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

April 23, 2014

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. An influence that fits some of 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

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 definitely not countercyclical. I show that the ratio of advertising to GDP remains constant in a year when employment remains constant and *falls* by about one percent for each percentage fall in employment in the previous year. 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 the advertising/GDP ratio should be neither procyclical (as it actually is) nor countercyclical (as the widening-profit-margin model implies). The baseline

model includes 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. Consider the ratio of the advertising/sales variable to the labor share. A regression of that ratio 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.

The paper follows an important branch of the advertising literature, launched by Nerlove and Arrow (1962), by treating 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 the results show 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 consider a number of potential variations around the basic specification in the paper. I use current-year and past-year employment as alternative measures of the business cycle. Both the advertising/GDP ratio and the labor share are essentially uncorrelated with current-year employment and quite positively and equally correlated with past-year employment. As a result, the basic finding based on the ratio of the two variables is the same whether using current-year employment or past-year. For the base case, I use a residual elasticity of 6, corresponding to a markup ratio of 1.2, but I show that the results are essentially unchanged for elasticities of 3 and 12, which span the reasonable range for the parameter as a description of the average for the U.S. economy. The results use data filtered to eliminate longer-run movements of the two key series. I use two filters that accomplish this separation. And I consider variations from the base-case depreciation rate.

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 misstating the depreciation rate of advertising, a topic I consider separately with negative conclusions.

The basic finding of this paper is unfavorable for the standard sticky-price macro model. As Rotemberg and Woodford (1999) explain, that model implies rising profit margins in recessions, as prices remain at pre-recession levels while costs decline. Sellers who perceive an inability to cut prices to profit-maximizing levels ought to use other tools to offset the decline in profit. Advertising is among those tools. The failure of advertising to rise when output falls suggests that recessions involve a more complicated process than the sticky-price model contemplates. In that process, sellers do not perceive a benefit from expanding sales

4

by cutting prices or by increasing advertising. Prices and advertising are unresponsive to the decline in output because change is unremunerative, not because sellers' hands are tied.

1 Related Research

1.1 Cyclical behavior of advertising

Advertising theory implies that the ratio of advertising spending to revenue is the logical variable for the purposes of this paper. Accordingly, the investigation here focuses on the cyclical behavior of that ratio. Most past research on the cyclical behavior of advertising has not examined the ratio of spending to revenue, but rather studies spending itself. Thus previous findings of procyclical advertising do not give a direct measure of the cyclical properties of the advertising/revenue ratio.

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. Bils (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 Bils interprets the model as pointing toward procyclical advertising.

Molinari and Turino (2009) document the strong positive correlation of advertising and GDP in the United States. They build a dynamic general-equilibrium model that includes advertising. Firms advertise to shift their demands outward. The model includes an exogenous process of variations in the residual elasticity of demand facing sellers. Advertising has a lasting effect modeled as in Nerlove and Arrow (1962). In the model, advertising responds positively to a markup shock. The main point of the paper is that advertising can amplify the response of key macro variables to driving forces.

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 Evidence on the sensitivity of advertising spending to the profit margin

Gurun, Matvos and Serub (2013) study advertising volume for subprime mortgages. They find large variation across geographic markets in profit margins and much more intensive advertising in markets with high margins. Their estimation strategies include an instrumental-variables estimator based on the geographic pattern of entry of Craigslist to the markets.

1.4 Cyclical changes in market power and profit margins

Macroeconomics has spawned a large literature on countercyclical market power. Bils (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. Bils made important adjustments based on cyclical variations in the incidence of overtime wages. Rotemberg and Woodford (1999) embraced Bils'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 (2013) and Nekarda and Ramey (2011) challenge the findings of countercyclical market power in favor of cyclically constant markups resulting from Bils's overstatement of the incidence and magnitude of overtime premiums.

Bils 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 during 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.

1.5   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 volume 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

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

The profit-maximizing price is

$$p^* = \frac{\epsilon}{\epsilon - 1} \, f \, c \tag{2}$$

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

$$p = m \, p^*. \tag{3}$$

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 (15) imply

$$p = m \, f \, \frac{\epsilon}{\epsilon - 1} \, c. \tag{4}$$

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.

8

2.1 Advertising

The first-order condition for advertising is

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

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

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

Substituting for p from equation (15) and for p^* from equation (2) restates the right-hand side in terms of exogenous influences:

$$R = \frac{\kappa A}{pQ} = \alpha \frac{(m-1)\epsilon + 1}{f \, m \, \epsilon}$$ (7)

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

$$R = \frac{\alpha}{\epsilon},$$ (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 Advertising capital

The variable A is the volume of advertising currently influencing demand. It should be distinguished from the current volume of advertising effort, a, because the effect of that effort lasts, in part, into future years. In other words, A is a capital stock, while a is gross

investment. 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

$$A_t = a_t + (1 - \delta)A_{t-1}. \tag{9}$$

Here δ is the rate of depreciation. The annual cost of the services of a unit of a stock of advertising over one year is

$$\kappa_t = \frac{r + \delta}{1 + r} \, v_t. \tag{10}$$

Here r is the annual real interest rate and v_t is the price of investment in advertising. Notice that this formula is $\kappa_t = v_t$ if there is complete depreciation within a year: $\delta = 1$.

2.3 Labor share

The second key variable is the labor share

$$\lambda = \frac{W}{pQ}. \tag{11}$$

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

$$\lambda = \frac{\gamma \, c \, Q}{pQ} = \gamma \frac{\epsilon - 1}{\epsilon} \frac{1}{f \, m} \tag{12}$$

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.

2.4 Solving for the wedges

From the propositions above,

$$\log R = (\epsilon - 1) \log m - \log f + \mu_R \tag{13}$$

and

$$\log \lambda = -\log m - \log f + \mu_\lambda, \tag{14}$$

10

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

$$\log m = \frac{\log R - \log \lambda}{\epsilon} + \mu_m \tag{15}$$

and

$$\log f = -\log \lambda - \frac{\log R - \log \lambda}{\epsilon} + \mu_f. \tag{16}$$

Here μ_m and μ_f are constant and slow-moving influences derived in the obvious way from μ_R and μ_λ. Given the value of the residual demand elasticity ϵ, the profit-margin wedge m and the product-market wedge f are observed time series.

2.5 The role of cyclical movements

The main goal of this paper is to 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 is procyclical and the labor share λ is close to non-cyclical. The expectation is that the wedges are both countercyclical—they measure forces that mediate reductions in employment when they rise.

2.6 Extracting the cyclical component of the measured wedges

This discussion follows Baxter and King (1999). The most intuitive way to describe the short-run and long-run properties of annual time series is in terms of *periodicity*, the number of years between one peak and and the next in a cyclical component. Short-run, high-frequency components have low periodicity, starting at two years, while long-run, low-frequency components have high periodicity. The ultimate long-run component, a constant, has infinite periodicity. On the other hand, the most convenient measure of frequency for the mathematics of time-series analysis is one normalized so that the lowest frequency is zero and the highest is π. Frequencies under this convention are often designated ω. The periodicity of a component at frequency ω is $2\pi/\omega$. The reason for this convention is that standard time-series analysis takes the history of a time series to be a weighted average of sine waves

and cosine waves. All the math derives from the fact that

$$e^{i\omega t} = \cos \omega t + i \sin \omega t. \tag{17}$$

Now consider a time series x_t. A linear filter is a lag polynomial $\phi(L)$. The time series $\hat{x}_t = \phi(L)x_t$, with adroit choice of $\phi(L)$, can emphasize business-cycle periodicities—ranging from once every two years to once every 5 years—and attenuate higher periodicities. The factor or *gain* applied to a periodicity with frequency ω is $|\phi(e^{i\omega})|$, the complex modulus of ϕ evaluated at $e^{i\omega}$. When the same filter is applied to the left- and right-hand variables of a regression, the overall gain has no effect on the regression. Consequently, the gain function can be normalized. I divide the gain by its maximum value over all periodicities.

Baxter and King discuss bandpass filters, constructed to have a gain close to a constant for low periodicities and close to zero for high periodicities. I do not use bandpass filters in this paper, because there is no sharp boundary between the periodicities of the business cycle and of medium and long-run components of aggregate variables. Simple filters are available with gains that decline smoothly with periodicity. I adopt Baxter and King's restriction to filters that have a gain of zero at infinite periodicity. These have the property that the sum of the coefficients is zero; that is, $\phi(1) = 0$. The filters I use here are:

$$\text{One-sided 3-coefficient: } \phi(L) = 1 - \psi L - (1 - \psi)L^2 \tag{18}$$
$$\text{Symmetric 5-coefficient: } \phi(L) = -(0.5 - \psi)L^2 - \psi L + 1 - \psi L^{-1} - (0.5 - \psi)L^{-2} \tag{19}$$

The parameter ψ influences the shape of the gain function in both filters. Among the one-sided filters, the best, in the sense of capturing cyclical periodicities and suppressing higher ones, is $\psi = 1$, the simple first difference. Among the two-sided filters, the best is $\psi = 0.35$, so $\phi(L) = -0.15L^2 - 0.35L + 1 - 0.35L^{-1} - 0.15L^{-2}$. Figure 1 shows the gain functions for the two chosen filters. The two-sided filter captures more of the periodicities associated with the business cycle and applies a lower gain to longer-term movements.

2.7 The business-cycle component

Not all of the movements of the advertising/GDP ratio and the labor share at the lower periodicities are associated with the business cycle. To measure the business-cycle component, I use the regression coefficients of the variables on the filtered employment rate (one minus the unemployment rate). In addition to the coefficient on the contemporaneous employment
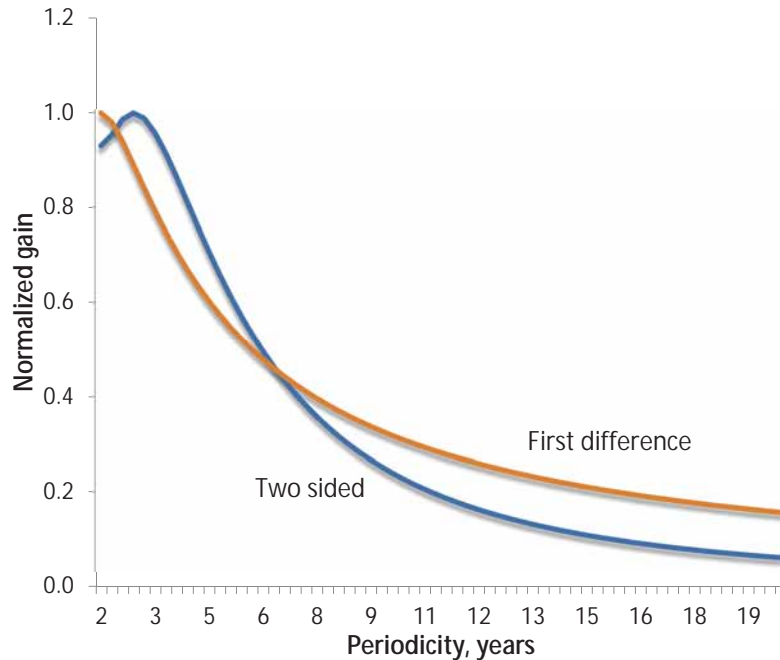
12

Figure 1: Gain Functions for Filters that Emphasize Cycle Movements and Suppress Longer-Term Movements in an Annual Time Series

rate, I calculate the coefficient on the employment rate of the previous year. Both the advertising ratio and the labor share are much more correlated with the lagged employment rate than with the current one.

3 The Advertising/GDP Ratio

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. A complete table downloaded from Galbi's website is in the backup file on my website.

For 2005 through 2010, the Census Bureau 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. These data are no longer available on the Bureau's website. The backup file contains a complete copy of the table from the Census. 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.

3.1 Depreciation rate

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. I take the real interest rate to be 5 percent per year.

3.2 Behavior of the advertising/output ratio, R

Figure 2 shows the ratio of advertising cost to private nominal GDP, stated as an index, with vertical bars shaded darker in years when the employment rate was low. The darker the vertical bars, the lower the *level* of employment (in contrast to recession bars, which would show *declines* in employment). The data run from 1950 through 2010. I calculate advertising expenditure by forming the capital stock of cumulative spending as described in subsection 2.2.

Figure 3 shows the first-difference-filtered index of advertising cost, in the same format as Figure 2. Its cyclical movements are more prominent after filtering.

To describe the cyclical properties of the advertising/GDP ratio, R, I calculate regressions of the filtered data on the filtered employment rate (one minus the unemployment rate). These appear in Table 1 for the contemporaneous and lagged employment rates. With the contemporaneous rate, the results for both filters agree that the ratio is slightly procyclical, but the hypothesis that it is noncyclical is easily accepted. With the lagged employment rate, the results show overwhelming evidence that the ratio is procyclical. But the issue is whether the ratio is as *countercyclical* as predicted by a model in which the profit margin widens in recessions. The right-hand column of the table gives the p-value for the hypothesis that the coefficient is -6, as it would be if the residual elasticity of demand were $\epsilon = 6$, a reasonable value, and if the wedge and the employment rate varied over the cycle as one percentage point of wedge for each percentage point of employment, also a reasonable value. For both filters and both choices of employment timing, the p value is infinitesimal—it calculates to exactly zero using 16-digit arithmetic. The evidence is overwhelmingly against the amount of
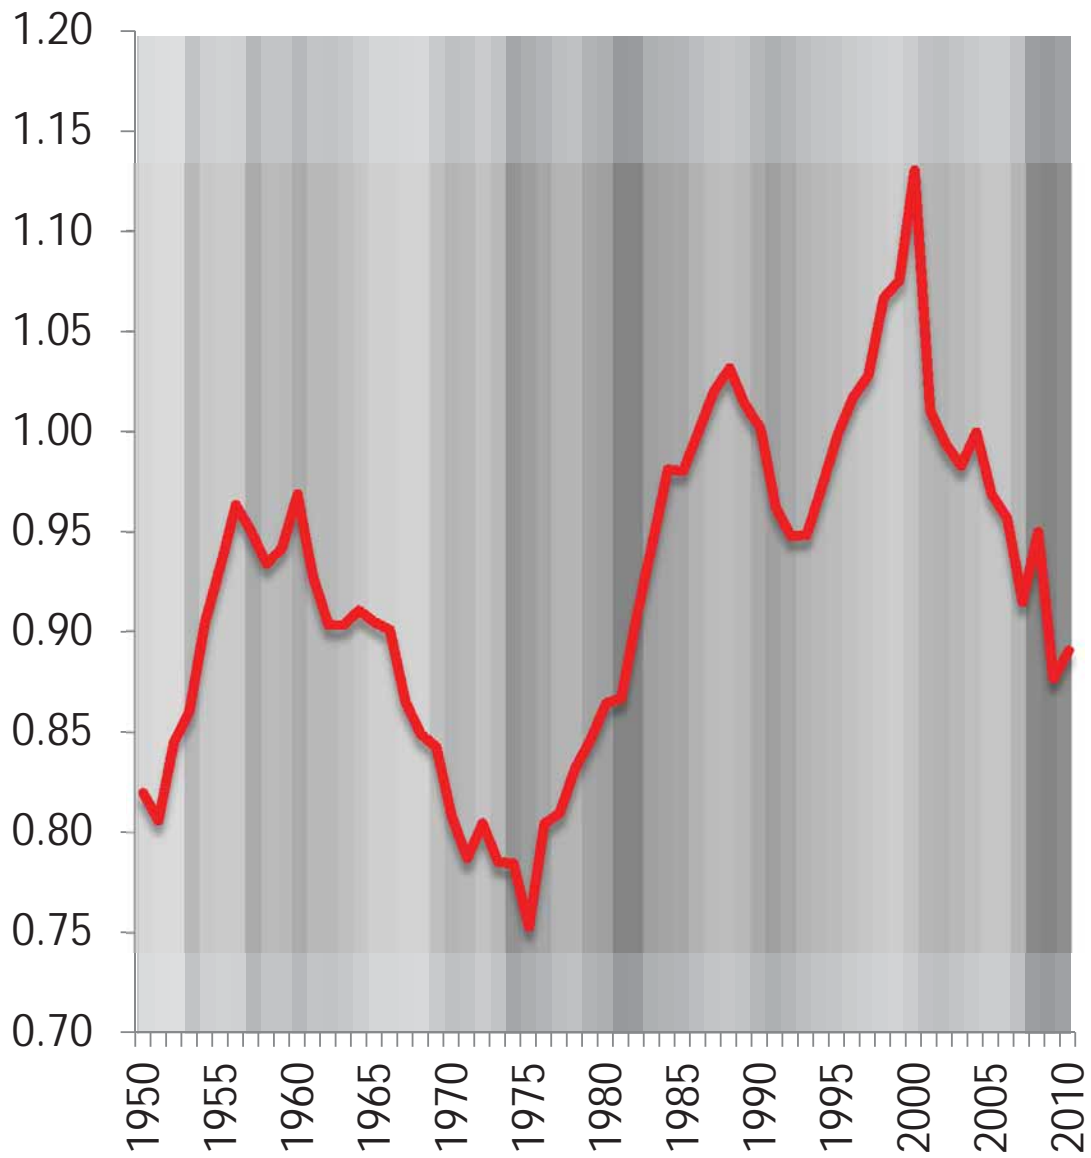
Figure 2: Index of the Ratio of Advertising Cost to Private GDP, with Shading in Inverse Proportion to the Employment Rate
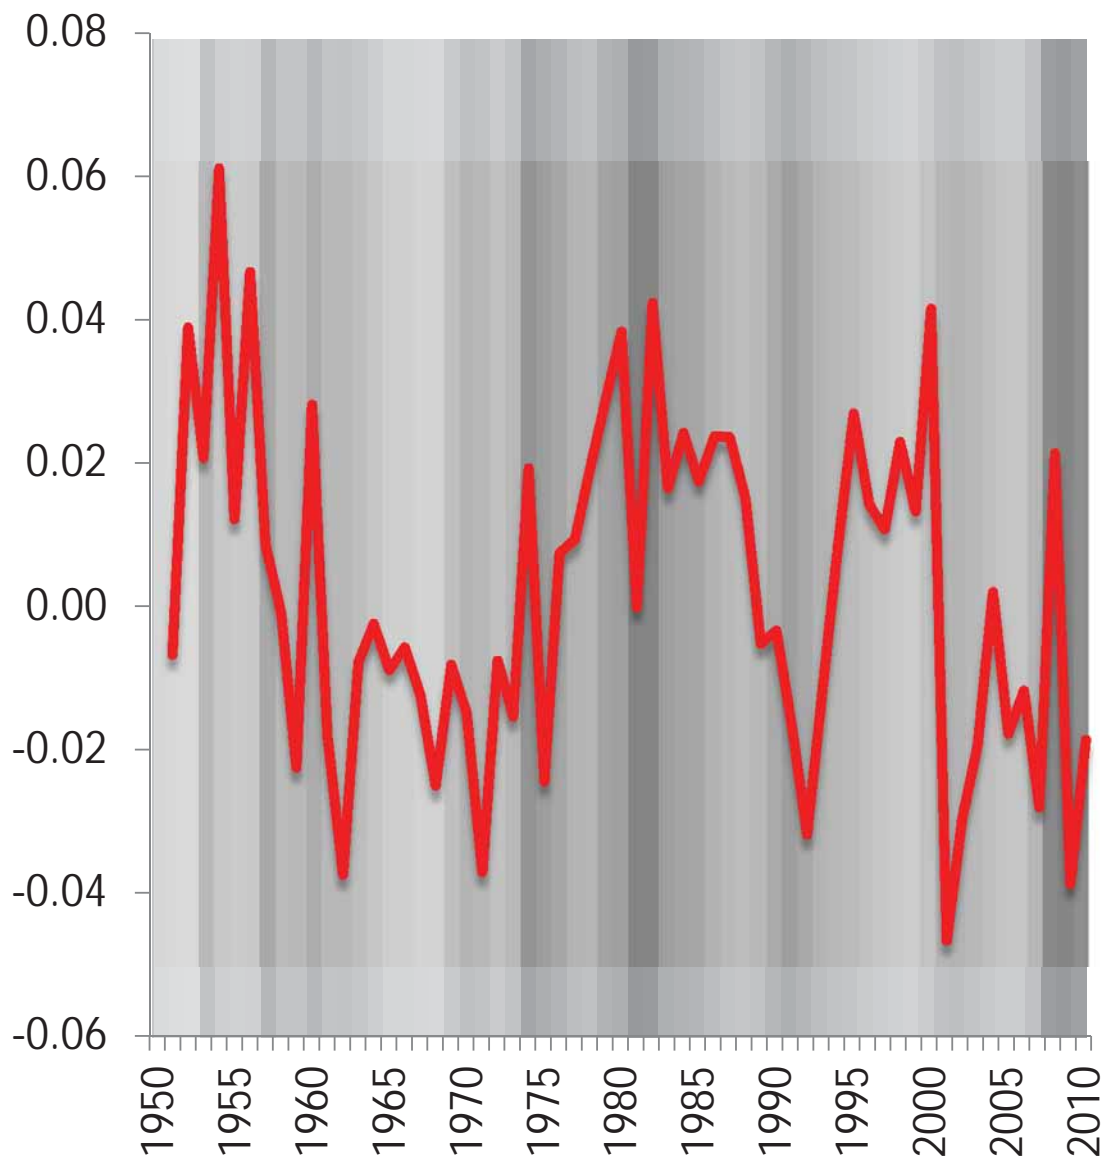
Figure 3: Filtered Index of the Ratio of Advertising Cost to Private GDP, with Shading in Inverse Proportion to the Employment Rate

Employment timing	Filter	Coefficient	Standard error	Years	Upper-tail p-value for coefficient = -6
Contemporaneous	First difference	0.21	(0.29)	1951-2010	0.00000
	Symmetric	0.10	(0.23)	1952-2008	0.00000
Lagged one year	First difference	0.85	(0.27)	1952-2010	0.00000
	Symmetric	0.82	(0.21)	1953-2008	0.00000

Table 1: Coefficients of Regressions of the Filtered Advertising/GDP ratio on the Employment Rate

countercyclical movement of the advertising ratio that a countercyclical profit-margin wedge would generate.

Though the evidence against highly countercyclical movements of the advertising ratio is overwhelming, that evidence does not completely rule out a large negative response to the profit-margin wedge m, because there may be confounding movements of the other wedge, f. 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 an 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 4 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.

Table 2 describes the cyclical movements of the labor share in the same format as Table 1 did for the advertising ratio. With contemporaneous employment, the results are easily compatible with the hypothesis of zero correlation. With lagged employment, the evidence is strong that the labor share is procyclical.
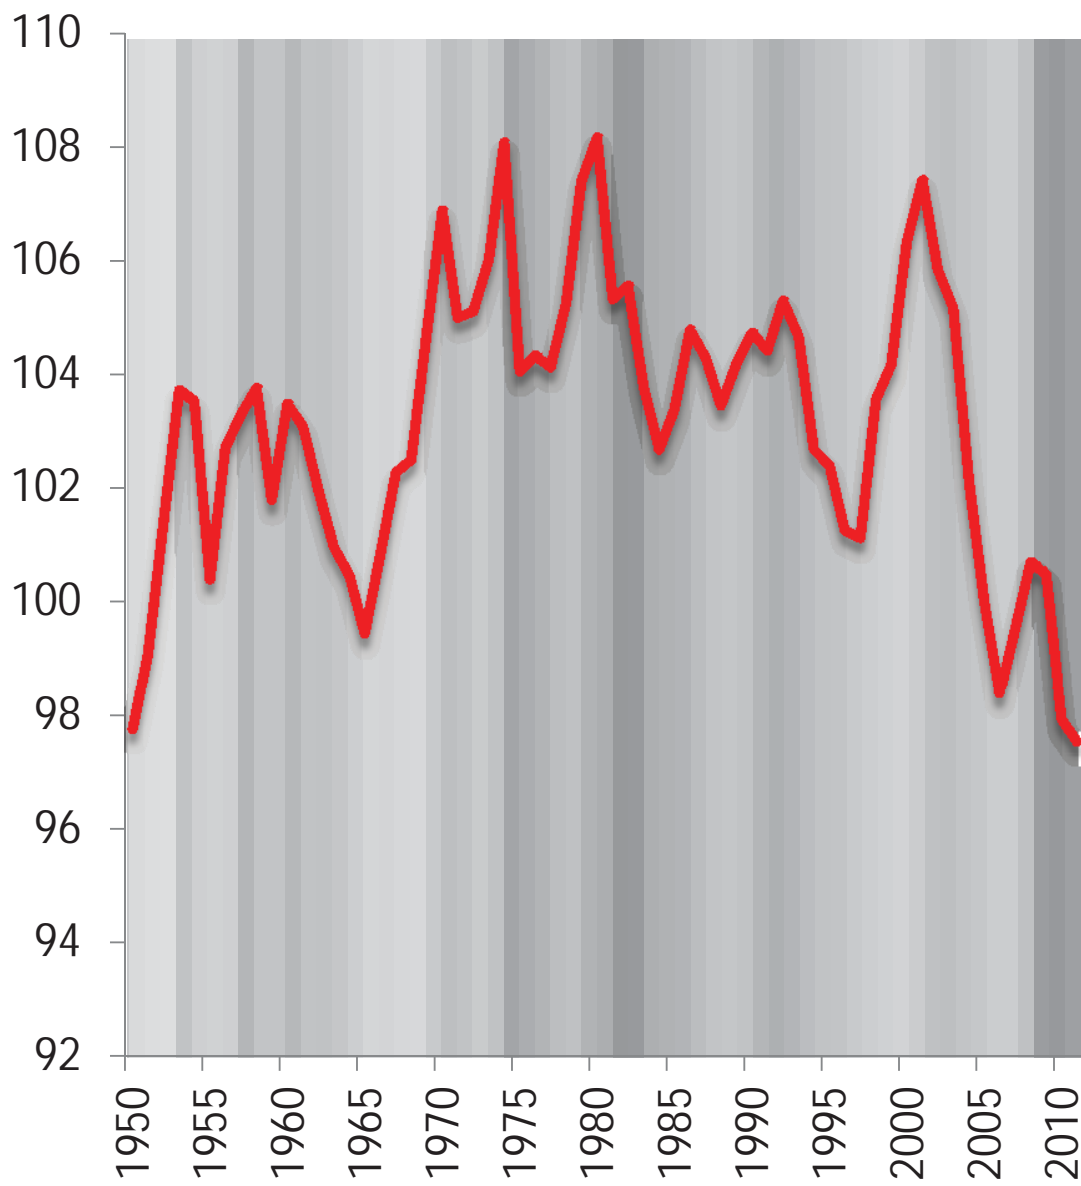
Figure 4: Index of the Labor Share of Non-Financial Corporate Income, with Shading in Proportion to the Employment Rate

Employment timing	Filter	Coefficient	Standard error	Years	Two-tail p-value for coefficient = 0
Contemporaneous	First difference	0.07	(0.18)	1951-2010	0.702
	Symmetric	0.05	(0.18)	1952-2008	0.784
Lagged one year	First difference	0.83	(0.15)	1952-2010	0.000
	Symmetric	0.82	(0.15)	1953-2008	0.000

Table 2: Coefficients of Regressions of the Filtered Labor Share on the Employment Rate

5 The Values of the Wedges and Their Effects on Employment

The values of the two wedges are

$$\log m = \frac{\log R - \log \lambda}{\epsilon} \tag{20}$$

and

$$\log f = -\log \lambda - \log m. \tag{21}$$

The calculated wedges inherit the filtering of the advertising ratio and the labor share, or equivalently, could be filtered after calculation from the unfiltered data. 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 5 shows the first-difference-filtered time series for the profit-margin wedge. Its movements are small and not visibly cyclical. Figure 6 shows the series for the product-market wedge using the same vertical axis scale. Its volatility is high and the increases are visible in a number of recessions.

Table 3 shows the coefficients of the regressions of the two filtered series for the profit-margin wedge m on the contemporaneous and lagged employment rate. All four combinations
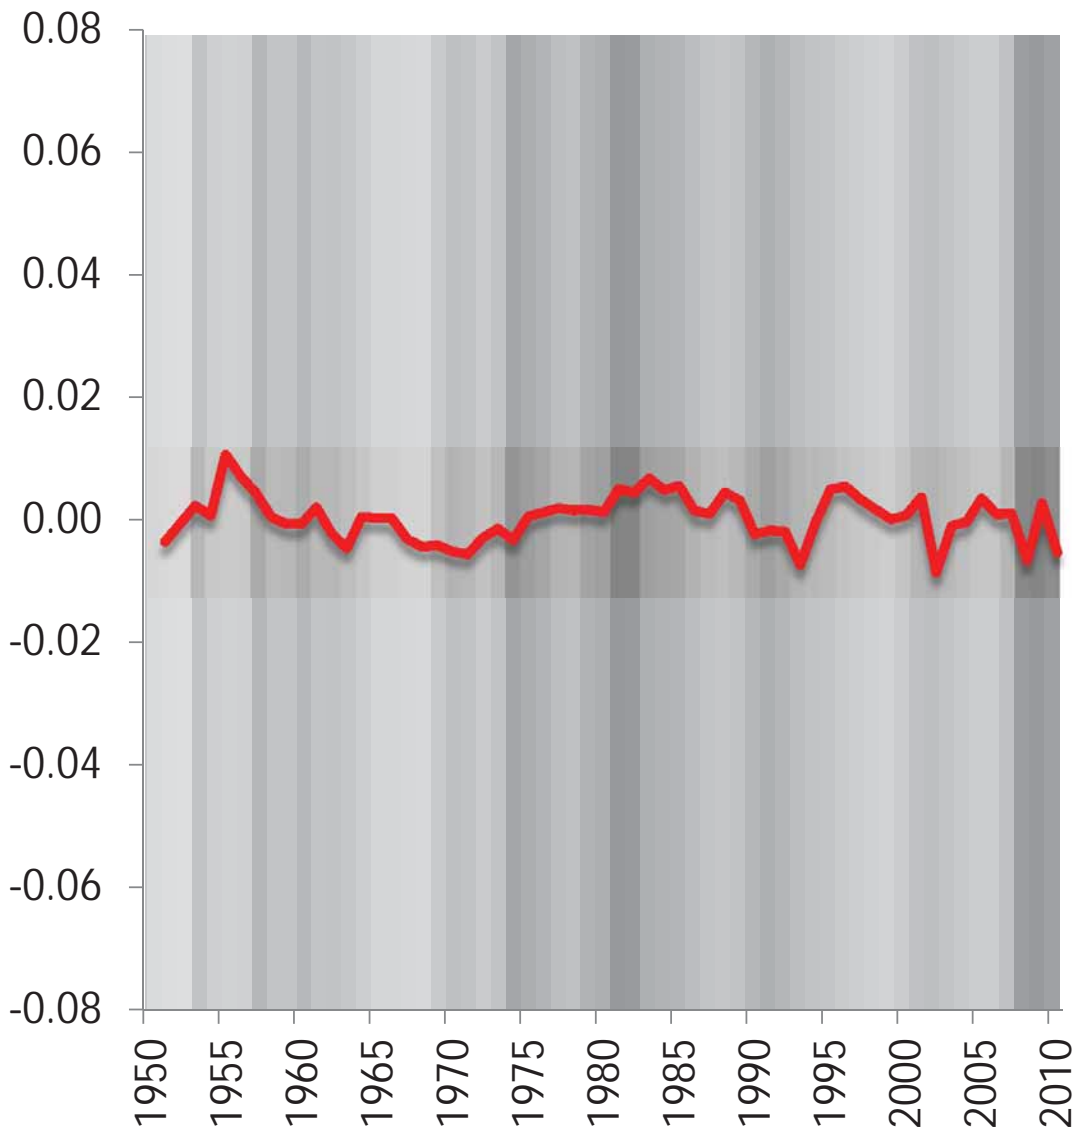
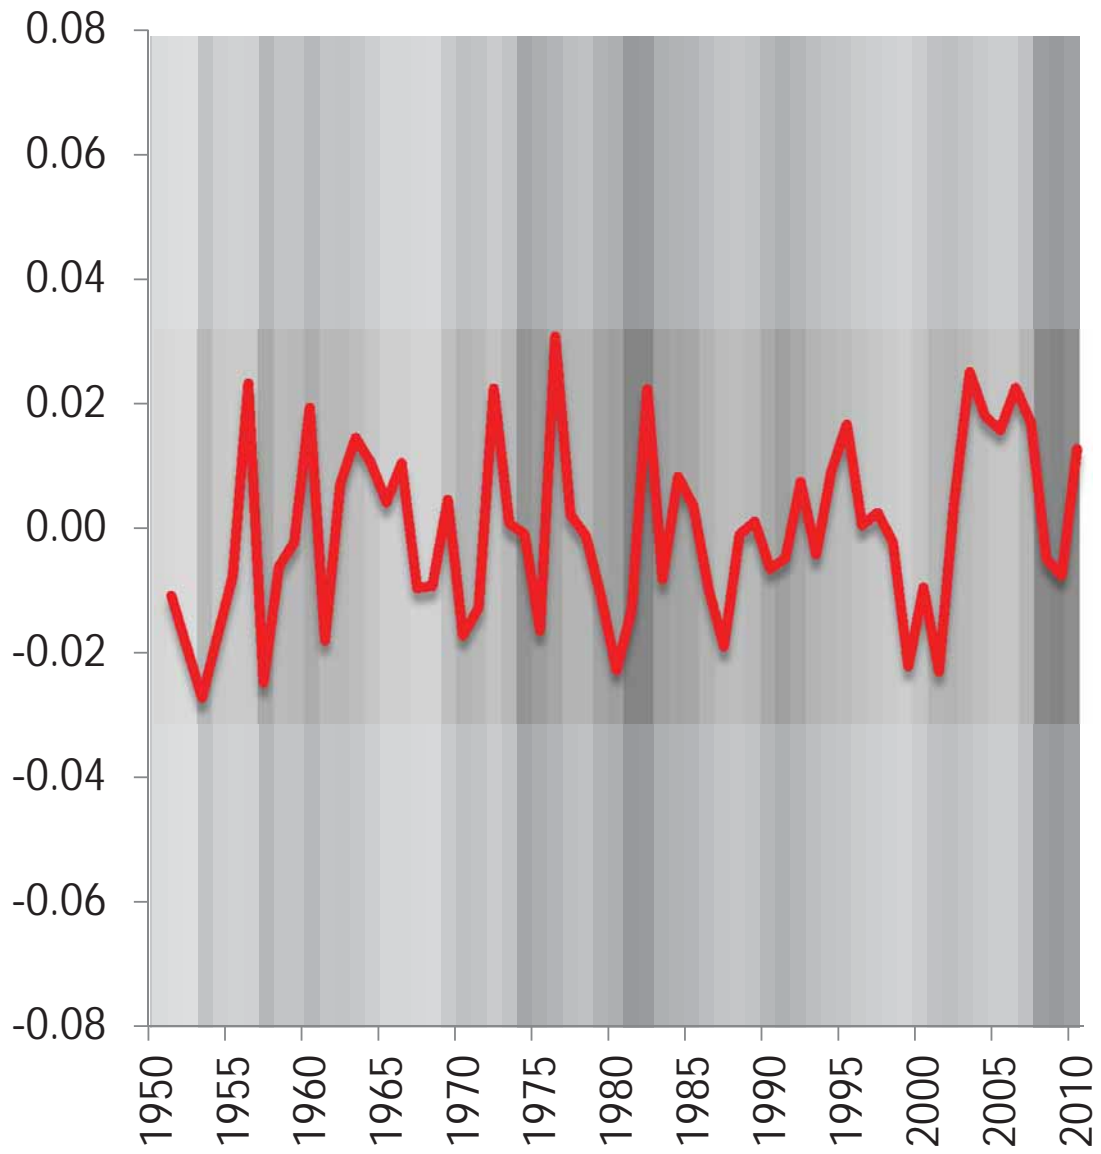Figure 5: Calculated Filtered Time Series for the Profit-Margin Wedge

Figure 6: Calculated Filtered Time Series for the Product-Market Wedge

Employment timing	Filter	Coefficient	Standard error	Years	Upper-tail p-value for coefficient = -0.1
Contemporaneous	First difference	0.02	(0.05)	1951-2010	0.004
	Symmetric	0.01	(0.04)	1952-2008	0.003
Lagged one year	First difference	0.00	(0.05)	1952-2010	0.014
	Symmetric	0.00	(0.04)	1953-2008	0.006

Table 3: Regressions of the Filtered Markup Wedge on the Employment Rate

Employment timing	Filter	Coefficient	Standard error	Years	Upper-tail p-value for coefficient = 0
Contemporaneous	First difference	-0.09	(0.18)	1951-2010	0.298
	Symmetric	-0.06	(0.17)	1952-2008	0.368
Lagged one year	First difference	-0.84	(0.14)	1952-2010	0.000
	Symmetric	-0.82	(0.14)	1953-2008	0.000

Table 4: Regressions of the Filtered Product-Market Wedge on the Employment Rate

confirm that cyclical movements of the wedge are tiny. The coefficients are quite small and accurately estimated. The p values for the hypothesis of a small countercyclical effect—a coefficient of -0.1—show strong evidence against even that small effect.

Table 4 shows the corresponding results for the product-market wedge f. Based on the contemporaneous employment rate, there is little sign of a cyclical component—the hypothesis of a zero coefficient is easily accepted. On the other hand, based on the lagged employment rate, there is a strong countercyclical component. The coefficient is large and negative for both filters and the hypothesis that it is zero is overwhelmingly rejected.

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 L. A three-way breakdown is

$$L_t = \theta \log m_t + \rho \log f_t + x_t, \tag{22}$$

where x_t captures all the other influences on employment. The coefficients θ and ρ are presumptively negative, 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 $\log m_t$ and $\log f_t$. But with outside information about the coefficients θ and ρ, 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 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 ρ 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

$$mf\frac{\epsilon}{\epsilon - 1}. \tag{23}$$

The two variables m and f enter with equal elasticities of minus 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. I take these to be the filtered values discussed earlier in the paper. I study 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

$$\mathrm{V}(L_t) = \theta \, \mathrm{Cov}(m_t, L_t) + \theta \, \mathrm{Cov}(f_t, L_t) + \mathrm{Cov}(x_t, L_t). \tag{24}$$

Divide by the variance of the employment rate to get

$$1 = \theta \, \frac{\mathrm{Cov}(m_t, L_t)}{\mathrm{V}(L_t)} + \theta \, \frac{\mathrm{Cov}(f_t, L_t)}{\mathrm{V}(L_t)} + \frac{\mathrm{Cov}(x_t, L_t)}{\mathrm{V}(L_t)}. \tag{25}$$

Note that this can be written more compactly as

$$1 = \theta\beta_m + \theta\beta_f + \beta_x, \tag{26}$$

where the βs are the coefficients of the corresponding variables regressed on L_t. These coefficients appear in Table 3 and Table 4.

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 the contribution of the residual. The presumption is that θ, β_m, and β_f are negative, so the contributions are positive. But nothing rules out a negative contribution, and some of the values of β_m are slightly, but not significantly, positive.

The calculations above depend on the parameter θ, the effect of wedges in general on the employment rate. Figure 7 shows how the calculation of the contributions depends on that parameter, using the values of $\beta_m = 0.0025$ from the third line of Table 3 (for lagged employment and the first-difference filter) and the corresponding $\beta_f = -0.84$ from Table 4. Recall that these results are based on $\epsilon = 6$. The horizontal axis is the effect of the wedge on employment, θ. At $\theta = -1$, the profit-margin wedge m accounts for -0.25 percent of the cyclical movements of the employment rate, the product-market wedge for 84 percent, and the other forces for the remaining 16 percent. With more negative values of θ, the product-market wedge accounts for an implausibly high fraction of the cyclical movements.

Table 5 reports the sensitivity of the results to the other determinants, with $\theta = -1$: the residual elasticity of demand, ϵ, the filter, and the choice of employment timing. The base case with $\epsilon = 6$ described above is in the middle of the table. A comparison across the columns of the table shows that the basic message of the paper—the weak role of the profit-margin wedge—holds for the wide range of values of the residual elasticity ϵ, for both filters, and for both choices of employment timing.

In the top line of the table, with $\epsilon = 3$, $\theta = -1$, contemporaneous employment, and the first-difference filter, the point estimate of the contribution of the profit-margin wedge is $\theta\beta_m = -0.05$ with a standard error of 0.09. This case suggests that the preponderance of the evidence is against any positive contribution of the profit-margin wedge, but a reasonable confidence interval would include some small positive values. The contribution of the product-market wedge is modestly positive, at 0.12, but its sign is statistically ambiguous. The residual elasticity implies extreme market power—the markup ratio is 1.5.
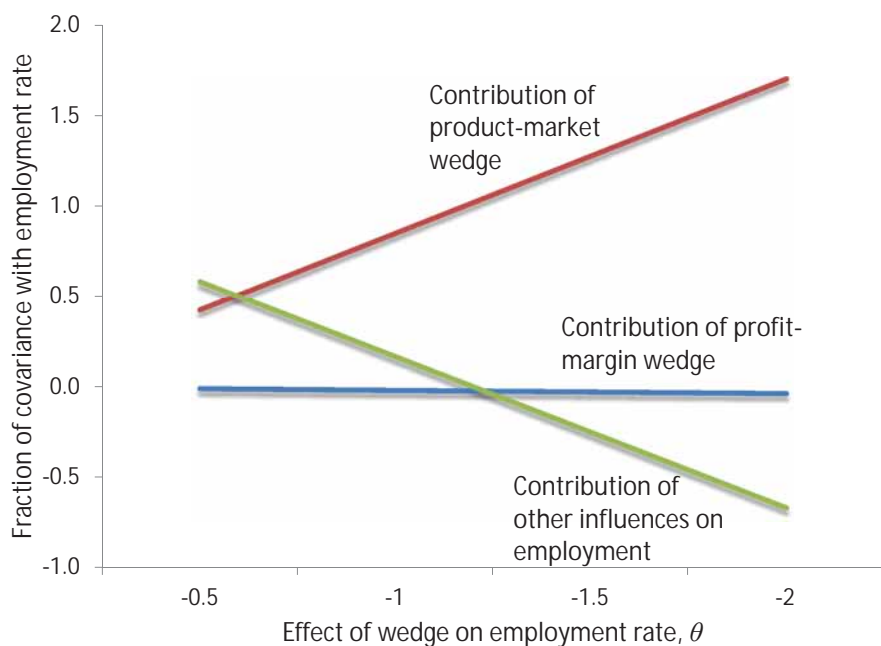
24

Fraction of covariance with employment rate

2.0

1.5

Contribution of product-market wedge

1.0

0.5

Contribution of profit-margin wedge

0.0

-0.5

Contribution of other influences on employment

-1.0

-0.5   -1   -1.5   -2

Effect of wedge on employment rate, θ

Figure 7: Contributions of Wedges to Employment Movements as Functions of the Parameter θ

Employment timing	Filter	*Implied contributions of wedges to cyclical movements in the employment rate*					
		ε, residual elasticity of demand					
		3		6		12	
		$\theta\beta_m$	$\theta\beta_f$	$\theta\beta_m$	$\theta\beta_f$	$\theta\beta_m$	$\theta\beta_f$
Contemporaneous	First difference	-0.05 (0.09)	0.12 (0.18)	-0.02 (0.05)	0.09 (0.18)	-0.01 (0.02)	0.08 (0.18)
	Symmetric	-0.02 (0.08)	0.07 (0.17)	-0.01 (0.04)	0.06 (0.17)	0.00 (0.02)	0.05 (0.18)
Lagged one year	First difference	-0.01 (0.09)	0.84 (0.15)	0.00 (0.05)	0.84 (0.14)	0.00 (0.02)	0.83 (0.14)
	Symmetric	0.00 (0.08)	0.82 (0.14)	0.00 (0.04)	0.82 (0.14)	0.00 (0.02)	0.82 (0.14)

Table 5: Implications of Alternative Values of the Residual Elasticity of Demand, with $\theta = -1$

In the middle of the top line, with $\epsilon = 6$, the point estimate of the fraction of employment movement explained by the profit-margin wedge is -0.02 and the fraction explained by the product-margin wedge is 0.09. Again, the contribution of both wedges is zero or small— neither is statistically unambiguously positive. At the right end of the top line, with $\epsilon = 12$, the fractions are -0.01 and 0.08. Once again, neither appears to an important determinant of employment volatility. These estimates of $\theta\beta_m$ have small standard errors, so they rule out any substantial role for the profit-margin shock. Results in the second line, for the symmetric filter, are quite similar, with slightly higher fractions explained by the profit-margin wedge.

In the lower panel of Table 5, based on associating the business cycle with lagged employment, the fraction of employment movement explained by the profit-margin wedge is estimated as zero to two decimal points in all but one case, where it is -0.01. Because the standard errors are small, these results give no support to the hypothesis that the profit-margin wedge has a meaningful effect on employment. On the other hand, the estimates of $\theta\beta_f$ imply a large role of the product-market wedge—it explains above 80 percent of the variance of filtered employment for all specifications, with fairly small standard errors.

The table demonstrates clearly that the finding of no important role for the product-margin wedge in employment fluctuations holds over the relevant range of values of the residual elasticity of demand, ϵ.

The table has conflicting results for the role of the product-market wedge, f, in employment fluctuations. In the upper panel, based on contemporaneous employment, the results give no meaningful support to the hypothesis of a positive role for the wedge, though the standard errors are larger than for the profit-margin wedge. In the lower panel, based on lagged employment, the results assign a large role to the product-market wedge.

5.2   Sensitivity to the depreciation rate for advertising

Table 6 shows how the results depend on the rate of depreciation of advertising, in the case of $\epsilon = 6$. The format is the same as for the preceding table. The top panel gives results based on contemporaneous employment. The left side is based on full depreciation within the year The contribution of the profit-margin wedge, $\theta\beta_m$, to the variance of employment is fairly small in magnitude, but unambiguously negative. On the right side, with slow depreciation at 30 percent per year, the contribution is again small in magnitude but is unambiguously positive. Estimates of the contribution of the product-market wedge, $\theta\beta_m$, are slightly neg-

Employment timing	Filter	Implied contributions of wedges to cyclical movements in the employment rate					
		δ, annual rate of depreciation					
		1		0.6		0.3	
		$\theta\beta_m$	$\theta\beta_f$	$\theta\beta_m$	$\theta\beta_f$	$\theta\beta_m$	$\theta\beta_f$
Contempo-raneous	First difference	-0.15 (0.07)	0.22 (0.17)	-0.02 (0.05)	0.09 (0.18)	0.11 (0.04)	-0.04 (0.18)
	Symmetric	-0.16 (0.06)	0.21 (0.17)	-0.01 (0.04)	0.06 (0.17)	0.14 (0.03)	-0.09 (0.17)
Lagged one year	First difference	0.14 (0.07)	0.69 (0.15)	0.00 (0.05)	0.84 (0.14)	-0.02 (0.04)	0.85 (0.14)
	Symmetric	0.17 (0.06)	0.65 (0.15)	0.00 (0.04)	0.82 (0.14)	-0.03 (0.04)	0.86 (0.14)

Table 6: Implications of Alternative Values of the Depreciation Rate

ative, but small in magnitude and have ambiguous signs. With lagged employment, in the bottom panel, the findings for the role of the profit-margin wedge include no case where its contribution is substantially positive and statistically unambiguous. Rather, the confidence intervals include only small positive values. The estimated role of the product-market wedge, $\theta\beta_f$, is large and statistically unambiguous, though its estimated value is in the 60-percent range for fast depreciation, in the left panel, rather than the 80-percent-plus values found for 60- and 30-percent depreciation.

The basic finding of the paper—that the profit-margin wedge is unimportant—easily survives in the range of depreciation rates from 1 to 0.3. The conclusion that the product-market wedge is quite important holds as well over this range, but only when lagged employment serves as the measure of the business cycle.

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 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. The support is conditional on the characterization of the business cycle in terms of the previous year's unemployment rate.

6 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 (for lagged employment). 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:

$$R = \frac{\kappa A}{pQ} = \frac{\alpha}{f_A \, f_Q \, m} \frac{(m-1)\epsilon + 1}{\epsilon} \tag{27}$$

From the derivation of equation (12), it is apparent that only the labor wedge f_L enters the formula for the labor share λ:

$$\lambda = \frac{W}{pQ} = \frac{1}{f_L \, f_Q \, m} \gamma \frac{\epsilon - 1}{\epsilon} \tag{28}$$

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 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:

$$\beta_R = d\log R/dL = (\epsilon - 1)\,\beta_m - \beta_f - \beta_A \tag{29}$$

and

$$\beta_\lambda = d\log \lambda/dL = -\beta_m - \beta_f - \beta_L. \tag{30}$$

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:

$$\beta_m = \frac{\beta_R - \beta_\lambda}{\epsilon} + \frac{\beta_A - \beta_L}{\epsilon} \tag{31}$$

and

$$\beta_f = -\frac{\beta_R + (\epsilon - 1)\beta_\lambda}{\epsilon} - \frac{\beta_A + (\epsilon - 1)\beta_L}{\epsilon}. \tag{32}$$

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

$$\hat{\beta}_m = \beta_m + \frac{\beta_L - \beta_A}{\epsilon} \tag{33}$$

and

$$\hat{\beta}_f = \beta_f + \frac{\beta_A + (\epsilon - 1)\beta_L}{\epsilon}. \tag{34}$$

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 cyclical response of the profit-margin wedge, β_m, is small, rather than large and countercyclical (negative) as implied by some macro models. Only a very large countercyclical coefficient for the advertising wedge (big negative β_A) or a very procyclical coefficient for the labor wedge (big positive β_L) would have much effect in concealing a large negative role for the profit-margin wedge through the bias term in the first equation.

A countercyclical advertising wedge would make advertising noncyclical by contributing a procyclical element to advertising that would conceal the countercyclical element arising from a countercyclical profit-margin wedge. Nothing comes to mind that would suggest such a phenomenon. If the countercyclical response β_A were 1.0 in magnitude, removing the bias move the inferred β_m down to –1/6, with $\epsilon = 6$.

A countercyclical measurement error in the labor share (positive 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 (33) 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 this paper, with the measurement of cyclical movements based on lagged employment, the labor share is quite procyclical without any adjustment of wages. There is a substantial difference between the correlation at business-cycle periodicities of the labor share and contemporaneous employment, on the one hand, and its correlation with lagged employment, on the other hand. The discrepancy demonstrates the need for further work on the cyclical properties of the labor share.

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 conclusion of the paper is that the contribution of the product-market wedge, β_f, is large, when lagged employment is taken as the measure of the cycle. This conclusion is more sensitive to consideration of additional wedges, notably the labor wedge. In equation (34), β_L enters with a coefficient $(\epsilon - 1)/\epsilon$, which is close to one. In the presence of a positive value of β_L, the implied negative value of β_f would be less negative and the implied contribution of the product-market wedge to employment movements correspondingly smaller.

7 Concluding Remarks

Figure 5 shows that the profit-margin wedge extracted from the advertising/GDP ratio R and the labor share λ has low volatility and no apparent cyclical movements. The wedge is close to uncorrelated with both this year's employment and last year's. The evidence against a countercyclical profit-margin mechanism for cyclical movements of employment seems strong.

The question then remains, what is the mechanism for employment fluctuations? This paper investigates the possible role of a product-market wedge that discourages economic activity in the same way that a sales tax does. The conclusion is that the product-market wedge f is *not* correlated with current-year employment change, but is strongly correlated with previous-year employment change. The wedge's adverse effect operates not in the year of a recessionary employment contraction, but rather in the following year. The product-market wedge is responsible for the fall in the advertising/GDP ratio R and for the decline in the labor share λ, in the aftermath of an employment contraction.

Because the profit-margin wedge has played an important role in fluctuations theory over the past two decades, the negative finding about that wedge helps focus the search for plausible business-cycle propagation mechanisms. The finding about the product-market wedge may also be useful, though it does not help understanding of the mechanism of the recession itself, only its aftermath.

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.

Baxter, Marianne and Robert G. King, "Measuring Business Cycles: Approximate Band-Pass Filters for Economic Time Series," *The Review of Economics and Statistics*, 1999, *81* (4), pp. 575–593.

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.

Gurun, Umit G., Gregor Matvos, and Amit Serub, "Advertising Expensive Mortgages," December 2013.

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 Differenciated 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," September 2013. Department of Economics, Princeton University.

Loecker, Jan De and Frederic Warzynski, "Markups and Firm-Level Export Status," *American Economic Review*, May 2012, *102* (6), 2437–71.

Molinari, Benedetto and Francesco Turino, "Advertising and Business Cycle Fluctuations," March 2009. Instituto Valenciano de Investigaciones Economicas, working paper WP-AD 2009-09.

Nekarda, Christopher J. and Valerie A. Ramey, "Industry Evidence on the Effects of Government Spending," *American Economic Journal: Macroeconomics*, January 2011, *3* (1), 36–59.

____ and ____ , "The Cyclical Behavior of the Price-Cost Markup," Working Paper 19099, National Bureau of Economic Research June 2013.

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.

Appendix

A   Full Derivation for Advertising

Optimal price:

$$p^* = \frac{\epsilon}{\epsilon - 1} f_Q c \tag{35}$$

Actual price:

$$p = m p^* \tag{36}$$

First-order condition for advertising:

$$\frac{\alpha}{A} Q \left(\frac{p}{f_Q} - c \right) = f_A \kappa \tag{37}$$

From above,

$$c = \frac{\epsilon - 1}{\epsilon} \frac{p}{m f_Q} \tag{38}$$

The first-order condition becomes

$$\frac{\alpha}{A} \frac{pQ}{f_Q} \frac{(m-1)\epsilon + 1}{\epsilon m} = f_A \kappa \tag{39}$$

and, finally,

$$R = \frac{\kappa A}{pQ} = \frac{\alpha}{f_A f_Q m} \frac{(m-1)\epsilon + 1}{\epsilon} \tag{40}$$

B   Full Derivation for Labor Share

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

$$\frac{f_L W}{cQ} = \gamma \tag{41}$$

Substitute for c:

$$f_L f_A m \frac{\epsilon}{\epsilon - 1} \frac{W}{pQ} = \gamma \tag{42}$$

so

$$\lambda = \frac{W}{pQ} = \frac{1}{f_L f_Q m} \gamma \frac{\epsilon - 1}{\epsilon} \tag{43}$$