

What the Cyclical Response of Advertising Reveals about Markups and other Macroeconomic Wedges *

Robert E. Hall
Hoover Institution and Department of Economics,
Stanford University
National Bureau of Economic Research
rehall@stanford.edu; stanford.edu/~rehall

March 21, 2013

Abstract

Theory suggests advertising should be remarkably sensitive to profit margins. Firms advertise to stimulate demand for their products. They advertise high-margin products aggressively and low-margin ones hardly at all. In modern macroeconomics, wedges are potent sources of fluctuations in employment. The profit margin or markup ratio is a leading example. In an important class of fluctuations models, profit margins rise in recessions and mediate the decline in employment. But a rise in profit margins should expand advertising by a lot. Really a lot. Advertising should be highly countercyclical. Instead, it is somewhat procyclical. The ratio of advertising spending to private GDP falls when the economy contracts. The behavior of advertising refutes the hypothesis that profit margins rise. But it is true that the labor share of income falls. Hence there must be another factor that lowers the labor share without raising profit margins. The only influence that fits the facts is a rise in a product-market friction or wedge that has the same effect as an increase in sales taxes. The cyclical behavior of advertising should point macroeconomics in a somewhat different direction in explaining employment fluctuations.

JEL D43 E12 E32

*The Hoover Institution supported this research. The research is also part of the National Bureau of Economic Research's Economic Fluctuations and Growth Program. I am grateful to Kyle Bagwell, Mark Bils, Chris Edmond, Valerie Ramey, Julio Rotemberg, Stephen Sun, and Michael Woodford for valuable comments. A file containing all of the data and calculations is available on my website.

Theorem: Let R be the ratio of advertising expenditure to the value of output. Let $-\epsilon$ be the residual elasticity of demand. Let m be an exogenous shift in the profit margin. Then the elasticity of R with respect to m is $\epsilon - 1$, which is a really big number.

After proving this theorem, which is a direct implication of the standard model of advertising, I dwell on its implications for an important issue in macroeconomics, the role of shifts in the profit margin. The basic idea is simple. In a slump, firms do not cut prices in answer to disappointing sales. If their costs are lower—because they have moved down their upward-sloping short-run marginal costs curves or because flexible-price factor markets now have lower prices—their profit margins are higher. The theorem says that they should expand advertising by substantial amounts. Consider the middle-of-the-road value for the residual elasticity of demand of 6, so that the ratio of price over marginal cost is $6/(6-1) = 1.2$. The ratio of advertising spending to GDP should rise by 5 times the proportional increase in that ratio. Advertising should be highly countercyclical. Firms should expand advertising aggressively in a slump.

In fact, advertising is procyclical. I show that the ratio of advertising to GDP *falls* by about one percent for each percentage fall in employment. Far from boosting advertising to recover business lost in a slump, firms cut advertising by a larger proportion than their loss of sales. The key finding, however, is that advertising is *not* highly countercyclical. I would have written this paper even if I had found advertising to be noncyclical or mildly countercyclical.

The thrust of standard advertising theory is that advertising should rise and fall in proportion to sales. The formula for the ratio is remarkably simple; it is the elasticity of sales with respect to advertising effort divided by the residual elasticity of demand. If the two elasticities are constants not influenced by the factors causing a slump, then advertising will be a constant fraction of sales. Macroeconomics has brought into play a mechanism not usually considered in advertising theory, namely that profit margins widen in slumps. That widening should result in a splurge of advertising in slumps.

The question at this point is what other factor could be operating to alter the standard property that implies that advertising should be neither procyclical (as it actually is) nor countercyclical (as the widening-profit-margin model implies). In my baseline model, I include a wedge that has the effect on a firm that a sales tax would. I call this a product-market

wedge. The term *wedge*, traditionally used in public finance to describe the effects of taxes, has come into wide use in macroeconomics to describe variables that intervene between two marginal values in theoretical efficiency conditions. The analogy to tax wedges is apt and the effects of macro wedges are similar to the effects of taxes.

The paper studies two key observed variables: (1) the ratio of advertising spending to revenue, and (2) the ratio of labor compensation to revenue (the labor share). Both the profit-margin wedge and the product-market wedge affect these variables. The elasticity of the advertising ratio with respect to the profit-margin wedge is $\epsilon - 1$, a number around 5. The elasticities of the advertising ratio with respect to the product-market wedge and of the labor share with respect to both wedges are all -1 . The fact that the profit-margin wedge has a large effect on the advertising ratio has a neat implication. Consider the ratio of the advertising/sales variable to the labor share. One property is that the elasticity of that ratio with respect to the product-market wedge is zero, because the wedge has the same effect on numerator and denominator. The second property is that the elasticity of the ratio with respect to the profit-margin wedge is the residual elasticity of demand, ϵ , say 6. These facts provide a clean identification of the role of the profit-margin wedge. That wedge should have a big positive effect on advertising in recessions, under the view that profit margins increase in recessions. A regression of the ratio of the advertising/sales variable to the labor share on employment should have a big negative coefficient that arises entirely from the margin effect and not at all from the product-market wedge. In reality, the regression coefficient is slightly positive and the confidence interval around it excludes any big negative effect. The finding casts serious doubt on the countercyclical profit-margin hypothesis.

On the other hand, the product-market wedge emerges as a fully consistent idea about the character of slumps. It says that rising frictions in recessions lower advertising and the labor share about equally, leaving the ratio of the two variables close to noncyclical. I avoid speculation in this paper about the source of the wedge.

I consider a number of potential variations around the basic specification in the paper. The first is to follow an important branch of the advertising literature and treat advertising expenditure as a form of investment. Because investment in, say, plant and equipment, is quite procyclical, this consideration might explain the findings despite a countercyclical margin—the procyclical effect from investment might be swamping the large countercyclical effect of the margin changes. But a reformulation of the model using standard investment

theory shows otherwise. A key factor in this finding is the high depreciation rate of advertising. A consensus of research on advertising is that around 60 percent of the effect of earlier advertising dissipates each year.

I also consider a model extended to include other cyclical shifts. These are (1) changes in productivity, (2) measurement error in the labor share, (3) measurement error in the capital share, and (4) measurement error in the price of advertising. I show that productivity and capital measurement errors have no effect on the measured values of the variables I study. Of course, they do affect other variables—the point is that they drop out of the ratios I consider. A plausible measurement error in the labor share—an idea I take seriously—has only a small effect on the key finding. Measurement error in the price of advertising could conceal part of its countercyclical movements but would have to be implausibly large to overturn the basic conclusion of the paper. The most likely form of such an error would come from failing to consider the investment component of advertising, a topic I consider separately with negative conclusions.

1 Related Research

1.1 Cyclical behavior of advertising

Borden (1942) noted the close correlation between advertising volume and an index of industrial production—see Simon (1970), Figures 2-11 and 2-12, who also cites a number of other sources confirming the correlation. Kaldor (1950) noted a similar correlation and Blank (1962) and Yang (1964) documented the correlation, without theoretical interpretation. Bills (1989), Table 1, presents regressions of the rate of change of real advertising expenditures on the rate of change of real GDP. A coefficient greater than one would indicate procyclical movements as that term is used in this paper. He uses data for the U.S. and Britain. In all cases the coefficients are positive and for more recent U.S. data and all British data, they exceed one. The model in the paper implies countercyclical market power for reasons similar to Edmond and Veldkamp (2009), discussed below, but Bills interprets the model as pointing toward procyclical advertising.

1.2 The level of market power

Positive advertising expenditure proves the existence of market power, for there is no incentive to advertise in perfectly competitive markets. Still, there is remarkably little consensus

on the extent of market power in the U.S. economy. The most recent survey of the subject appears to be Bresnahan (1989). His summary, in Table 17.1, reports residual elasticities in the range from 1.14 to 40, for industries from coffee roasting to banking. Many subsequent studies, mainly for consumer packaged goods, have appeared since the publication of Bresnahan's survey. I am not aware of any attempt to distill a national average from studies for individual products. Hausman, Leonard and Zona (1994), for example, study the demand for beer and find residual elasticities (holding the prices of competing beers constant) in the range from 3.5 to 5.9. Most research does not try to reconcile residual elasticities estimated from demand equations with data on price/marginal cost ratios from producers, though Bresnahan discusses this topic extensively. De Loecker and Warzynski (2012) use firm-level data from Slovenia in a producer-side framework and find average markups of about 1.2, corresponding to a residual elasticity of demand of 6, the value I take in my base case.

1.3 Cyclical changes in market power and profit margins

Macroeconomics has spawned a large literature on countercyclical market power. Bills (1987) launched the modern literature that studies cyclical variation in the labor share. My interpretation of that literature is that it measures not variations in profit margins but rather in the labor share, because these are not the same thing in the presence of the product-market wedge that I consider. Bills made important adjustments based on cyclical variations in the incidence of overtime wages. Rotemberg and Woodford (1999) embraced Bills's adjustments in a survey chapter that explains how New Keynesian models explain cyclical variations in output and employment through variations in market power resulting from sticky prices and flexible cost. Nekarda and Ramey (2010) and Nekarda and Ramey (2011) challenge the findings of countercyclical market power in favor of cyclically constant markups resulting from Bills's overstatement of the incidence and magnitude of overtime premiums.

Bills and Kahn (2000) argue that marginal cost is procyclical and thus profit margins are countercyclical because firms internalize the fluctuations in their employees' disamenity of work effort. In slumps, the marginal disamenity of effort is low, because effort itself is low. In an expansion, as effort rises, its marginal burden on workers rises and marginal cost of production rises accordingly, even if cash payments to workers do not rise in proportion to the marginal burden. They use this hypothesis to explain the otherwise puzzling behavior of inventory investment. Firms allow inventory levels to decline persistently below normal dur-

ing booms and above normal in slumps, which would only make sense if marginal production costs are high in booms and low in slumps.

Chevalier and Scharfstein (1996) develop and estimate a model in which capital-market frictions influence pricing decisions at the retail level. In slumps, firms that are financially constrained disinvest in customers by setting prices at higher than normal margins over marginal cost.

Edmond and Veldkamp (2009) look at the issues of market power from the consumer’s perspective. They find that rising dispersion of income distribution lowers residual elasticities in slumps. Firms respond by setting prices further above marginal cost.

The literature on cyclical changes in market power is complementary to the ideas in this paper. In many of the accounts in the existing literature, the question becomes acute: Why does advertising not expand in slumps when the residual elasticity falls?

Kaplan and Menzio (2013) is an interesting new paper in which the product market becomes more competitive in slumps, because the unemployed shop more intensively than the employed. Their theoretical model is consistent with the findings of this paper that advertising is procyclical. The model does not consider other wedges as potential mediating forces of fluctuations. Rather, its calibration has a sufficiently strong adverse effect of unemployment on incentives for hiring that it generates multiple equilibria, so recessions are times when the economy transits from a good equilibrium to a bad one.

1.4 Cyclical fluctuations in product-market wedges

I am not aware of any empirical work on this topic.

2 Theory

Suppose that the residual demand facing a firm is a constant-elastic function of the firm’s price p , the average \bar{p} of its rivals’ prices, its own advertising A , and the average of its rivals’ advertising \bar{A} , with elasticities $-\epsilon$, $\bar{\epsilon}$, α , and $-\bar{\alpha}$. The marginal cost of production is c and the cost of a unit of advertising is κ . Although customers pay p for each unit of output, the firm receives only p/f , where f is a product-market friction or wedge that depresses the price the firm receives. The factor f may be above or below 1. The firm’s objective is

$$\max_{p,A} \left(\frac{p}{f} - c \right) p^{-\epsilon} \bar{p}^{\bar{\epsilon}} A^{\alpha} \bar{A}^{-\bar{\alpha}} - \kappa A. \quad (1)$$

The profit-maximizing price is

$$p^* = \frac{\epsilon}{\epsilon - 1} f c \quad (2)$$

and in symmetric equilibrium, $\bar{p} = p$ and $\bar{A} = A$. For some reason—possibly price stickiness—the firm actually sets the price

$$p = m p^*. \quad (3)$$

The *profit-margin wedge*, m , may be above or below 1. If $m > 1$, the firm keeps the added profit per unit sold though it loses profit overall from the reduced volume. The reverse occurs if $m < 1$.

Equation (2) and equation (3) imply

$$p = m f \frac{\epsilon}{\epsilon - 1} c. \quad (4)$$

The variable part of the markup of price p over marginal cost c is the product of the two wedges, $m f$. The profit-margin wedge has implications stressed in Rotemberg and Woodford (1999) and is the way that sticky prices affect real allocations, as those authors explain. On the other hand, the wedge f also appears in equation (1), where it has the effect of taking away the margin increase from the firm, so an increase in f does not raise profit. Consequently, the two wedges have quite different effects. Later in the paper I will demonstrate that authors thinking they are measuring the profit-margin wedge m by studying labor's share of total cost are actually measuring the compound wedge $m f$, under the assumptions of this model.

2.1 Advertising

The first-order condition for advertising is

$$\frac{\alpha}{A} Q \left(\frac{p}{f} - c \right) = \kappa. \quad (5)$$

Rearranging and dividing both sides by p yields an expression for the ratio of advertising expenditure to revenue:

$$\frac{\kappa A}{p Q} = \alpha \frac{p/f - c}{p}. \quad (6)$$

Substituting for p from equation (3) and for p^* from equation (2) restates the right-hand side in terms of exogenous influences:

$$R = \frac{\kappa A}{p Q} = \alpha \frac{(m - 1)\epsilon + 1}{f m \epsilon} \quad (7)$$

Absent the special influences captured by f and m , that is, with $f = m = 1$, the advertising/revenue ratio is

$$R = \frac{\alpha}{\epsilon}, \quad (8)$$

a standard result in the advertising literature, first derived by Dorfman and Steiner (1954). See Bagwell (2007) for an impressively complete review of the literature on the economics of advertising.

From these equations, two useful results follow:

Proposition Rm: The elasticity of the advertising ratio R with respect to the profit-margin wedge m at the point $f = m = 1$ is $\epsilon - 1$.

Proposition Rf: The elasticity of the advertising ratio with respect to the wedge f is -1 .

Proposition Rm is the centerpiece of the paper—advertising is highly sensitive to the profit-margin wedge. If markups rise in a slump, firms should increase efforts aggressively to attract new customers and retain existing ones, because selling to them has become more profitable.

2.2 Labor share

The second measure is the labor share

$$\lambda = \frac{W}{pQ}. \quad (9)$$

Here W is the firm's total wage bill including all forms of compensation. Under the assumptions of Cobb-Douglas technology with labor elasticity γ and cost minimization, the wage bill is $\gamma c Q$, so

$$\lambda = \frac{\gamma c Q}{pQ} = \gamma \frac{\epsilon - 1}{\epsilon} \frac{1}{f m} \quad (10)$$

Two additional results then follow immediately:

Proposition λm : The elasticity of the labor share λ with respect to the profit-margin wedge m is -1 .

Proposition λf : The elasticity of the labor share with respect to the product-market wedge f is -1 .

2.3 Solving for the wedges

From the propositions above,

$$\log R = (\epsilon - 1) \log m - \log f + \mu_R \quad (11)$$

and

$$\log \lambda = -\log m - \log f + \mu_\lambda, \quad (12)$$

where μ_R and μ_λ are constant and slow-moving influences apart from m and f .

Solving this pair of equations for $\log m$ and $\log f$ yields

$$\log m = \frac{\log R - \log \lambda}{\epsilon} + \mu_m \quad (13)$$

and

$$\log f = -\log \lambda - \frac{\log R - \log \lambda}{\epsilon} + \mu_f. \quad (14)$$

Here μ_m and μ_f are constant and slow-moving influences derived in the obvious way from μ_R and μ_λ .

2.4 The role of cyclical movements

The main goal of this paper is to study movements of the observed time series R and λ and make inferences about the cyclical movements of the inferred wedges m and f , especially to quantify their contributions to the business cycle. Throughout the paper, I measure the business cycle by the employment rate, the fraction of the labor force holding jobs (one minus the unemployment rate). Variables are procyclical if they move positively with the employment rate and countercyclical if they move negatively. The data show that the advertising/sales ratio R and the labor share λ are quite procyclical. The expectation is that the wedges are both countercyclical—they measure forces that mediate reductions in employment when they rise.

3 Cyclical Movements of Advertising/GDP Ratio and the Labor Share

For many years, Robert J. Coen of the ad agency Erickson-McCann published a compilation of data on advertising expenditure. I was unable to find any surviving original copy of his data. Douglas Galbi posted a copy of Coen's estimates through 2007 in his blog, along with

estimates for early years from other sources. Galbi also provides links to Coen’s data sources, but the only one still active is for the data on newspapers. Further information about the sources appears in the appendix.

For 2005 through 2010, the Census Bureau has published revenue data for NAICS industry 51, the information sector, which includes the advertising industries. I define advertising as the sum of newspapers, magazines, broadcasting, and Internet. In the three years that the Census figures overlap Coen’s, the latter is 1.38 times the former. I take the figures for 2008 through 2010 to be this factor times the Census figure.

Figure 1 shows the ratio of advertising spending to private nominal GDP, stated as an index, with vertical bars shaded darker in years when the employment rate was low.

I form

$$R = \frac{\kappa A}{pQ} \quad (15)$$

as the ratio of advertising expenditure to private nominal GDP. Recall that the model for this variable is

$$\log R = (\epsilon - 1) \log m - \log f + \mu_R. \quad (16)$$

To characterize the cyclical behavior of R , I use the regression

$$\log R_t = \beta_R L_t + \mu_{R,t}. \quad (17)$$

Here the coefficient β_R multiplying the employment rate L_t measures the procyclicality of R . $\mu_{R,t}$ is a residual with substantial low-frequency movements unrelated to the business cycle. I model the low-frequency variation μ_R in a number of ways. The first is as a random walk:

$$\mu_{R,t} = \omega_R + \frac{1}{1 - \mathbb{L}} \eta_t, \quad (18)$$

where ω_R is a constant, \mathbb{L} is the lag operator, and $\eta_{R,t}$ is a white-noise innovation. The second is as a deterministic linear trend plus a highly serially correlated random component:

$$\mu_R = \omega_R + \tau_R t + \frac{1}{1 - \rho \mathbb{L}} \eta_t, \quad (19)$$

where τ_R is the linear trend. I omit the second specification if the estimated ρ exceeds one.

The third and fourth models take the slow-moving component to be a filtered version of R_t , say \bar{R}_t , and estimate

$$\log(R_t/\bar{R}_t) = \omega_R + \beta_R L_t + \eta_{R,t}. \quad (20)$$




Figure 1: Index of the Ratio of Advertising Spending to Private GDP, with Shading in Negative Proportion to the Employment Rate

<i>Left-hand variable</i>	<i>Coefficients and standard errors</i>					
	<i>Employment rate</i>		<i>Trend</i>		<i>Autoregressive term</i>	
Change in log of all types of advertising/GDP	0.97	(0.37)	-	-	-	-
Log of all types of advertising/GDP	0.97	(0.39)	0.000	(0.004)	0.93	(0.05)
HP cycle component	0.41	(0.23)	-	-	-	-
Bandpass cycle component	0.36	(0.16)	-	-	-	-
TV ads	0.40	(0.71)	0.014	(0.002)	0.71	(0.03)
Radio ads	0.37	(0.55)	0.008	(0.002)	0.81	(0.04)
Newspaper ads	1.19	(0.51)	-	-	-	-
Direct mail	-0.42	(0.49)	0.011	(0.008)	0.96	(0.04)
Other advertising	1.74	(0.54)	-0.011	(0.001)	0.77	(0.08)

Note: The right-hand variable in the first row is in first differences. Regressions include constants, not shown.

Table 1: Cyclical Behavior of the Advertising/GDP Ratio

The filter absorbs any trend in R_t , so I do not include an explicit trend. For the filters, I use Hodrick-Prescott with smoothing parameter 6.25 and a bandpass filter that passes cycle periods of 2 through 8 years.

The data run from 1950 through 2010. In addition to total advertising, I also estimate equations for the major components: TV, radio, newspapers, direct mail, and other. Table 1 shows the results. In all cases but one, the estimated cyclical response—the coefficient on the employment rate—is positive, generally statistically unambiguously positive. The advertising/GDP ratio is plainly procyclical—it rises when employment rises. The relationship is much stronger for the random-walk and AR specifications of the low-frequency component—the HP and bandpass filters plainly absorb a good deal of the variation in the advertising/GDP ratio that would otherwise be associated with employment.

I reiterate that the finding that advertising is procyclical is not central to the point of the paper. Rather, the key finding is that advertising is not strongly countercyclical, as it would be if the profit margin were countercyclical. The findings in Table 1 are not dispositive on the unimportance of fluctuations in the profit margin, however, because there is a possibility that a strongly procyclical effect from the product-market wedge is masking a countercyclical effect from the margin. To deal with this issue, I turn to a study of the cyclical movement of the labor share.

4 The Labor Share

The Bureau of Labor Statistics publishes a quarterly index of the labor share of non-financial corporations at bls.gov/lpc, series PRS88003173, starting in 1947. The limitation to corporations is desirable because there is no reliable basis for dividing proprietary income into labor and capital components. Figure 2 shows the data with bars shaded in proportion to the employment rate.

The labor share rises briefly but sharply at the beginning of a contraction, apparently because of labor hoarding. Similarly, it falls at the beginning of expansions, possibly because more intense work effort is the initial response to an increase in demand. To account for these transitory dynamics, I use the following specification to describe the cyclical movements of the labor share:

$$\log \lambda_t = \omega_\lambda + \beta_\lambda L_t + \phi \left[L_t - \frac{1}{N} (L_{t-1} + \dots + L_{t-N}) \right] + \mu_{\lambda,t}. \quad (21)$$

Here the coefficient β_λ on the employment rate L_t captures the cyclical behavior apart from the labor-hoarding aspect. ϕ is the coefficient of the term for labor-hoarding dynamics, ω_λ is a constant, and $\mu_{\lambda,t}$ is a slow-moving component of the labor share arising from other sources. I use the same four specifications for $\mu_{\lambda,t}$ as for the advertising/GDP ratio, R .

Table 2 shows the results of estimating the four specifications of the equation for the labor share. The transitory part of the specification uses $N = 4$ lagged quarters of the employment rate. The top row uses the random-walk specification and the second row the AR(1) specification of the slow-moving component. Both find that the labor share is procyclical, with a semi-elasticity of 0.70. That estimate is reasonably precise, with a standard error of 0.18. The estimate confirms the presence of a transitory component in the




Figure 2: Index of the Labor Share of Non-Financial Corporate Income, with Shading in Proportion to the Employment Rate

<i>Left-hand variable</i>	<i>Coefficients and standard errors</i>					
	<i>Employment rate</i>		<i>Departure of employment rate from recent past</i>		<i>Autoregressive term</i>	
Change in log labor share	0.70	(0.18)	-1.24	(0.17)	-	-
Log labor share	0.70	(0.18)	-1.23	(0.17)	0.957	(0.019)
HP cycle component	0.27	(0.06)	-1.39	(0.14)	-	-
Bandpass cycle component	0.27	(0.06)	-1.33	(0.13)	-	-

Note: The two right-hand variables in the first row are in first differences.

Regressions include constants, not shown

Table 2: Cyclical Behavior of the Labor Share

opposite direction. The immediate response of the share to an increase in employment—the difference between the first and second coefficient—is actually positive.

The third and fourth lines in Table 2 use filtering methods on the share to account for its slower movements. They find semi-elasticities about half as large. The reported standard error is probably an overstatement of the reliability, as it does not account for the preliminary filtering.

I conclude that both the advertising/GDP ratio and the labor share are procyclical. This finding points strongly in the direction of an influence of the product-market wedge that trumps any influence of the profit-margin wedge, given that the latter would stimulate a powerful countercyclical response of advertising.

5 The Values of the Wedges and Their Effects on Employment

The empirical counterparts of the two wedges, including their low-frequency components, are

$$\log M = \frac{\log R - \log \lambda}{\epsilon} \quad (22)$$

and

$$\log F = -\log \lambda - \log M. \tag{23}$$

Constructing these variables requires a value of the residual elasticity of demand ϵ . As I noted earlier, though market power is an important topic in many branches of applied microeconomics and is the subject of a large literature, the results of empirical research are inconclusive with respect to any single value for ϵ that would typify the aggregate economy. That said, most economists would probably place the typical value of the residual elasticity of demand in the range from 3 to 20, corresponding to profit margins of 33 down to 5 percent of price. I will present results for $\epsilon = 6$, which corresponds to a markup ratio of $\epsilon/(\epsilon - 1) = 1.2$, along with a discussion of results for lower and higher amounts of market power.

Figure 3 and Figure 4 show the values of the two wedges again shaded in proportion to the employment rate. The darker the vertical bars, the lower the *level* of employment (in contrast to recession bars, which would show *declines* in employment). The basic findings of the paper are plainly visible in these figures. Equation (22) shows that the calculated value of the profit-margin wedge subtracts the log of the labor share from the log of the advertising/GDP ratio, which removes the common element of the two. That element is the log of the product-market wedge, $\log F$. What is left has almost no cyclical movement. On the other hand, $\log F$, calculated from equation (23) by using the calculated value of $\log M$ to remove the influence of that wedge from the labor share, is countercyclical: Times of low employment coincide with high values of F . Recall that the coefficients in the middle panel of Table 2 imply a transitory increase in the labor share in an employment contraction and thus a transitory decline in F . As soon as this dissipates, the higher level of F that prevails in times of low employment becomes visible.

5.1 Role of the two wedges in employment volatility

The main goal of this research is to quantify the contributions of $\log m$ and $\log f$ to the movements of the employment rate. A three-way breakdown is

$$L_t = -\theta \log m_t - \delta \log f_t + x_t, \tag{24}$$

where x_t captures all the other influences on employment. The coefficients θ and δ are presumptively positive, because both wedges discourage employment. This equation is not a regression with x_t playing the role of the disturbance, because x_t is surely correlated with




Figure 3: Calculated Time Series for the Profit-Margin Wedge, Including Low-Frequency Component




Figure 4: Calculated Time Series for the Product-Market Wedge, Including Low-Frequency Component

$\log m_t$ and $\log f_t$. But with outside information about the coefficients θ and δ , it is possible to decompose the movements of L_t into those attributable to each of the three components on the right-hand side. The econometric issue of identification does not arise here, because no coefficients are estimated.

Current macroeconomic theory characterizes the effects of aggregate driving forces in terms of wedges, notably m , which plays a key role in the New Keynesian model's transmission mechanisms to account for cyclical movements in employment and aggregate activity, as explained, for example, in Rotemberg and Woodford (1999). Wedges are intermediate variables, not exogenous driving forces, in most macro models.

The first helpful insight from macro theory is that the two coefficients θ and δ should have essentially the same value, say θ . Theory suggests that all wedges combine to generate a single master wedge separating the marginal product of labor from the marginal value of time. The producer's contribution to the wedge is the ratio of the price paid by the consumer to the producer's cost. From equation (4), the ratio is

$$mf \frac{\epsilon}{\epsilon - 1}. \quad (25)$$

The two variables m and f enter with equal elasticities of one.

Second, Hall (2009) suggests that the employment rate responds to the master wedge with a semi-elasticity of somewhat more than 1 in absolute value. I take $\theta = 1$ as the main case, but examine the consequences of lower and higher values.

The next step is to measure the contributions of $\theta \log m_t$, $\theta \log f_t$, and x_t to the movements of the employment rate. The two measured series $\theta \log M_t$ and $\theta \log F_t$ contain slower-moving components, attributable to forces such as drift in the elasticity α of demand with respect to advertising and drift in the elasticity γ of output with respect to labor. These components are prominent in Figure 3 and Figure 4. The paper is about cyclical movements. To focus on these, I use a combination of two tools. The first, as in the earlier sections on the cyclical movements of R and λ , is to apply a filter to the data that attenuates the lower frequencies. The filter involves quasi-differences, such as

$$\hat{m}_t = \log M_t - \rho \log M_{t-1}. \quad (26)$$

For $\rho = 1$, the filter that is most aggressive in isolating higher frequencies, the gain is zero at the lowest frequency. After filtering, the relation among the variables is

$$\hat{L}_t = -\theta \hat{m}_t - \theta \hat{f}_t + \hat{x}_t. \quad (27)$$

The second tool defines the cycle in terms of the movements of the employment rate, an idea that runs through the entire paper and receives its sharp definition here. Rather than looking at variances of the components, which treat all movements equally, I look at covariances with the employment rate. Covariances filter out movements not related to the cycle, given my definition that equates the cycle to movements in employment. An added benefit of this approach is that covariances are additive. The decomposition is

$$V(\hat{L}_t) = -\theta \text{Cov}(\hat{m}_t, \hat{L}_t) - \theta \text{Cov}(\hat{f}_t, \hat{L}_t) + \text{Cov}(\hat{x}_t, \hat{L}_t). \quad (28)$$

Finally, I divide by the variance of the employment rate to get

$$1 = \theta \frac{-\text{Cov}(\hat{m}_t, \hat{L}_t)}{V(\hat{L}_t)} + \theta \frac{-\text{Cov}(\hat{f}_t, \hat{L}_t)}{V(\hat{L}_t)} + \frac{\text{Cov}(\hat{x}_t, \hat{L}_t)}{V(\hat{L}_t)}. \quad (29)$$

Note that this can be written more compactly as

$$1 = \theta \beta_m + \theta \beta_f + \beta_x, \quad (30)$$

where the β s are the coefficients of the negatives of the corresponding variables regressed on \hat{L}_t . Earlier in the paper I mentioned that the regression of the log of the advertising/GDP ratio divided by the labor share on employment was highly informative about the cyclical role of the profit-margin wedge, m . The coefficient, β_m , is exactly the coefficient from that regression with the sign reversed (so that its presumed value is positive). Similarly, β_f is the regression of the negative of the product-market wedge on the employment rate, again with a positive presumed value. The quantity $\theta \beta_m$ is the contribution of the profit-margin wedge; $\theta \beta_f$ is the contribution of the product-market wedge; and the remainder, $1 - \theta \beta_m - \theta \beta_f$ is contribution of the residual. Because θ is positive and the two coefficients are presumptively positive, the contributions are presumptively positive (I say presumptively because it turns out that β_m is actually, paradoxically, slightly but not significantly negative).

Earlier I estimated the coefficients of the filtered values of the advertising/GDP ratio R_t and the labor share λ_t on similarly filtered employment. These are β_R and β_λ . Because regression is a linear operation and the coefficients β_m and β_f are linear functions of those coefficients, the desired coefficients can be calculated directly from the earlier values:

$$\beta_m = \frac{\beta_\lambda - \beta_R}{\epsilon} \quad (31)$$

and

$$\beta_f = \frac{\beta_R - (\epsilon - 1)\beta_\lambda}{\epsilon}, \quad (32)$$

These are, for the first-difference specification ($\rho = 1$) and $\theta = 1$,

$$\beta_m = \frac{0.70 - 0.97}{6} = -0.04 \quad (33)$$

and

$$\beta_f = \frac{0.97 + 5 \times 0.70}{6} = 0.75. \quad (34)$$

The profit-margin wedge m is slightly *procyclical*, contrary to expectation, while the product-market wedge f is robustly *countercyclical*, as expected.

Calculation of the standard errors of these estimates requires the covariance of the estimates of β_R and β_λ . Because the first uses annual data and the second quarterly data, a bit of an econometric issue arises in calculating the covariance, as described in the appendix. The resulting standard errors are 0.07 and 0.17. Thus the role of the profit-margin wedge m is tightly circumscribed around zero and the role of the product-market wedge is unambiguously substantially *countercyclical*.

The calculations above depend on the parameter θ , the effect of wedges in general on the employment rate. Figure 5 shows how the calculation of the contributions depends on that parameter, using values of \hat{m}_t and \hat{f}_t calculated with $\epsilon = 6$ and $\rho = 1$. The horizontal axis is the effect of the wedge on employment, θ . At $\theta = 1$, the profit-margin wedge m accounts for -4 percent of the cyclical movements of the employment rate, the product-market wedge for 75 percent, and the other forces for the remaining 29 percent. For higher values of θ , the product-market wedge accounts for an implausibly high fraction of the cyclical movements.

Table 3 reports the sensitivity of the results to two other determinants, the residual elasticity of demand, ϵ , and the quasi-differencing coefficient, ρ . The base case described above is enclosed in a box. Removing the quasi-difference filtering completely, shown in the first row, results in an estimate of the procyclicality of the advertising/GDP ratio, β_R , of about half its value in the base case, but still positive. However, its confidence interval is wide and includes *countercyclical* values. I use the term *procyclical* loosely here because in the absence of any quasi-differencing, low frequencies have as big a role in the estimate as do cyclical frequencies. The labor share is cyclically neutral with zero quasi-differencing. The rest of the results in the first row of the table show that, without concentrating on cyclical frequencies, the procedure assigns small *procyclical* roles to the profit-margin wedge and small *countercyclical* roles to the product-market wedge. Most of the movements of the employment rate are assigned to the residual when the large low-frequency movements are included in the measured wedges.

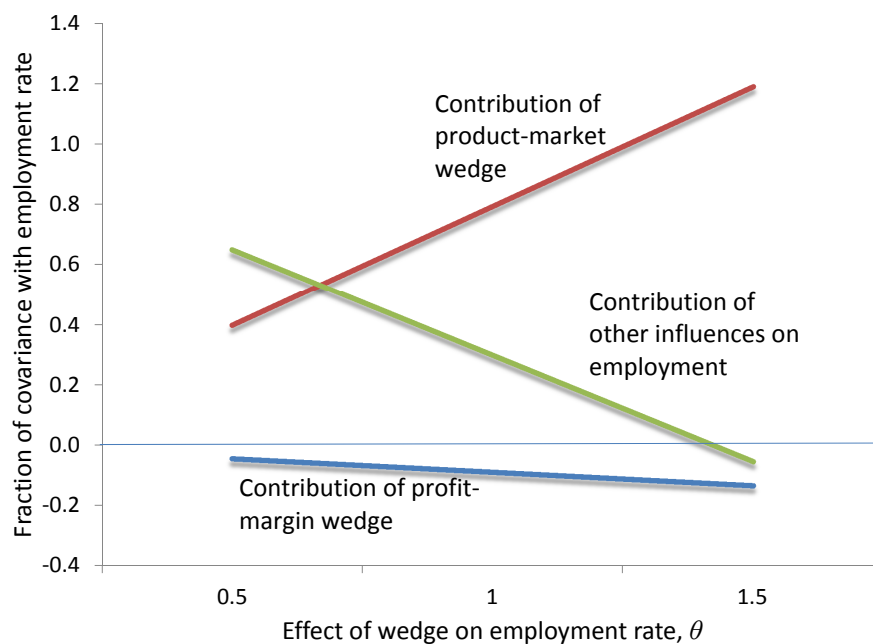


Figure 5: Contributions of Wedges to Employment Movements as Functions of the Parameter θ

ρ , annual quasi-difference coefficient	Estimated coefficients		Implied contributions of wedges to cyclical movements in the employment rate					
			ε , residual elasticity of demand					
	β_R	β_λ	3		6		12	
			β_m	β_f	β_m	β_f	β_m	β_f
0	0.46 (0.73)	0.06 (0.10)	-0.13 (0.25)	0.20 (0.25)	-0.07 (0.12)	0.13 (0.15)	-0.03 (0.06)	0.10 (0.11)
0.5	0.74 (0.57)	0.45 (0.14)	-0.09 (0.20)	0.55 (0.21)	-0.05 (0.10)	0.50 (0.15)	-0.02 (0.05)	0.48 (0.14)
1	0.97 (0.37)	0.70 (0.18)	-0.09 (0.13)	0.79 (0.18)	-0.04 (0.07)	0.75 (0.17)	-0.02 (0.03)	0.73 (0.17)

Note: the figure in the first column is the annual coefficient; the quarterly coefficient is that figure raised to the power 1/4.

Table 3: Implications of Alternative Quasi-Differencing Coefficients and Residual Elasticities of Demand

The bottom row of Table 3 shows that the basic message of the paper—the strong countercyclical role of the product-market wedge and the weak (and counterintuitive) role of the profit-margin wedge—holds for the wide range of values of the residual elasticity ϵ .

In no case does the profit-margin wedge make any substantial contribution. With $\epsilon = 6$, $\theta = 1$, and $\rho = 1$, the point estimate of the contribution is $\beta_m = -0.04$ with a standard error of 0.07. With confidence probability 0.75, the results show that the contribution is zero or below. The confidence probabilities are 0.92 that $\beta_m \leq 0.05$ and 0.98 that $\beta_m \leq 0.1$. The evidence against any meaningful role for the profit-margin wedge is quite strong. Results for other values of the residual elasticity ϵ are equally strong. Without filtering in favor of cyclical frequencies ($\rho = 0$), the point estimate is the same (slightly negative) but the standard error is 0.13, so the confidence probabilities are not as strong evidence against a small positive value of β_m .

On the other hand, the evidence in favor of positive values of the contribution β_f of the product-market wedge is strong in all cases. The point estimate with $\epsilon = 6$, $\theta = 1$, and $\rho = 1$ is 0.75 with a standard error of 0.17. The confidence probability that $\beta_f \geq 0.4$ is 0.98 and that $\beta_f \geq 0.6$ is 0.81. As I noted earlier, this evidence depends on the restriction to cyclical frequencies.

5.2 Potential overstatement of the procyclicality of the labor share

The procyclical behavior of the labor share, as demonstrated in this paper and in a good deal of earlier work, has an important role in the conclusions of the paper. The two key equations mapping the observed cyclical behavior of advertising and the labor share into the underlying relations described by β_m and β_f can be written

$$\beta_m = \frac{\beta_\lambda - \beta_R}{\epsilon} \tag{35}$$

and

$$\beta_f = \beta_\lambda - \beta_m. \tag{36}$$

In the first equation, a substantially procyclical labor share (positive value of β_λ) mostly offsets the opposite role of the highly procyclical advertising ratio, R , and, in any case, the difference is divided by ϵ , resulting in a modest amount of cyclical variation in m , as reflected in a small value of β_m , suggesting that the profit-margin effect is slightly procyclical, contrary to its expected countercyclical effect. In the second equation, because β_m is small,

the inferred value of the cyclical coefficient for the product-market wedge f , β_f , is close to the coefficient measuring the procyclicality of the labor share, β_λ , which is big. The bottom line is a small role for the profit-margin wedge m and a big countercyclical role for the product-market wedge f .

Nekarda and Ramey (2010) and Nekarda and Ramey (2011) challenge the finding of a procyclical labor share. They conclude that the labor share, properly measured, is constant over the cycle. These authors, along with most of their predecessors, frame the question as the cyclical behavior of the markup ratio, so that it might appear that they estimate the cyclical movement of the profit margin wedge, β_m , rather than the cyclical movements captured by β_λ , in the notation of this paper. However, the earlier literature does not consider the cyclical role of the product-market wedge f introduced in this paper. Hence it is appropriate to consider their results as bearing on the labor share rather than the profit margin. Equation (12) in Nekarda and Ramey (2010) describes the relation between the labor share and the markup ratio. The variable considered in their work and in much of the earlier work is the reciprocal of the labor share, adjusted for overtime wages and for the elasticity of substitution. The latter adjustment has an effect only if the technology is not Cobb-Douglas. Their equation (13) is close to the one in this paper, except that the cycle indicator is not based on the employment rate.

If Nekarda and Ramey are correct that $\beta_\lambda = 0$, the equations above become

$$\beta_m = \frac{-\beta_R}{\epsilon} \quad (37)$$

and

$$\beta_f = -\beta_m. \quad (38)$$

With $\beta_R = 0.97$ as in the base case and $\epsilon = 6$, the values are $\beta_m = -0.16$ and $\beta_f = 0.16$, both with standard errors of 0.06. The basic conclusions of the paper are qualitatively the same—the implied role for the profit-margin wedge is paradoxically procyclical rather than countercyclical, as wedges are generally thought to be in general. The implied role for the product-market wedge f is countercyclical, as it should be, but quite a bit smaller than in the base case. The conclusion that the profit-margin wedge does not have a countercyclical role remains, but now it is statistically unambiguously *procyclical*, which fits with the model of Kaplan and Menzio (2013) but not with standard macro models with wedges. The implied role for product-market wedge is quite a bit smaller than before, though still

statistically unambiguous. The calculation confirms that resolution of the debate over the cyclical movements of the properly measured labor share is a top priority in understanding cyclical fluctuations.

5.3 Conclusion about the role of profit-margin wedge and product-market frictions

The findings point in the direction that β_m is close to zero. The suggestion that the data do not support countercyclical profit margins is not new. But the companion finding is new—that the data strongly support the hypothesis that firms encounter some kind of friction during slumps that makes them behave as if they were paying a higher tax on their output.

6 Advertising as Investment

One potential reason that advertising is procyclical is that advertising may be an investment that has a lasting benefit, extending beyond the period of the expenditure itself. Other types of investment—inventories and plant-equipment in particular—are quite procyclical, in the sense used in this paper, that the ratio of expenditure to revenue declines in a recession. In this section, I investigate whether the investment character of advertising is concealing a countercyclical movement of advertising spending driven by countercyclical margins.

The empirical literature on the effects of advertising has reached a reasonably strong consensus that most of the effect of advertising on sales occurs within a year—see Bagwell (2007), pages 1726 to 1728, for cites, and Corrado, Hulten and Sichel (2009) for a recent quantification and additional cites. The latter paper places the annual depreciation rate of advertising capital at 60 percent.

Nerlove and Arrow (1962) developed the theory of investment in depreciable advertising along the same lines as Jorgenson’s (1963) famous model of investment in plant and equipment. The stock of advertising, A_t , evolves according to

$$A_t = a_t + (1 - \delta)A_{t-1}. \tag{39}$$

Here a_t is purchases of new advertising and δ is the rate of depreciation.

The price of a unit of advertising—now thought of as the services of a unit of a stock of advertising over one year—is

$$\kappa_t = \frac{r + \delta}{1 + r} p_t. \tag{40}$$

	β_R	β_λ	β_m	β_f
Base case	0.97 (0.37)	0.70 (0.18)	-0.04 (0.07)	0.75 (0.17)
Capital	0.21 (0.29)	0.70 (0.18)	0.08 (0.07)	0.62 (0.17)

Table 4: Comparison of Results for the Base Case and an Alternative with Advertising Treated as Capital

Here r is the annual real interest rate, δ is the annual rate of depreciation, and p is the price of investment in advertising, which I take to be the price of goods and services in general. Notice that this formula is $\kappa_t = p_t$ if there is complete depreciation within a year: $\delta = 1$.

The ratio of advertising expenditure to revenue is $\frac{\kappa_t A_t}{p_t Q_t}$. To form this ratio from the data, I calculate the advertising stock A_t from the recursion, equation (39), using data on a_t formed as the Coen data on advertising expenditure deflated by the price index for private GDP, starting in 1929. I calculate κ using $\delta = 0.6$ and $r = 0.05$.

Table 4 compares the results from treating advertising as investment with those shown earlier in the paper, in the base case ($\epsilon = 6$ and $\rho = 1$). The value of the cyclicity coefficient for the advertising/GDP ratio is considerably lower (about 0.2 instead of nearly 1), confirming that some of the procyclicality of the ratio arises from investment factors. But the ratio remains procyclical, not countercyclical as it would be if the profit-margin wedge were an important factor in its movements. The cyclicity coefficient β_λ remains the same, of course. The contributions of the two wedges are not changed much. The one for the profit-margin wedge, β_m , takes on the right positive sign, but the coefficient is small and sufficiently accurately estimated to rule out large negative values. The one for the product-market wedge, β_f , is a little smaller than in the base case, but remains robustly large. The basic finding of the paper—that the profit-margin wedge is unimportant and the product-market wedge is quite important—easily survives the switch to the investment approach. This conclusion does rest on the depreciation rate of 0.6.

7 Other Influences

To this point, the paper has compared two forces that affect the advertising/sales ratio R and the labor share λ . These are the profit-margin wedge m and the product-market wedge

f . The evidence favors a small role for the margin wedge and a large role for the product-market wedge. A rise in the product-market wedge during slumps explains both the decline in R during slumps and the decline in λ . A natural question is whether other influences could have the same effect. In this section I argue that the set of other influences is quite limited.

To consider the effects of other influences, I extend the model to include the following:

- A Hicks-neutral productivity index, h
- A labor wedge or measurement error, f_L
- A capital wedge or measurement error, f_K
- An advertising wedge or measurement error, f_A

For clarity, I refer to the product-market wedge f as f_Q in this section. For the three new f wedges, I assume that the firm pays an amount per unit that is the wedge times the reported price; for example, the firm pays an actual wage bill of $f_L W$ when the reported wage bill is W . Marginal cost c is now a function of h , f_L and f_K . But in the derivation of the advertising spending/revenue ratio R , leading up to equation (7), neither c nor its determinants make their way into R . The new wedge f_A does affect the ratio in the extended model:

$$R = \frac{\kappa A}{pQ} = \frac{\alpha}{f_A f_Q m} \frac{(m-1)\epsilon + 1}{\epsilon} \quad (41)$$

From the derivation of equation (10), it is apparent that only the labor wedge f_L enters the formula for the labor share λ :

$$\lambda = \frac{W}{pQ} = \frac{1}{f_L f_Q m} \gamma \frac{\epsilon - 1}{\epsilon} \quad (42)$$

The appendix contains complete derivations for R and λ .

These conclusions follow:

- The Hicks-neutral productivity index h and the capital wedge or measurement error f_K affect neither the advertising/sales ratio R nor the labor share λ .
- The new wedge f_A affects R with an elasticity of -1 and the new wedge f_L affects λ with an elasticity of -1 ; the margin wedge m remains the only wedge that has a high elasticity.

- The advertising wedge or measurement error, f_A , lowers R in the same way that f_Q does.
- The labor wedge or measurement error, f_L , lowers λ in the same way that f_Q does.
- Equal values of f_A and f_L have the same effect as f_Q of the same value.

The relations between the estimated coefficients β_R and β_λ and the underlying coefficients describing the negatives of the cyclical responses of the profit-margin wedge (β_m), the product-market wedge (β_f), the advertising wedge (β_A), and the labor wedge (β_L), become:

$$\beta_R = \text{Semi-elasticity of } R \text{ with respect to } L = -(\epsilon - 1) \beta_m + \beta_f + \beta_A \quad (43)$$

$$\beta_\lambda = \text{Semi-elasticity of } \lambda \text{ with respect to } L = \beta_m + \beta_f + \beta_L \quad (44)$$

With four unknown coefficients and two equations, none of the unknowns is identified. Conditional on values of the new coefficients β_A and β_L , the values of the coefficients for the wedges studied earlier are:

$$\beta_m = \frac{\beta_\lambda - \beta_R}{\epsilon} + \frac{\beta_A - \beta_L}{\epsilon} \quad (45)$$

and

$$\beta_f = \frac{\beta_R + (\epsilon - 1)\beta_\lambda}{\epsilon} - \frac{\beta_A + (\epsilon - 1)\beta_L}{\epsilon}. \quad (46)$$

I let $\hat{\beta}_m$ and $\hat{\beta}_f$ be the estimates discussed earlier in the paper, computed from the first term in each of the equations above. The bias in these estimates in the presence of the new wedges is revealed by rewriting the equations as

$$\hat{\beta}_m = \beta_m - \frac{\beta_A - \beta_L}{\epsilon} \quad (47)$$

and

$$\hat{\beta}_f = \beta_f + \frac{\beta_A + (\epsilon - 1)\beta_L}{\epsilon}. \quad (48)$$

In both equations, the bias from the advertising wedge is attenuated by division by the residual elasticity of demand ϵ . The same is true of the labor wedge in the first equation. The key conclusion of the paper is that the contribution of the profit-margin wedge, β_m , is small and slightly procyclical (negative), rather than large and countercyclical (positive) as in some macro models. Only a very large countercyclical coefficient for the advertising wedge (big positive β_A) or a very procyclical coefficient for the labor wedge (big negative

β_L) would have much effect in concealing a large role for the profit-margin wedge through the bias term in the first equation.

A countercyclical advertising wedge would imply that the properly measured ratio of advertising spending to sales was less procyclical or even countercyclical, which would raise the implied value of the contribution of the profit-margin wedge β_m . Nothing comes to mind that would suggest such a phenomenon. If the countercyclical response β_A is less than 1.0, removing the bias would add at most 0.2 to the contribution of the profit-margin wedge to fluctuations in employment.

A procyclical labor wedge (negative value of β_L) is a more likely source of bias in my estimate of the effect, β_m , of the profit-margin wedge. Bils (1987) introduced the idea that mismeasurement of wages resulted in an understatement of the procyclical behavior of the labor share that a countercyclical profit-margin wedge m would induce. Rotemberg and Woodford (1999) promoted Bils's findings. The idea is that the marginal wage rises relative to the average wage because of overtime pay. The marginal wage is the appropriate concept for calculating the labor share. Again, I note that Nekarda and Ramey challenge the magnitude of Bils's overtime adjustment. In any case, as equation (47) shows, the adjustment is divided by ϵ in the framework of this paper, so the effect of such an adjustment on the key measure, β_m , is necessarily small in the framework of this paper.

Bils, Rotemberg-Woodford, and Nekarda-Ramey start with the finding that the standard measure of the labor share, before overtime adjustment, is cyclically neutral. To support the hypothesis that the procyclical behavior of the labor share indicated a countercyclical profit-margin wedge, Bils introduced the wage adjustment. By contrast, in this paper, the labor share is quite procyclical even without any adjustment of wages. The reason is that I measure the cyclical movements of the share after considering the transitory effects of labor hoarding—see Table 2. The substantial role for the transitory countercyclical element of the movements of the labor share demonstrates the need for further work on the cyclical properties of the labor share. However, because the division by ϵ in equation (47), only very large cyclical corrections of the labor share would shake the conclusion of this paper that β_m is quite small.

I conclude that the extended model is not successful in identifying a plausible source of variation that overcomes the high positive elasticity of advertising with respect to the profit margin. Imputing a substantial countercyclical error in measuring the price of advertising is

implausible. A countercyclical error in measuring the labor share may be plausible but so is a procyclical error associated with labor hoarding.

The second major conclusion of the paper is that the contribution of the product-market wedge, β_f , is large. This conclusion is more sensitive to consideration of additional wedges, notably the labor wedge. In equation (48), β_L enters with a coefficient $(\epsilon - 1)/\epsilon$, which is close to one. Earlier I noted that accepting the proposition from Nekarda and Ramey that the correct value of β_λ is zero, because of a countercyclical measurement error with $\beta_L = 0.7$, implies that the contribution of the product-market wedge is only 0.16, far below the value of 0.75 in the base case in this paper.

8 Concluding Remarks

Macroeconomists have become fond of invoking rising profit margins to explain the puzzle of slumps. I think it will be important for these model-builders to bring the behavior of advertising into the variables under consideration. It's hard to overcome the implication that advertising is really, really sensitive to profit margins. The obvious conclusion from the failure of advertising to explode in recessions is that profit margins remain about the same when the economy contracts.

If so, we need to redouble efforts to track down the sources of poor economic performance in long-lasting slumps. Sticky wages—a topic unrelated to any of the results in this paper—may be one important source. But the results of this paper suggest that product-market frictions or wedges deserve further investigation as well.

References

- Bagwell, Kyle, “The Economic Analysis of Advertising,” in M. Armstrong and R. Porter, eds., *Handbook of Industrial Organization, Volume 3*, Elsevier, 2007, chapter 28, pp. 1701–1844.
- Bils, Mark, “The Cyclical Behavior of Marginal Cost and Price,” *American Economic Review*, December 1987, *77* (5), 838–855.
- , “Pricing in a Customer Market,” *Quarterly Journal of Economics*, November 1989, *104* (4), 699–718.
- and James Kahn, “What Inventory Behavior Tells Us about Business Cycles,” *American Economic Review*, June 2000, *90* (3), 458–481.
- Blank, David M., “Cyclical Behavior of National Advertising,” *The Journal of Business*, 1962, *35* (1), 14–27.
- Borden, Neil Hopper, *The Economic Effects of Advertising*, Chicago: Irwin, 1942.
- Bresnahan, Timothy F., “Empirical Studies of Industries with Market Power,” in Richard Schmalensee and Robert D. Willig, eds., *Handbook of Industrial Organization, Volume II*, Amsterdam: North-Holland, 1989, chapter 17, pp. 1011–1057.
- Chevalier, Judith A. and David S. Scharfstein, “Capital-Market Imperfections and Countercyclical Markups: Theory and Evidence,” *American Economic Review*, 1996, *86* (4), 703–725.
- Corrado, Carol, Charles Hulten, and Daniel Sichel, “Intangible Capital and U.S. Economic Growth,” *Review of Income and Wealth*, 2009, *55* (3), 661–685.
- Dorfman, Robert and Peter O. Steiner, “Optimal Advertising and Optimal Quality,” *American Economic Review*, 1954, *44* (5), 826–836.
- Edmond, Chris and Laura Veldkamp, “Income Dispersion and Counter-Cyclical Markups,” *Journal of Monetary Economics*, 2009, *56*, 791–804.

- Hall, Robert E., "Reconciling Cyclical Movements in the Marginal Value of Time and the Marginal Product of Labor," *Journal of Political Economy*, April 2009, 117 (2), 281–323.
- Hausman, Jerry, Gregory Leonard, and J. Douglas Zona, "Competitive Analysis with Differentiated Products," *Annals of Economics and Statistics / Annales d'conomie et de Statistique*, 1994, (34), 159–180.
- Jorgenson, Dale W., "Capital Theory and Investment Behavior," *American Economic Review*, 1963, 53 (2), 247–259.
- Kaldor, Nicholas, "The Economic Aspects of Advertising," *The Review of Economic Studies*, 1950, 18 (1), 1–27.
- Kaplan, Greg and Guido Menzio, "Shopping Externalities and Self-Fulfilling Unemployment Fluctuations," Working Paper 18777, National Bureau of Economic Research February 2013.
- Loecker, Jan De and Frederic Warzynski, "Markups and Firm-Level Export Status," *American Economic Review*, May 2012, 102 (6), 2437–71.
- Nekarda, Christopher J. and Valerie A. Ramey, "The Cyclical Behavior of the Price-Cost Markup," June 2010. University of California, San Diego.
- and — , "Industry Evidence on the Effects of Government Spending," *American Economic Journal: Macroeconomics*, January 2011, 3 (1), 36–59.
- Nerlove, Marc and Kenneth J. Arrow, "Optimal Advertising Policy under Dynamic Conditions," *Economica*, 1962, 29 (114), 129–142.
- Rotemberg, Julio J. and Michael Woodford, "The Cyclical Behavior of Prices and Costs," in John Taylor and Michael Woodford, eds., *Handbook of Macroeconomics, Volume 1B*, Amsterdam: North-Holland, 1999, chapter 16, pp. 1051–1135.
- Simon, Julian L., *Issues in the Economics of Advertising*, Urbana: University of Illinois Press, 1970.
- Yang, Charles Yneu, "Variations in the Cyclical Behavior of Advertising," *Journal of Marketing*, 1964, 28 (2), 25–30.

A Full Derivation for Advertising

Optimal price:

$$p^* = \frac{\epsilon}{\epsilon - 1} f_Q c \quad (49)$$

Actual price:

$$p = mp^* \quad (50)$$

First-order condition for advertising:

$$\frac{\alpha}{A} Q \left(\frac{p}{f_Q} - c \right) = f_A \kappa \quad (51)$$

From above,

$$c = \frac{\epsilon - 1}{\epsilon} \frac{p}{mf_Q} \quad (52)$$

The first-order condition becomes

$$\frac{\alpha}{A} \frac{pQ}{f_Q} \frac{(m-1)\epsilon + 1}{\epsilon m} = f_A \kappa \quad (53)$$

and, finally,

$$R = \frac{\kappa A}{pQ} = \frac{\alpha}{f_A f_Q m} \frac{(m-1)\epsilon + 1}{\epsilon} \quad (54)$$

B Full Derivation for Labor Share

With cost minimization and Cobb-Douglas technology, labor cost is a fixed share γ of total cost:

$$\frac{f_L W}{cQ} = \gamma \quad (55)$$

Substitute for c :

$$f_L f_A m \frac{\epsilon}{\epsilon - 1} \frac{W}{pQ} = \gamma \quad (56)$$

so

$$\lambda = \frac{W}{pQ} = \frac{1}{f_L f_Q m} \gamma \frac{\epsilon - 1}{\epsilon} \quad (57)$$

C Data Sources

Table 5 gives further information about sources of advertising data.

<i>Type of advertising</i>	<i>Source</i>	<i>Years</i>	<i>Reference</i>
Newspapers	Newspaper Association of America	1950-2011	http://www.naa.org/Trends-and-Numbers/Advertising-Expenditures/Annual-All-Categories.aspx
All types	Douglas Galbi	1900-2007	http://purplemotes.net/2008/09/14/us-advertising-expenditure-data/
All types	Census Bureau	1994-2004	http://www.census.gov/prod/www/abs/bus-services.html
		2005-2010	http://www.census.gov/services/index.html

Table 5: Sources for Advertising Data

D Covariance of Estimates of β_R and β_λ

The covariance matrix of the vectors of estimates, say $\hat{\theta}_R$ and $\hat{\theta}_\lambda$, is

$$\mathbb{E}_{X_R, X_\lambda} [(\hat{\theta}_R - \theta_R)(\hat{\theta}_\lambda - \theta_\lambda)'] = (X_R' X_R)^{-1} X_R' \mathbb{E} (\eta_R \eta_\lambda') X_\lambda (X_\lambda' X_\lambda)^{-1} \quad (58)$$

I assume that the covariance matrix of the disturbances η_R and η_λ is $I \otimes (\sigma_{R,\lambda} \iota)$, where I is a $T_R \times T_R$ identity matrix, \otimes is the Kronecker product, $\sigma_{R,\lambda}$ is the covariance of $\eta_{R,t}$ and any $\eta_{\lambda,t'}$ with t' in the same year as t , and ι is a 1×4 vector of ones. In words, both disturbances are serially uncorrelated and have only a contemporaneous cross-correlation, which is equal across the 4 quarters of the year. Under this assumption, the covariance matrix of the estimates is

$$\sigma_{R,\lambda} (X_R' X_R)^{-1} X_R' \bar{X}_\lambda (X_\lambda' X_\lambda)^{-1}, \quad (59)$$

where \bar{X}_λ is a matrix of annual sums of the variables in X_λ . The desired covariance of β_R and β_λ is the appropriate element of this matrix.