-\varepsilon}}.$$

Intermediate Good Firms

The i th intermediate good is produced by a monopolist using a technology that is the limiting case of the flexible price model without capital:

$$Y_t^i(q_t) = \frac{A_t}{q_t} N_t^i(q_t). \quad (35)$$

Here, $N_t^i(q_t)$ denotes the labor employed by the i^{th} monopolist who is producing a product of quality q . If prices were flexible, the optimal price for the i th monopolist

would be given by the usual mark-up formula:

$$P_t^i(q_t) = \frac{\varepsilon}{\varepsilon - 1} \frac{w_t}{A_t} q_t.$$

However, producers are subject to Calvo-style pricing frictions. We assume that monopolists post a pricing schedule that is linear in q_t :

$$P_t^i(q_t) = \mu_t^i q_t.$$

The monopolist can re-optimize the slope of the pricing schedule, μ_t^i , with probability $1 - \xi$. With probability ξ , the firm has to post the same price schedule as in the previous period:

$$P_t^i(q_t) = \mu_{t-1}^i q_t.$$

We denote by $\tilde{\mu}_t^i$ the optimal price-quality schedule for firms that have the opportunity to re-optimize μ_t^i at time t . Since only a fraction $1 - \xi$ of the firms have this opportunity, the aggregate price level is given by:

$$P(q_t) = \mu_t q_t, \tag{36}$$

where

$$\mu_t = \left[(1 - \xi) (\tilde{\mu}_t^i)^{1-\varepsilon} + \xi \mu_{t-1}^{1-\varepsilon} \right]^{\frac{1}{1-\varepsilon}}. \tag{37}$$

Firm i chooses $\tilde{\mu}_t^i$ to maximize its discounted profits, given by:

$$E_t \sum_{j=0}^{\infty} \beta^j \lambda_{t+j} \left[P_t^i(q_{t+j}) Y_t^i(q_{t+j}) - w_{t+j} N_t^i(q_{t+j}) \right], \tag{38}$$

subject to the Calvo price-setting friction, the production function, and the demand function for $Y_t^i(q_t)$.

Given that the price schedule is linear in quality, the demand function for $Y_t^i(q_t)$ can be written as:

$$\tilde{\mu}_t^i = \frac{P(q_t)}{q_t} \left[\frac{Y(q_t)}{Y^i(q_t)} \right]^{\frac{1}{\varepsilon}}. \tag{39}$$

Since the price schedule chosen in period t is only relevant along paths in which the firm cannot reoptimize its schedule, the firm's problem is given by:

$$E_t \sum_{j=0}^{\infty} \beta^j \xi^j \lambda_{t+j} q_{t+j}^{1-\varepsilon} [P(q_{t+j})]^\varepsilon Y(q_{t+j}) \left[(\tilde{\mu}_t^i)^{1-\varepsilon} - w_{t+j} \frac{(\tilde{\mu}_t^i)^{-\varepsilon}}{A_{t+j}} \right].$$

Following the usual procedure to solve Calvo-style models, we obtain the following modified Phillips Curve:

$$\hat{\pi}_t = \frac{(1 - \beta\xi)(1 - \xi)}{\xi} \theta \hat{N}_t + \beta \hat{\pi}_{t+1}, \quad (40)$$

and the following intertemporal Euler condition:

$$\hat{N}_{t+1} - \hat{N}_t = \frac{-rr + r_{t+1} - \hat{\pi}_{t+1}}{\theta}. \quad (41)$$

It is useful to compare these two equations with those associated with a version of the model with no quality choice:

$$\hat{\pi}_t = \frac{(1 - \beta\xi)(1 - \xi)}{\xi} \hat{N}_t + \beta \hat{\pi}_{t+1}, \quad (42)$$


$$\hat{N}_{t+1} - \hat{N}_t = -rr + r_{t+1} - \hat{\pi}_{t+1}. \quad (43)$$

Comparing these two sets of equations we see that, since $\theta < 1$, the model with quality choice produces a higher response to monetary shocks than the standard model. This difference in amplification is illustrated in Figure 2.²⁷

To understand this difference, it is useful to consider first a flexible-price version of the model without quality choice. In this model, if the central bank raises the nominal interest rate, the price level falls and expected inflation rises, leaving the real interest rate unchanged. As a result, the change in the nominal interest rates has no effect on real variables.

²⁷We use the same parameter values for β , θ , and ν as in the flexible price model. We set $\varepsilon = 0.75$ so firms optimize their price schedule on average every fourquarters. We use a Taylor rule with a 1.5 coefficient on inflation.

Figure 2: Impulse Response Functions for a Negative Monetary Shock



Now consider the model with sticky prices but no quality choice. When the central bank raises the nominal interest rate, only a few firms can lower prices and so the price decline is spread over time. As a result, the real interest rate rises. This rise in the real interest rate makes households want to consume less today and more in the future. The current demand for consumption falls and, since employment is demand determined, hours fall.

The key difference between the model with and without quality choice is that the former exhibits a stronger response of the labor supply to shocks. As a consequence, the wage rate has to fall more to clear the labor market than in the standard model. The rate of inflation becomes more negative than in the standard model, as firms lower prices in response to the lower labor costs. This higher rate of deflation implies that the real interest rate is higher in the model with quality choice than in the standard model. This higher real interest rate is associated with a larger fall in consumption, as households postpone consumption to take advantage of the high real interest rate. The result is a larger fall in employment in the quality-choice model than in the standard model.

5 Conclusion

In this paper, we show that when consumers suffer a reduction in their income, they trade down in the quality of the goods and services they consume. 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 22 to 36 percent of the decline in employment during the recent recession.

We introduce quality choice in both flexible and sticky macro models. In these models, consumers change the quality of what they consume over the course of the business cycle. This behavior amplifies the effects of both real and monetary shocks.

We find that introducing quality choice improves the performance of business cycle

models along two dimensions. First, it generates comovement in labor in consumption and investment goods sectors. Second, it generates an endogenous countercyclical labor wedge that improves the ability of the model to explain the behavior of hours worked in the data.

6 References

Abel, Andrew B. “Asset prices under habit formation and catching up with the Joneses,” *The American Economic Review*: 38-42, 1990.

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.

Calvo, Guillermo A. “Staggered prices in a utility-maximizing framework,” *Journal of Monetary Economics* 12.3: 383-398, 1983.

Christiano, Lawrence J., and Terry J. Fitzgerald “The business cycle: it’s still a puzzle,” *Economic Perspectives*, Federal Reserve Bank of Chicago 22: 56-83, 1998.

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,” forthcoming, *Journal of Political Economy*, 2015.

Nakamura, Emi and Jón Steinsson. “Five facts about prices: A reevaluation of menu cost models,” *The Quarterly Journal of Economics*: 1415-1464, 2008.

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

Shimer, Robert “Convergence in macroeconomics: the labor wedge,” *American Economic Journal: Macroeconomics* 1.1: 280-97 ,2009.

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 16 presents some description of the data used to analyze quality shifts in expenditure in six 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 16: Data Sample Description

Sectors	Data Source	Number of Firms (II)	2007 Annual Sales	
	(I)		of Average Firm (\$m) (III)	% of U.S. Sector (IV)
Accommodation	Compustat and company reports	11	3,088	15%
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

six-digit 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.²⁸

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 Indirect effects of trading down on employment

The total change in employment accounted for by trading down is the sum of two components: the direct effect of quality shifts on employment within the same sector and the indirect effect on employment in other sectors that provide intermediate inputs to the sector that experienced trading down.

Indirect effects can arise from differences in the required total inputs per unit of output across quality tiers. For instance, suppose that low-quality restaurants require fewer inputs produced by other (non-restaurant) sectors than higher-quality restaurants. In this case, a shift from high to low quality not only affects employment within the restaurant sector, but also reduces the amount of inputs that need to be produced by other sectors and subsequently labor in those sectors.

We measure these indirect effects of trading down on employment in a way that is comparable to how the BEA constructs its “Total requirements tables.” These tables describe how a change in inputs required by one sector affects the outputs of other

²⁸Note some establishments are excluded because we only include items that are recorded at the modal unit of measure within the 6-digit product category.

sectors. We compute the following expression,

$$\sum_{j \in \mathbf{J}} \sum_{l \in \mathbf{J}} \left[\frac{N_j}{\text{Output}_j} \right]_{2012} \cdot [\Delta \text{Output in } j \text{ due to quality shifts within } l]. \quad (1)$$

where the term $[N_j/\text{Output}_j]_{2012}$ refers the number of employees per dollar value of output in sector j in 2012. The second term refers to the change over 2007-2012 in the value of output produced by sector j as inputs for sector l due to quality-related shifts in market shares within sector l . The second term is computed as

$$g_{lj,2012} \cdot \left[\text{Output}_{l,2012} \cdot \sum_{i(l) \in \mathbf{Q}} \Delta \text{Market share}_{i(l)} \cdot \left(\frac{\text{Input}_{i(l)}}{\text{Output}_{i(l)}} \right)_{2012} \right]. \quad (2)$$

The first term, $g_{lj,2012}$, denotes the change in output of sector j given a dollar change in total inputs required by sector l (from any sector) in 2012. The second (bracketed) term represents the magnitude (in dollars) of the change in total inputs required by sector l related to shifts in the market shares of the quality tiers within sector l . The product of these two terms gives us the change in output required by sector j due to shifts in quality within sector l .

We obtain data for the first term of equation (2) directly from the BEA “Total Requirements Industry by Industry” table, and construct the empirical counterpart of the second (bracketed) term using data from Compustat, the BLS and the BEA. We use Compustat data to compute $\Delta \text{Market share}_{i(l)}$. The term $(\text{Input}_{i(l)}/\text{Output}_{i(l)})$ denotes the revenue-weighted average input-output ratio for firms in quality tier i within sector l in 2012. We construct the empirical counterparts of the inputs and outputs of each firm to be consistent with the definitions in the BEA Total Requirements tables. The dollar value of output is computed as

$$\text{Output}_{i,2012} = \text{Operating expenses}_{i,2012} - \text{Cost of goods sold}_{i,2012}.$$

The value of the inputs is defined as

$$\text{Input}_{j,2012} = \text{Operating expenses}_{j,2012} - \text{Cost of goods sold}_{j,2012} - \text{Labor costs}_{j,2012}.$$

The operating expenses and cost of goods sold are obtained from Compustat. The coverage of firms labor costs in Compustat is relatively sparse because firms are not required to include this item in their annual reports. We use instead firms' labor costs computed as the number of employees reported at a firm-level (from Compustat) multiplied by the sector-wide average annual wage (from the BLS). The change in the input-output ratio is then multiplied by the dollar value of output in sector l in 2012, to give the change in the dollar value of output produced by sector j due to quality shifts in sector l .

D Quality as a normal good

In this appendix, we show that quality is a normal good in our model. We start with the household's first-order condition for quality:

$$\log [C(q_t)] = \frac{1}{1 - \theta} \frac{q_t}{P(q_t)} P'(q_t). \quad (3)$$

In the static model, income is given by:

$$Y = C(q)P(q).$$

Using equation (3) we obtain:

$$\log(Y) = \frac{1}{1 - \theta} \frac{q_t}{P(q_t)} P'(q_t) + \log [P(q)].$$

Suppose that income rises. Using the form of the production function, we can rewrite the last equation:

$$\log(Y) = \log \left(\frac{1}{A} \right) + \frac{\rho - 1}{\rho} \log \left[\alpha^{\frac{1}{1-\rho}} \left(\frac{1}{q} \right)^{\frac{\rho}{1-\rho}} w^{\frac{\rho}{\rho-1}} + (1 - \alpha)^{\frac{1}{1-\rho}} r^{\frac{\rho}{\rho-1}} \right] + \frac{1}{1 - \theta} \frac{1}{1 + \left(\frac{1-\alpha}{\alpha} \right)^{\frac{1}{1-\rho}} \left(\frac{r}{wq} \right)^{\frac{\rho}{\rho-1}}}.$$

It is easy to see that both the term,

$$\log \left[\alpha^{\frac{1}{1-\rho}} \left(\frac{1}{q} \right)^{\frac{\rho}{1-\rho}} w^{\frac{\rho}{\rho-1}} + (1 - \alpha)^{\frac{1}{1-\rho}} r^{\frac{\rho}{\rho-1}} \right],$$

and the term

$$\frac{1}{1-\theta} \frac{1}{1 + \left(\frac{1-\alpha}{\alpha}\right)^{\frac{1}{1-\rho}} \left(\frac{r}{wq}\right)^{\frac{\rho}{\rho-1}}},$$

are increasing with quality. So, the optimal level of q is increasing in Y .

E Quality and labor intensity

In this Appendix we study the conditions under which labor intensity is increasing in quality. We use the firm's first-order conditions to write employment/sales as:

$$\frac{N}{P(q)C} = \frac{\left[\frac{w}{r} \frac{1-\alpha}{\alpha}\right]^{1/(\rho-1)}}{\frac{r}{1-\alpha} \left\{ \alpha \left[\frac{w}{r} \frac{1-\alpha}{\alpha}\right]^{\rho/(\rho-1)} + (1-\alpha)q^{1/(1-\rho)} \right\}}.$$

This expression implies that labor intensity rises with quality only when $\rho < 0$.

It is interesting to consider two additional measures of labor intensity, the labor-capital ratio and the labor share, even though we cannot construct empirical counterparts to these measures. Using the first-order conditions for the firms' problem we can write these measures as:

$$\frac{N}{K} = \left[\frac{1-\alpha}{\alpha} \frac{w}{r} \right]^{1/(\rho-1)} q^{\rho/(\rho-1)},$$

$$\frac{wN}{P(q)C} = \frac{1}{1 + (w/r)^{\frac{\rho}{1-\rho}} [(1-\alpha)/\alpha]^{\frac{1}{1-\rho}} q^{\frac{\rho}{1-\rho}}}.$$

The condition $\rho < 0$ is also necessary for these measures of labor intensity to be increasing in q .

F Balanced growth

In this appendix we show that a modified version of the flexible price model is consistent with balanced growth. The economy's planner's problem is:

$$\max U = \sum_{t=0}^{\infty} \beta^t \left\{ \frac{q_t^{1-\theta}}{1-\theta} \frac{(C_t/X_t)^{1-\sigma} - 1}{1-\sigma} - X_t^{1-\theta} \phi \frac{N_t^{1+\nu}}{1+\nu} \right\},$$

s.t.

$$C_t = A_t \left[\alpha \left(\frac{X_t^2 N_t^C}{q_t} \right)^\rho + (1 - \alpha) (K_t^C)^\rho \right]^{\frac{1}{\rho}}, \quad (4)$$

$$K_{t+1} = A_t \left[\alpha (X_t N_t^I)^\rho + (1 - \alpha) (K_t^I)^\rho \right]^{\frac{1}{\rho}} + (1 - \delta) K_t, \quad (5)$$

$$N_t^C + N_t^I = N_t,$$

$$K_t^C + K_t^I = K_t.$$

Making the utility function compatible with balanced growth requires two modifications. The first is to scale the disutility of labor by $X_t^{1-\theta}$. Without this modification, labor effort increases over time. We can interpret $X_t^{1-\theta}$ as representing technical progress in home production. The second modification is that $C_t^{1-\sigma}$ needs to be replaced with $(C_t/X_t)^{1-\sigma}$. This modification resembles Abel's (1990) external habit formulation.

In order for quantities to grow at a constant rate in the steady state we need, as usual, labor-augmenting technical progress. If the production function for consumption takes the form:

$$C_t = A_t \left[\alpha \left(\frac{X_t N_t^C}{q_t} \right)^\rho + (1 - \alpha) (K_t^C)^\rho \right]^{\frac{1}{\rho}},$$

C_t grows in the steady state at the same rate as X_t but the quality of the goods consumer, q_t , remains constant. In order for both C_t and q_t to grow at the same rate as X_t we need labor-augmenting technical progress to depend on X_t^2 as in equation (4).

It is easy to see that the resource constraints (4) and (5) are consistent with C_t , q_t , and K_t growing at the same rate as X_t .

To show that the modified model is consistent with balanced growth, we use the first-order conditions for the planner's problem. Combining the first-order condition

for C_t and q_t we obtain:

$$\frac{(C_t/X_t)^{1-\sigma} - 1}{1-\sigma} = \frac{1}{1-\theta} \frac{q_t}{C_t} A_t \left[\alpha \left(\frac{N_t^C}{q_t/X_t} \right)^\rho + (1-\alpha) (K_t^C/X_t)^\rho \right]^{\frac{1-\rho}{\rho}} \alpha (N_t^C)^\rho (X_t/q_t)^{1+\rho}.$$

Combining the first-order condition for N_t and C_t we obtain:

$$\left(\frac{X_t}{q_t} \right)^{1-\theta} \phi N_t^\nu = \frac{q^{1-\theta}}{1-\theta} \left(\frac{C_t}{X_t} \right)^{-\sigma} A_t \left[\alpha \left(\frac{N_t^C}{q_t/X_t} \right)^\rho + (1-\alpha) \left(\frac{K_t^C}{X_t} \right)^\rho \right]^{\frac{1-\rho}{\rho}} \alpha \left(\frac{1}{q_t/X_t} N_t^C \right)^\rho (N_t^C)^{-1}.$$

Both equations are consistent with C_t , q_t , and K_t growing at the same rate as X_t .