

Federal Reserve Bank of Minneapolis
Research Department Staff Report 517

October 2015

Appendix: Resurrecting the Role of the Product Market Wedge in Recessions*

Mark Bilal

University of Rochester and NBER

Peter J. Klenow

Stanford University and NBER

Benjamin A. Malin

Federal Reserve Bank of Minneapolis

*The views expressed herein are those of the authors and not necessarily those of the Federal Reserve Bank of Minneapolis or the Federal Reserve System.

This appendix provides a detailed description of our data, a more thorough explanation of some calculations, and some robustness results.

A1. The Aggregate Wedge

A1.1. Representative Agent Wedge

Variables used to construct the RAW include:

- $\frac{y_t}{n_t}$: (Real) Output per hour; BLS, Labor Productivity & Costs, Business Sector.
- c_t : (Real) Nondurables and services consumption per adult equivalent; NIPA consumption data, adjusted for indirect taxes following Prescott (2004) and McDaniel (2007). Adult-equivalent population = (Population ≥ 16) + 0.5(Population ≤ 15).
- n_t : Hours worked per capita; Hours worked from BLS (LPC, Business Sector), and population is Civilian Pop 16+.
- $\tau_t = ((\tau_t^c + \tau_t^n)/(1 + \tau_t^c))$, where τ_t^c is the average tax rate on consumption, following McDaniel (2007), and τ_t^n is the average marginal labor tax rate, using NBER TaxSim to extend Barro-Redlick (2011) through 2012.

A1.2. Extensive and Intensive Margin Wedges

Some variables (for example, y_t/n_t and c_t) used to construct the IMW and EMW are the same as used for the RAW. Additional variables include:

- h_t : Average weekly hours worked (per worker); BLS, LPC, Business Sector.
- v_t : Vacancies (per capita); Pre-1995 is help-wanted index, and post-1995 is Barnichon's (2010) spliced series of help-wanted and JOLTS. Population is 16+.

- m_t : Matches (per capita); Post-1994 from Fallick-Fleischman (2004), and pre-1994 is backcast using data on unemployment and vacancies, following Blanchard and Diamond (1989).

All variables are seasonally adjusted.

The calibration is described in the text with the exception of ψ , the fixed (utility) cost of employment. One can derive an expression for ψ by combining the steady-state optimality conditions for the extensive and intensive margins and assuming the EMW and IMW are the same in steady state. The result is $\psi \equiv \frac{h^{1+1/\eta}}{\eta+1} \left[1 - (\eta + 1) [1 - \beta(1 - \delta)] \left[\frac{\kappa v}{\phi m} + \gamma \right] \right]$.

The EMW includes expectational terms in S_t , for example, $\mathbb{E}_t \left\{ \frac{u'(c_{t+1})}{u'(c_t)} \frac{y_{t+1}/n_{t+1}}{y_t/n_t} \right\}$. We construct these using three-variable, four-lag VARs consisting of real GDP growth, aggregate (log) hours worked, and the respective expectational term. We estimate the VAR using data over the entire sample period, and then use the estimated coefficients to construct time series of the expectational terms.

Finally, we constructed the EMW using alternative data. We assumed $\frac{\beta u'(c_{t+1})}{u'(c_t)} = \frac{1}{1+r_{t+1}}$ and measured the (ex post) real interest rate, r_{t+1} , as the three-month T-bill rate less (realized) core PCE inflation at $t + 1$. Our results change little. The cyclical elasticity of the EMW with respect to GDP was -1.89 (s.e. 0.28) and with respect to aggregate hours was -1.54 (0.15).

A1.3. Aggregate Wedge Decomposition

The decomposition requires wage measures. For our baseline (labeled AHE), we assume $\frac{w_t n_t}{p_t y_t}$ is the labor share of income as measured in the BLS's LPC Business Sector. Because we also have a series for labor productivity $\frac{y_t}{n_t}$, we can back out the average real wage in the economy.

Kudlyak (2014) estimated the semi-elasticities of average hourly earnings, new hire wages, and the user cost of labor, respectively, to the unemployment rate. We use these estimated elasticities, along with the time series of unemployment and our (baseline) average wage measure, to construct time

Figure A1: Weekly Hours: Self-Employed vs. Wage Earners



series for new hire wages and the user cost of labor.

A2. Self-Employed

Figures A1–A4 display the time series of the variables that underlie our estimates of the cyclicity of the all-worker and self-employed wedges, respectively. Figure A1 displays HP-filtered (log) indices for hours per week for both the self-employed and wage earners, while Figure A2 presents the same comparison for annual hours. Figure A3 presents HP-filtered aggregate labor productivity, self-employed income per hour, and income per hour for the unincorporated self-employed. We construct the time series in these three figures using data from the March CPS, as described in Section 3 of the main paper.

Figure A4 presents our consumption series for the self-employed together with aggregate consumption, both HP-filtered. To construct consumption for the self-employed, we use the Consumer Expenditure Surveys (CE) from 1987 through 2012 to get a quarterly series for the growth rate of consumption of the

Figure A2: Annual Hours: Self-Employed vs. Wage Earners



Figure A3: Alternative Productivity Measures

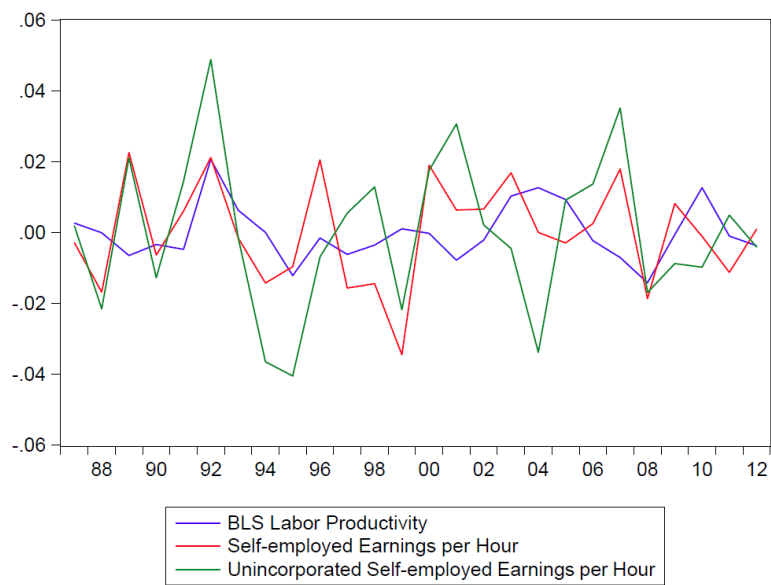



Figure A4: Alternative Consumption Measures



self-employed relative to that for a representative sample of households in the CE. The relative growth rate, in turn, is integrated to obtain a series for relative consumption of the self-employed, indexed to the beginning of 1987. We add this relative estimate to NIPA aggregate consumption to arrive at an estimate of the cyclical nature of consumption for the self-employed. The following paragraphs describe the construction of the quarterly growth rates of consumption for self-employed workers.

The CE has been an ongoing quarterly survey since 1980, with about 5,000 households interviewed each quarter. Households are asked about their detailed expenditures for the previous three months. Each household is surveyed up to four consecutive quarters, allowing construction of up to three observations on quarterly growth. We focus on expenditures on nondurables and services, which we construct by aggregating individual categories that are clearly not durables by NIPA standards. We include expenditures on housing: for renters this is captured by household rent; for home owners it reflects the owner's estimate of its rental value (rental equivalence). The categories we can classify as nondurables and services constitute about two-thirds of household

expenditures. We deflate these expenditures by the GDP deflator for nondurables and services. Individual growth rates across any two quarters are calculated by the midpoint formula to reduce the impact of extreme values.

During each of the first and fourth quarterly interviews on expenditures, households are surveyed about the work experience of their household members during the past 12 months. We focus on the work history in the latter survey, as the work history over the prior 12 months conforms to the time frame for reported expenditures. (For a small number of households, we fill in for missing employment information from responses collected in earlier quarters.) We create a sample of workers from the CE households, including all members that meet our sample requirements. These requirements are chosen to mimic our treatment of the CPS data: (i) individuals must be between ages 20-70; (ii) they must report working at least 10 weeks during the year, at a workweek of 10 hours or more when working; (iii) we exclude workers in the top or bottom 9.6 percent of the income distribution and the top 1.2 percent of hours per week. These last exclusions are chosen to match those we made on the CPS data, dictated by its top-coding of income and hours. We make two other sample restrictions in order to measure quarterly growth rates of household consumption. We exclude households in the top and bottom 1 percent of expenditures in any quarter in order to eliminate top-coded expenditures and outliers. We exclude households that exhibited a change in household size across the quarters that are the basis for the growth rate. In all calculations we employ the CE sampling weight that is designed to make the sample representative of the U.S. civilian noninstitutionalized population.

We classify workers as self-employed, as opposed to wage earners, if they report that the job for which they received most income was self-employment and, in fact, at least 95 percent of their reported income over the past 12 months is from (nonfarming) self-employment. This conforms well to our definition in Section 3 of the main paper based on CPS data. We do not observe consumption at the individual level (e.g., for a self-employed member versus a

wage-earning member). Thus, we have to make the simplifying assumption that households equate consumption across members. For example, if a household has one self-employed worker and one wage earner, then that household contributes two members to our overall sample and one member to our self-employed sample. But the growth rate in consumption in any quarter will be the same for both members of that household. We have 11,849 quarterly observations on consumption growth that apply for self-employed workers, which equals 115 per quarter on average.

A3. Intermediates

We first derive an industry-level, intensive margin total wedge using more general technology and preferences than we used for our baseline results. We then derive the industry-level, *extensive margin* total wedge. Finally, we describe the data used in our calculations.

The *gross output* production function implies a marginal product of labor on the intensive margin of

$$mpn_{it}^{int} = \alpha(1 - \theta) \left(\frac{y_{it}}{v_{it}} \right)^{\frac{1}{\varepsilon}} \left(\frac{v_{it}}{n_{it}} \right)^{\frac{1}{\omega}} (z_{v,it} z_{n,it})^{\frac{\omega-1}{\omega}} e_{it}.$$

For our baseline, $\varepsilon = \omega = 1$, this simplifies to $mpn_{it}^{int} = \alpha(1 - \theta) \frac{y_{it}}{n_{it}} e_{it}$.

Our baseline used the following preferences:

$$\mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \left\{ \frac{c_t^{1-1/\sigma}}{1-1/\sigma} - \nu \sum_i \left[\left(\frac{h_{it}^{1+1/\eta}}{1+1/\eta} + \psi \right) e_{it} \right] \right\},$$

so the marginal rate of substitution of consumption for an extra hour per worker in industry i is $mr s_{it}^{int} = \nu h_{it}^{1/\eta} c_t^{1/\sigma} e_{it}$.

Thus, our baseline industry- i (intensive margin) labor wedge is (up to an

additive constant)

$$\begin{aligned} \ln(\mu_i^{int}) &= \ln\left(\frac{p_i m p n_i^{int}}{p m r s_i^{int}}\right) = \ln\left(\frac{p_i y_i}{p n_i}\right) - \left[\frac{1}{\sigma} \ln(c) + \frac{1}{\eta} \ln(h_i)\right] \\ &= \ln\left(\frac{p_i \frac{v_i}{n_i}}{p \frac{v}{n}}\right) + \ln\left(\frac{y_i}{v_i}\right) - \frac{1}{\eta} \ln\left(\frac{h_i}{h}\right) + \ln\left(\frac{m p n_i^{int}}{m r s_i^{int}}\right), \end{aligned} \quad (1)$$

where $m p n_t^{int} \equiv \alpha(1 - \theta) \frac{v_t}{n_t} e_t$ and $m r s_t^{int} \equiv \nu h_t^{1/\eta} c_t^{1/\sigma} e_t$ are based on aggregate data. For $\varepsilon, \omega \neq 1$, it is straightforward to see how the total wedge would be altered. Specifically, for $\varepsilon < 1$, the total wedge becomes less countercyclical if gross output is more procyclical than value added.

Note that our preferences assume separability across labor supply in different industries. This seems reasonable because workweeks are person-specific. But, we could consider alternative preferences, say,

$$\mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \left\{ \frac{c_t^{1-1/\sigma}}{1-1/\sigma} - \nu \left(\frac{h_t^{1+1/\eta}}{1+1/\eta} + \psi \right) e_t \right\},$$

where $h_t \equiv \frac{\sum_i h_{it} e_{it}}{\sum_i e_{it}}$ and $e_t \equiv \sum_i e_{it}$. In this case, $m r s_{it}^{int} = \nu h_t^{1/\eta} c_t^{1/\sigma} e_{it}$. The industry- i labor wedge is thus (for baseline technology, $\varepsilon = \omega = 1$)

$$\ln(\mu_i^{int}) = \ln\left(\frac{p_i \frac{v_i}{n_i}}{p \frac{v}{n}}\right) + \ln\left(\frac{y_i}{v_i}\right) + \ln\left(\frac{m p n_i^{int}}{m r s_i^{int}}\right).$$

Under these preferences, labor supply is perfectly substitutable across industries and only aggregate labor supply, e and h , matters. The labor wedge no longer needs an adjustment for industry-specific workweeks. Table A1 shows how replacing industry-specific workweeks with aggregate average weekly hours worked affects our results (compare with Table 8 in the main paper). Manufacturing industries exhibit more procyclical workweeks, and the labor wedge is thus less countercyclical for manufacturing. On the other hand, it is more countercyclical for nonmanufacturing and all industries. (For the latter, recall that the 60 industries covered by KLEMS are not necessarily

Table A1: Cyclicalty of (Common MRS) Intensive Margin Total Wedge

	Elasticity wrt	
	GDP	Total Hours
All Industries	-1.10 (0.26)	-0.72 (0.13)
Manufacturing	-0.55 (0.39)	-0.35 (0.20)
Non-Manufacturing	-1.25 (0.24)	-0.82 (0.12)

Note: Each entry is from a separate regression. Annual data are from 1987 to 2012 for 60 industries (1,560 industry-year observations): 18 manufacturing and 42 nonmanufacturing. All variables are in logs and HP-filtered. Regressions include industry fixed effects and use industry average value-added shares as weights. Standard errors are clustered by year.

representative of the entire economy.)

Moving to the extensive margin, $mpn_{it}^{ext} = mpn_{it}^{int} \frac{h_{it}}{e_{it}}$ and $mrs_{it}^{ext} = mrs_{it}^{int} \frac{\Omega_{it}}{h_{it}^{1/\eta}} \frac{h_{it}}{e_{it}}$. The industry- i extensive margin labor wedge is

$$\ln(\mu_i^{ext}) = \ln\left(\frac{p_i mpn_i^{ext}}{p mrs_i^{ext}}\right) - S_i = \ln(\mu_i^{int}) - \ln\left(\frac{\Omega_i}{h_i^{1/\eta}}\right) - S_i,$$

or, for our baseline case, it is

$$\ln(\mu_i^{ext}) = \ln\left(\frac{p_i \frac{v_i}{n_i}}{p \frac{v}{n}}\right) + \ln\left(\frac{y_i}{v_i}\right) - \ln\left(\frac{\Omega_i}{\Omega}\right) + \ln(\mu^{ext}) - (S_i - S). \quad (2)$$

Because of data limitations (i.e., vacancies and matches are not available at the industry level), we assume S_{it} differs across industry only because of industry-specific workweek movements. That is, $S_{it} = \left[\frac{h_i/h_{it}}{h/h_t}\right] S_t$. Table A2 displays the cyclicalty of the extensive margin total wedge. The results are similar to Table A1.

To construct the industry-level total wedge and the intermediates-based product market wedge, some variables (e.g., c_t) are the same as used earlier in

Table A2: Cyclicity of (Common MRS) Extensive Margin Total Wedge

	Elasticity wrt	
	GDP	Total Hours
All Industries	-1.14 (0.52)	-0.86 (0.28)
Manufacturing	-0.59 (0.66)	-0.49 (0.35)
Non-Manufacturing	-1.28 (0.49)	-0.96 (0.26)

Note: Each entry is from a separate regression. Annual data are from 1987 to 2012 for 60 industries (1,560 industry-year observations): 18 manufacturing and 42 nonmanufacturing. All variables are in logs and HP-filtered. Regressions include industry fixed effects and use industry average value-added shares as weights. Standard errors are clustered by year.

the paper. Additional variables include:

- p_t : Price deflator for nondurables and services consumption; Tornqvist index of NIPA implicit price deflators for nondurables and services.
- $p_{it}, y_{it}, n_{it}, p_{mit}m_{it}$: Respectively, the gross output deflator, real gross output, hours worked, and expenditures on intermediates (Tornqvist index of materials, energy and services) by industry from BLS KLEMS.
- h_{it} : Average weekly hours worked (per worker); ratio of hours worked (from BLS KLEMS) to industry-specific employment (calculated with data underlying BLS LPC dataset).

A4. Other Nonwage Decompositions

A4.1. Advertising

Hall (2014) considers a simple theory of advertising (further simplified here), in which a firm's objective is

$$\max_{p,A} (p - mc) \frac{A^\alpha}{p^\epsilon} - \kappa A,$$

where p is the firm's price, A its advertising volume, mc the marginal cost of production, κ the cost of a unit of advertising, and $-\epsilon$ and α are the elasticities of demand with respect to price and advertising. The first-order condition for advertising yields an expression for the ratio of advertising expenditure to revenue:

$$\frac{\kappa A}{pQ} = \alpha \left[1 - \frac{1}{p/mc} \right]. \quad (3)$$

Hall's finding that the advertising expenditure share of revenue is acyclical, combined with equation (3), suggests that markups are also acyclical.

But, as stated in the main text, if advertising spending displays a constant elasticity impact on consumers' reservation prices, rather than on quantity demanded, that implication no longer holds.¹ The firm's objective then becomes

$$\max_{p,A} (p - mc) \left(\frac{p}{A^\alpha} \right)^{-\epsilon} - \kappa A,$$

and optimal advertising requires

$$\frac{\kappa A}{pQ} = \alpha\epsilon \left[1 - \frac{1}{p/mc} \right]. \quad (4)$$

In this case, an increase in the price elasticity of demand lowers the price markup — that is, $p/mc = \epsilon/(\epsilon - 1)$ — but has no effect on the advertising share. Why? The reduced benefit of advertising, from the decline in p/mc , is exactly

¹Assume a fixed population and that individual i 's willingness to pay for a good is given by $x_i = ZA^\alpha\Omega_i$, where Z is an aggregate shifter, A is advertising for the good, and Ω_i is the individual preference. If Ω_i is distributed basic Pareto, $f(\Omega_i) = \epsilon\Omega_i^{-(1+\epsilon)}$ for $\Omega_i \geq 1$, then demand for the good is $Z^\epsilon A^{\alpha\epsilon} p^{-\epsilon}$, where p is the price of the good.

canceled by the more elastic response of sales to that advertising.

A4.2. Inventories

Here, we show how data on work-in-process (WIP) inventories can be used to infer a price markup. Following Christiano (1988), we assume a production function that uses WIP inventories as one of its inputs. For a firm in industry i ,

$$y_{it} = g(z_{it}, n_{it}, k_{it})q_{it}^{\varphi_{it}},$$

where y_{it} denotes output, q_{it} is beginning-of-period inventories, and z_{it} , n_{it} , and k_{it} are TFP, hours worked, and capital, respectively. The elasticity of output with respect to inventories, φ_{it} , is allowed to vary across both industry and time. The law of motion for inventories is assumed to be

$$q_{i,t+1} = (1 - \delta_q)q_{it} + y_{it} - y_{it}^f,$$

where δ_q is the depreciation rate of inventories, $y_{it}^f \geq 0$ is output of finished goods, and $q_{i,t+1} \geq 0$. That is, total output y_{it} is the sum of gross investment in WIP inventories and finished-good output. The latter includes both final sales and (gross) investment in finished-goods inventories, but it is not necessary to separate these two for our purposes.

An optimizing firm minimizes the expected present discounted cost of producing a given path of finished goods. One perturbation on its cost-minimizing strategy would be to produce an additional unit of output in the form of WIP inventories at time t and then reduce production just enough at $t + 1$ — that is, by $\left(1 - \delta_q + \varphi_{i,t+1} \frac{y_{i,t+1}}{q_{i,t+1}}\right)$ — to keep inventories unaffected at $t + 2$ forward. At an optimum,

$$\frac{mc_{it}}{p_t} = \mathbb{E}_t \left[M_{t,t+1} \frac{mc_{i,t+1}}{p_{t+1}} \left(1 - \delta_q + \varphi_{i,t+1} \frac{y_{i,t+1}}{q_{i,t+1}} \right) \right], \quad (5)$$

where $\frac{mc_{it}}{p_t}$ is the (real) marginal cost of production and $M_{t,t+1}$ is the firm's

discount factor. In words, the firm equates the marginal cost of output to its marginal benefit, which is reduced future production costs.² Because the industry product market wedge is $\mu_{it}^p \equiv \frac{p_{it}}{mc_{it}}$, we can write (5) as

$$\frac{p_{it}/p_t}{\mu_{it}^p} = \mathbb{E}_t \left[M_{t,t+1} \frac{p_{i,t+1}/p_{t+1}}{\mu_{i,t+1}^p} \left(1 - \delta_q + \varphi_{i,t+1} \frac{y_{i,t+1}}{q_{i,t+1}} \right) \right]. \quad (6)$$

We assume the stochastic discount factor is given by $M_{t,t+1} \equiv \beta \frac{u'(c_{t+1})}{u'(c_t)}$ and that the joint conditional distribution of $\frac{u'(c_{t+1})}{u'(c_t)}$, $\frac{1}{\mu_{i,t+1}^p}$, $\frac{p_{i,t+1}}{p_{t+1}}$, and $1 - \delta_q + \varphi_{i,t+1} \frac{y_{i,t+1}}{q_{i,t+1}}$ is log-normal and homoskedastic.³ We then take logs of equation (6) and get (up to a constant)

$$\ln(\mu_{it}^p) \approx \ln \left(\frac{p_{it}}{p_t} \right) + \mathbb{E}_t \left\{ -\ln \left(\frac{u'(c_{t+1})}{u'(c_t)} \right) + \ln(\mu_{i,t+1}^p) - \ln \left(\frac{p_{i,t+1}}{p_{t+1}} \right) - \frac{\varphi_{i,t+1} y_{i,t+1}}{1 - \delta_q q_{i,t+1}} \right\},$$

where $\frac{\varphi_{i,t+1} y_{i,t+1}}{1 - \delta_q q_{i,t+1}} \approx \ln \left(1 + \frac{\varphi_{i,t+1} y_{i,t+1}}{1 - \delta_q q_{i,t+1}} \right)$. Iterating forward for $\ln(\mu_{i,t+s}^p)$ and using $u'(c_t) = c_t^{-1/\sigma}$ yields the inventory-based product market wedge:

$$\ln(\mu_{it}^p) \approx -\frac{1}{\sigma} \ln(c_t) + \ln \left(\frac{p_{it}}{p_t} \right) - \mathbb{E}_t \sum_{s=1}^{\infty} \frac{\varphi_{i,t+s} y_{i,t+s}}{1 - \delta_q q_{i,t+s}} + \text{constant terms}. \quad (7)$$

The intuition for equation (7) is as follows. Suppose the economy is in a recession in period t , so the log marginal utility of consumption, $-\ln(c_t)/\sigma$, is high. If the firm's price markup and relative price are not cyclical, then (7) says the path of future output-to-inventory ratios must be high. That is, the firm should be depleting future WIP inventories in order to push output out the door

²Note that an optimizing firm will *always* produce to the point that the marginal value of an extra unit of output equals its marginal cost. In a model in which the firm can adjust sales at the margin, the marginal value of output is simply marginal revenue. If the firm cannot adjust sales, the additional unit of output is held as an inventory and valued accordingly. The value of a finished-good inventory is the expected discounted revenue it generates when it is eventually sold. The value of a WIP inventory, on the other hand, is that the firm enters the next period with a larger stock of WIP inventories.

³As explained by Campbell (2003), log-normality implies the log of an expectation can be expressed as an expectation of the log plus a variance term. The conditional homoskedasticity means the variance term is not time-varying.

today and boost consumption.

Alternatively, if the expected path of output-to-inventory ratios is not cyclical, then for equation (7) to hold, the firm's real marginal cost (mc_{it}/p_t) must be low in recessions. In turn, either the product market wedge (μ_{it}^p) is high or the firm's relative price (p_{it}/p_t) is low in recessions. That is, if firms do not deplete inventory investment in recessions, one explanation is that product market distortions keep the firm's price high relative to its marginal cost.

To measure the product market wedge according to equation (7), we turn to NIPA, which provides quarterly and monthly measures of inventories, sales, and sales price deflators by industry. We define industry output as sales plus the change in (total) inventories, and we use quarterly data from 1987 to 2012 for comparison with previous sections.⁴ WIP inventories are available for 22 (roughly two-digit) manufacturing industries, but the industry classification changed from the SIC to NAICS in 1997. To create consistent industry definitions, we aggregate some industries, leaving 14 sectors.⁵

To calibrate the parameters in equation (7), we first note that inventory-to-output ratios exhibited significant low-frequency movement over our sample period. We thus let φ_{it} vary over time and set $\varphi_{it} = \left[\frac{1}{\beta} - (1 - \delta_q) \right] \frac{\bar{q}_{it}}{\bar{y}_{it}}$, where $\frac{\bar{q}_{it}}{\bar{y}_{it}}$ is a quadratic trend fitted to the inventory-output ratio.⁶ Our quarterly calibration sets $\beta = 0.996$ and $\delta_q = 0.01$. As a result φ_{it} , which measures the share of output attributable to inventories, is quite low, about 0.2 percent, on average.

Constructing the inventory-based wedge requires computing, at each point

⁴Specifically, the output-to-WIP-inventory ratio, $\frac{y_{it}}{q_{it}}$, and price deflator for (industry) sales, p_{it} , are taken from the NIPA Underlying Detail Tables, Real Inventories and Sales.

⁵We use a Tornqvist index to construct chain-weighted growth rates of real sales, real inventories, and price deflators for the combined industries. For bridging across the 1996-97 break, we made two assumptions. For inventories, we assume the industry shares of nominal inventories do not change between December 1996 and January 1997. (This is feasible, since the inventory data are reported for both classifications in 1997, but there is no such overlap for the sales.) For sales, we assume the growth rate in the nominal inventory-to-shipments ratio is the same as that of the real inventory-to-shipments ratio (in January 1997). The former is constructed using data from the Census M3 survey, which has a consistent NAICS industry classification across 1996-97.

⁶This specification for φ_{it} ensures that equation (6) holds in (detrended) steady state.

in time, the sum of expected future output-to-inventory ratios. We estimate industry-specific, three-variable, 12-(monthly)-lag VARs consisting of real GDP growth, aggregate (log) hours worked, and the industry-specific output-to-inventory ratio. The latter two variables are quadratically detrended. We estimate the VARs using data over the entire sample period, and then use the estimated coefficients to produce a time series for the expected sum of future output-to-inventory ratios.⁷

Figure A5 plots the weighted-average industry μ^p against GDP. As shown, the wedge is quite countercyclical. This is also true if we define the cycle in terms of hours worked. Figure A6 plots μ^p again, but now aggregated to an annual frequency and plotted against the weighted-average manufacturing-industry total wedge constructed in Section 4 of the main paper rather than against GDP. The product market wedge accounts for most of the cyclical variation in the total wedge.

We next run regressions of the industry-level wedge on the cycle

$$\log(\mu_{it}^p) = \alpha_i + \beta^p \log(cyc_t) + \varepsilon_{it},$$

where the weights are the industry's average share of output and standard errors are clustered by period. Table A3 displays the results at an *annual* frequency for comparison with the total wedge.⁸ The strongly countercyclical μ^p (-0.70 elasticity with respect to GDP) accounts for nearly all of the cyclicity in the total wedge (-0.73).

Finally, we have used WIP inventories for our calculations because these

⁷We considered a second approach to calculating the expected sum of future output-to-input ratios, which involved truncating the sum at either four or eight quarters and calculating the (ex post) realized sum. Because we project the constructed wedge on the time- t business cycle, using the (ex post) realized values is valid for our purposes. It does require using a one-sided HP-filter for the business cycle, so the difference between expected and realized values of the output-to-inventory ratios is orthogonal to the time- t cycle. This second approach produced results for the wedge that were very similar to the VAR approach.

⁸The quarterly elasticities are more precisely estimated: -0.80 (s.e. 0.12) with respect to GDP and -0.33 (0.08) with respect to hours.


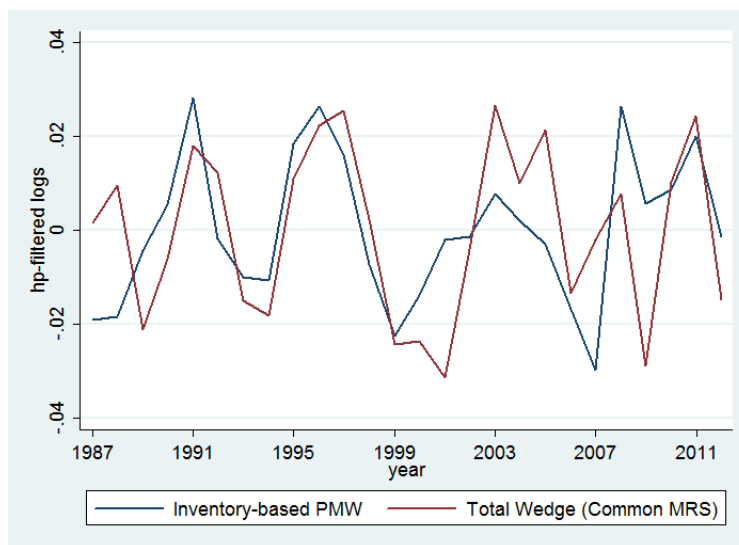
Figure A5: Cyclicity of Inventory-Based μ^P Figure A6: Inventory-Based μ^P vs. Total Wedge μ 

Table A3: Cyclicalty of Inventory-Based μ^P

	Elasticity wrt	
	GDP	Total Hours
Product market wedge	-0.70 (0.26)	-0.26 (0.15)
Marginal utility of consumption	-1.33 (0.06)	-0.76 (0.08)
Relative price	0.83 (0.20)	0.60 (0.11)
Expected output/inventory path	0.21 (0.05)	0.10 (0.04)

Note: Each entry is from a separate regression. Annual data are from 1987 to 2012 for 14 manufacturing industries (364 industry-years). Variables are in logs and HP-filtered. Regressions include industry fixed effects and use industry average value-added shares as weights. Standard errors are clustered by year. See equation (7) for the wedge components.

align most closely with the theory, which posits a role for inventories in *production*. Christiano (1988) argues for total inventories (i.e., including materials, WIP, and final goods inventories), noting that labor inputs can be conserved by transporting materials in bulk and holding finished inventories. For robustness, we redo our calculations using total inventories instead of WIP inventories. The results are fairly similar to those reported in Table A3: the cyclical elasticity of the product market wedge is -0.56 with respect to GDP and -0.22 with respect to hours.⁹

References

Barnichon, Regis, “Building a Composite Help-Wanted Index,” *Economics Letters*, 2010, 109 (3), 175–178.

⁹Using total inventories enables one to consider industries outside of manufacturing (e.g., wholesale and retail trade).

- Barro, Robert J. and Charles J. Redlick, “Macroeconomic Effects from Government Purchases and Taxes,” *Quarterly Journal of Economics*, 2011, 126 (1), 51–102.
- Blanchard, Olivier J. and Peter Diamond, “The Beveridge Curve,” *Brookings Papers on Economic Activity*, 1989, 1, 1–76.
- Campbell, John Y., “Consumption-Based Asset Pricing,” in G.M. Constantinides, M. Harris, and R. Stulz, eds., *Handbook of the Economics of Finance*, Vol. 1B, 2003, chapter 13, pp. 803–887.
- Christiano, Lawrence J., “Why Does Inventory Investment Fluctuate So Much?,” *Journal of Monetary Economics*, 1988, 21 (2-3), 247–280.
- Fallick, Bruce and Charles A. Fleischman, “Employer-to-Employer Flows in the U.S. Labor Market: The Complete Picture of Gross Worker Flows,” May 2004. FEDS Working Paper 2004-34.
- Hall, Robert E., “What the Cyclical Response of Advertising Reveals about Markups and Other Macroeconomic Wedges,” April 2014. Unpublished manuscript.
- Kudlyak, Marianna, “The Cyclicity of the User Cost of Labor,” *Journal of Monetary Economics*, 2014, 68, 53–67.
- McDaniel, Cara, “Average Tax Rates on Consumption and Investment Expenditures and Labor and Capital Income in the OECD 1950-2003,” 2007.
- Prescott, Edward C., “Why Do Americans Work So Much More Than Europeans?,” *Federal Reserve Bank of Minneapolis Quarterly Review*, 2004, 28 (1), 2–13.