Deciphering the fall and rise in the net capital share: accumulation, or scarcity?*

Matthew Rognlie†

Spring 2015

Abstract

In the postwar era, developed economies have experienced two substantial trends in the net capital share of aggregate income: a rise during the last several decades, which is well-known, and a fall of comparable magnitude that continued until the 1970s, which is less well-known. Overall, the net capital share has increased since 1948, but when disaggregated this increase comes entirely from the housing sector: the contribution to net capital income from all other sectors has been zero or slightly negative, as the fall and rise have offset each other. Several influential accounts of the recent rise emphasize the role of increased capital accumulation, but this view is at odds with theory and evidence: it requires empirically improbable elasticities of substitution, and it presumes a correlation between the capital-income ratio and capital share that is not visible in the data. A more limited narrative that stresses scarcity and the increased cost of housing better fits the data. These results are clarified using a new, multisector model of factor shares.

1 Introduction

How is aggregate income split between labor and capital? Ever since Ricardo (1821) pronounced it the "principal problem of Political Economy," this question of distribution has puzzled and inspired economists.

*Final draft for the Spring 2015 Brookings Papers on Economic Activity. For helpful comments and conversations I thank my discussants Robert Solow and Brad DeLong, the editors David Romer and Justin Wolfers, the participants at the Spring 2015 Brookings Panel on Economic Activity, as well as Adrien Auclert, Loukas Karabarbounis, Stephen Murphy, Thomas Piketty, Alp Simsek, Ludwig Straub, Ivan Werning, and Gabriel Zucman.

†MIT Department of Economics.

Views differ. In one popular interpretation, the division between labor and capital remains remarkably stable over time: Keynes (1939) called this "one of the most surprising, yet best-established, facts in the whole range of economic statistics," and Kaldor (1957) immortalized it as one of the stylized facts of economic growth. In contrast, another tradition has emphasized variation in income shares: Solow (1958) was famously skeptical, disputing the labor share's status as "one of the great constants of nature." Recently, Solow's view has experienced a resurgence, with the labor share apparently trending downward. Elsby, Hobijn and Sahin (2013) carefully document this decline for the US, and Karabarbounis and Neiman (2014b) describe a broad, worldwide retreat of labor income in favor of capital.

Some influential recent narratives of this shift adopt what I call the *accumulation view*: capital's share has risen, and will continue to rise, because of capital accumulation. According to Piketty (2014) and Piketty and Zucman (2014), several forces are driving up aggregate savings relative to income, and the resulting growth in the ratio of the capital stock to income has led to a rise in capital's share. Alternatively, Karabarbounis and Neiman (2014b) stress the role of falling prices for investment goods; in their account, lower prices lead to more aggregate capital investment and ultimately more capital income. Although these two narratives specify different initial shocks, the subsequent channel is common to both: accumulation of capital through investment leads to growth in capital income, because the rising quantity of capital is not fully offset by a fall in the returns per unit of capital.

This paper argues against the accumulation view, on both empirical and theoretical grounds. Empirically, it reveals that the long-term increase in the capital's net share of income in large developed countries has consisted entirely of housing. Outside of housing, capital's rise in recent decades has merely reversed a substantial earlier fall, and in neither direction have there been a parallel movement in the value of capital—all facts that are difficult to reconcile with the accumulation view.

From a more theoretical perspective, the accumulation view is only successful when the elasticity of substitution between labor and capital is sufficiently high. Clarifying the distinction between elasticities gross and net of depreciation, this paper argues that the elasticity required is much higher than the existing literature suggests (particularly in the Piketty (2014) case).

Moving beyond the canonical one-sector model to a multisector model that explicitly acknowledges some important dimensions of capital heterogeneity—for instance, the distinction between housing and non-housing, as well as the distinction between equipment and structures—I continue to find little support for either version of the accumulation view. Instead, a more viable (albeit incomplete) explanation of recent trends is that residential investment has become more expensive, and land scarcer. Although this has *lowered* the quantity of housing, there has been a more than offsetting rise in net rents per unit of housing, pushing up the contribution of housing to capital's net share of income. In short, the data and theory support a *scarcity view*: the net capital share is rising in part because some forms of capital are becoming relatively more scarce, not more abundant.

I begin the paper with a look at the evidence on factor income shares for large developed economies over the postwar period 1948–2010. Several conceptual issues are crucial, especially the distinction between gross and net shares. Although both gross and net concepts are worthwhile when interpreted properly, I argue that the net viewpoint—much less common among recent entries in the literature—is more directly applicable to the discussion of distribution and inequality, because it reflects the resources that individuals are ultimately able to consume. I also restrict attention to the private sector, and in light of the severe measurement difficulties for proprietor's income identified by Elsby et al. (2013) and others, I apply the net shares from the corporate sector to the non-housing sector as a whole.

This measurement reveals a striking discrepancy in the long-term behavior of gross and net shares, echoing the claims of Bridgman (2014). It shows that the net capital share generally fell from the beginning of the sample through the mid-1970s, at which point the trends reversed. In the long run, there is a moderate increase in the aggregate net capital share, but this owes entirely to the housing sector. Indeed, housing's average portion of the aggregate net capital share rose from roughly 3% to 9% over the sample period, even as the private sector fell from 23% to 20%. This essential role of housing is notably absent from previous discussions of the factor distribution of income, and represents an important new contribution of this paper. It parallels a large (though less dominant) role for rising housing wealth in the aggregate wealth-income ratio, which has been documented by Piketty and Zucman (2014), Bonnet, Bono, Chapelle and Wasmer (2014), and others. Although these two trends are sometimes conflated, their alignment is not preordained: in fact, section 5.2 finds that a shock to savings should push them in opposite directions.

Outside of housing, there is a pronounced U-shape in the net capital share, with a steep fall in the 1970s and a more recent recovery. At shorter horizons, there is also a strong cyclical element, as long acknowledged by observers ranging from Mitchell (1913) to Rotemberg and Woodford (1999). To gauge whether the long-term fall and rise is consistent with the accumulation view, I contrast it with the time series for the capital-income ratio, finding that there is little similarity between the two. Using US data on the value of the three major components of non-housing capital—equipment, structures, and land—I perform a simple decomposition of the net capital share into returns on these components, plus a residual that can be interpreted as representing firm markups over cost. Markups are responsible for most of the change in shares, in both directions; in particular, accumulation of equipment or structures cannot explain the recent rise.

With these facts in mind, I next ask whether the accumulation view is viable theoretically. First, I look at the canonical single-sector model with a production function $F(K, N)$ that combines capital and labor. Here, for both the Piketty (2014) and Karabarbounis and Neiman (2014b) narratives, the key parameter is the elasticity of substitution for F. One important oversight in past discussions, however, has been the distinction between gross and net F: the elasticity for gross production is always higher than the elasticity for net. The Piketty (2014) hypothesis—accumulation through aggregate savings driving up the net capital share—is only viable if the *net* elasticity of substitution is greater than 1, which I argue is out of line with most existing evidence. The related conjecture of rising $r - g$ requires even more unlikely levels of substitutability. By contrast, the Karabarbounis and Neiman (2014b) hypothesis only calls for a *gross* elasticity above 1, which I argue is more plausible but still unlikely.

Given the limitations of the single-sector model, to better confront the data and formulate an alternative to the accumulation view I build a multisector model that incorporates key distinctions between sectors (housing and non-housing) and types of capital (equipment, structures, and land). When calibrated to match the structure of the US economy, the model continues to contradict Piketty (2014). For any choice of lower-level elasticities near the range suggested by the literature, an increase in savings results in a lower net capital share. By contrast, the mechanism in Karabarbounis and Neiman (2014b) remains theoretically viable when labor and equipment are close substitutes, but it works by increasing the value of equipment relative to total income, which is not consistent with the time series evidence.

The multisector model offers better support for the *scarcity view*. If, as most evidence suggests, consumers' demand for housing is sufficiently inelastic, the rising price of residential investment and growing scarcity of land can account for most of the growth in housing's portion of capital income. Although this does not resolve all aspects of the time series—especially the fall and rise in the corporate sector—it does explain a sizable portion of he long-term contribution of housing.

Before the recent preeminence of the accumulation view, there were varied attempts to explain a falling labor share. Elsby et al. (2013) highlight the role of offshoring, while other papers emphasize additional structural and institutional forces.[1] This literature does not

[1]See, for instance, Azmat, Manning and Reenen (2012), who address the role of privatization, and

apply directly here, since it uses gross concepts rather than net. Nonetheless, given the diverse accounts that have been proposed, it is no surprise that *this* paper fails to find a single mechanism that can explain the recent behavior of factor shares in its entirety.

The paper proceeds as follows. Section 2 discusses the conceptual basis of factor shares and provides evidence on the postwar path of the net capital share among G7 economies, including a decomposition that isolates the role of housing. Section 3 uses a simple decomposition to analyze the trends in net capital share further, restricting itself to the US to make use of more detailed data on capital stocks. Section 4 examines the canonical single-sector model, clarifying the different between net and gross elasticities. Section 5 integrates data and theory by building a multisector model, which refutes the accumulation view more definitively and supports the scarcity view as a partial alternative.

2 Evidence on factor income shares in developed countries

2.1 Conceptual issues

The notion of a "labor" or "capital" share is not monolithic; there are several ways to define and measure these concepts, and different choices lead to strikingly different interpretations of the data.

Decomposing gross value added. In the national accounts, the gross value added of a sector at market prices—the value of its gross output, minus the intermediate inputs used in production—can be divided[2] into three components: labor income (which includes both wages and supplementary compensation), taxes on production, and gross capital income (usually called "gross operating surplus" in the national accounts). Since the second component, taxes on production, does not accrue to either labor or capital, when analyzing the distribution of income between factors it is often convenient to subtract this component, leaving us with gross value added at *factor cost*. This can then be divided entirely into labor and gross capital shares, which sum to 1. Since I focus in this paper on the division of income between capital and labor, I will generally use this approach.

It is important to recognize that the split of value added between labor and capital is only the *initial* distribution. Labor income goes both to wages and to supplementary benefits, and a sizable share of wage income is subsequently paid to the government in

Arpaia, Perez and Pichelmann (2009), who draw attention to capital-skill complementarity.

[2] This decomposition potentially applies at many levels of aggregation: for instance, the "sector" may be the entire domestic economy, in which case gross value added at market prices is called gross domestic product (GDP).

taxes. Capital income is ultimately apportioned between many recipients, including the government (in the form of corporate and proprietor income taxes) and both debt and equity investors.

For instance, consider a sawmill. The gross value added at factor cost is the difference between its sales of lumber and the cost of logs, excluding taxes on production. Once all compensation of employees at the sawmill is subtracted, the remainder is its gross capital income. Some of this capital income will be paid to lenders in the form of interest, some will be paid to the government in taxes on profits, and the rest may be retained on the balance sheet of the sawmill or distributed as dividends to shareholders. Gross capital income is thus a very broad concept, encompassing funds that are ultimately paid out to many different recipients—it is unaffected, for instance, by the split in financing between debt and equity.[3]

Gross versus net: concepts. An alternative to gross value added is *net* value added, which subtracts depreciation. This can be divided into labor and *net* capital income, the latter being gross capital income minus depreciation. Whether a gross or net measure is more appropriate depends on the question being asked: the allocation of gross value added between labor and gross capital more directly reflects the structure of production, while the allocation of net value added between labor and net capital reflects the ultimate command over resources that accrues to labor versus capital.

For instance, in an industry where most of the output is produced by short-lived software, the gross capital share will be high, evincing the centrality of capital's direct role in production. At the same time, the net capital share may be low, indicating that the returns from production ultimately go more to software engineers than capitalists—whose return from production is offset by a loss from capital that rapidly becomes obsolete.

Both measures are important: indeed, a rise in the gross capital share in a particular industry is particularly salient to an employee whose job has been replaced by software, and it may proxy for an underlying shift in distribution *within* aggregate labor income—for instance, from travel agents to software engineers. The massive reallocation of gross income in manufacturing from labor to capital, documented by Elsby et al. (2013), has certainly come as unwelcome news to manufacturing workers. But when considering the ultimate breakdown of income *between* labor and capital, particularly in the context

[3]This invariance can be very useful in analyzing trends—for instance, when high inflation pushes up nominal interest rates, a large share of capital income is often paid to bondholders in the form of nominal interest. As Modigliani and Cohn (1979) memorably observed in the context of late-1970s inflation, this causes recorded profits to dramatically understate true profits, since they do not reflect the gain from real depreciation in nominal liabilities.

of concern about distribution in the aggregate economy, the net measure is likely more relevant. This point is affirmed by Piketty (2014), who uses net measures; the welfare relevance of net concepts is elucidated by Weitzman (1976).

Gross versus net shares: measurement and history. Historically, the study of income shares has spanned both gross and net concepts: indeed, the famous quote by Keynes (1939) about the stability of labor's share referred to data on net shares, as did Kaldor (1957)'s influential stylized fact.

More recently, however, the vast majority of work on the topic—including Karabarbounis and Neiman (2014b)'s well-known documentation of the declining global labor share—has examined gross shares. To a large extent, this is because gross shares are easier to measure and interpret: as economists since Kalecki (1938) have observed, net income inherently involves a somewhat arbitrary computation of depreciation. High-quality data on gross shares is available for more countries, more years, and more levels of aggregation within each country.

Recently, debate has intensified about the empirical importance of this distinction. Bridgman (2014) argues that the inclusion of depreciation—and, to a lesser extent, taxes on production—in the denominator of the labor share has caused economists to greatly overstate the magnitude and novelty of the labor share's decline. Augmenting their global dataset with information on depreciation, Karabarbounis and Neiman (2014a) argue to the contrary that gross and net labor shares have mainly moved together, and that moving from gross to net shares at most moderately attenuates the downward trend. In my data analysis, I will focus on net shares, finding that the concerns in Bridgman (2014) are valid, especially in the years preceding the start of the Karabarbounis and Neiman (2014a) sample.

Mixed income and other concerns. The distinction between gross and net is not the only concern when computing income shares. Another crucial problem is how to allocate "mixed" income—income earned by the self-employed that is recorded in the national accounts as going to capital. The central difficulty is that this income includes both returns to labor and returns to the capital investments made by the self-employed, with no data available to disentangle the two. This was an essential question for early students of the labor share in the US: as Johnson (1954) and others pointed out, the dramatic rise in workers' share of income in the first half of the twentieth century was in large part due to the shift from entrepreneurial income (often on farms) to formal labor income.

One solution is to disregard the entrepreneurial sector of the economy, limiting attention (for instance) to the labor share within the corporate sector. In any attempt to measure the labor share for the economy as a whole, however, some approach to dividing mixed income must be chosen—and this choice can matter a great deal. Indeed, Elsby et al. (2013) demonstrate that the "headline" measure provided by the BLS most likely exaggerates the decline in the US gross labor share, due to weaknesses in its approach to imputing labor income for the self-employed. This approach assumes that the self-employed receive the same average compensation per hour as all other workers—an imputation that, although popular and tractable, has some unlikely implications for the US data.

Alternative approaches to dividing mixed income, discussed by Gollin (2002), take several forms: they may do a more sophisticated estimation of labor income for the self-employed based on personal characteristics, or assume that the entrepreneurial sector has the same division between labor and capital as either some other sector or the economy as a whole. I follow Piketty and Zucman (2014) in adopting a form of the latter imputation, assuming that the non-corporate sector (excluding housing) has the same net capital share as the corporate sector.

Finally, another difficult point is the treatment of general government, as well as any other sectors whose output is valued in the national accounts "at cost"—meaning that gross value added is set equal to labor and depreciation costs—rather than by the market. Here, net capital income equals zero by construction; regardless, it is unclear what net capital income would mean in the context of government.

2.2 Income shares in the G7

To better understand the recent evolution of factor shares, I turn to a panel with national accounts data from the G7—which consists of the US, Japan, Germany, France, the UK, Italy, and Canada, currently the seven largest advanced economies by nominal GDP. Most of the data for the panel is derived from the Piketty and Zucman (2014) database, which in turn is taken directly from each country's national accounts publications.

Although this is a much narrower selection of countries than in the global panels of Karabarbounis and Neiman (2014a,b), it has several offsetting advantages. Most importantly, it covers a longer timespan: five countries have data starting in 1960 or earlier, and three countries have data starting in 1950 or earlier.[4] By contrast, Karabarbounis and

[4] The full set of start dates is 1948 (France, UK, US), 1955 (Japan), 1960 (Canada), 1990 (Italy), and 1991 (Germany). Data for France, the UK, and the US is available starting even earlier, but I focus on 1948 onward because that is when the necessary data starts becoming available for my subsequent, more detailed exercise for the US in section 3. This also keeps the focus on postwar dynamics, detached from the sizable

Neiman (2014a,b)'s dataset starts in 1975, and for many small and developing countries data only starts becoming available much later. Since the net labor share in most countries was close to its postwar peak in the mid-1970s, this offers an incomplete view of the overall trend. The dataset here also permits greater disaggregation, particularly along a dimension that will turn out to be crucial (housing versus the rest of the economy). By focusing on developed economies, it loses some generality but stays closer to the contemporary debate about inequality and income distribution, which has mostly dealt with the developed world.

Estimated average shares. To summarize the evolution over time of various income share measures $s_{i,t}$, I follow Karabarbounis and Neiman (2014a,b) by running panel regressions of the form

$$s_{i,t} = \phi_i + \alpha_t + \epsilon_{i,t}$$

for countries i and years t. I then display the yearly fixed effects α_t, normalizing them such that the fixed effect for the first year of the sample, α_{1948}, equals the average share in the dataset in 1948.[5] I run both unweighted and weighted regressions; the weight for a country is its share of the sample's aggregate GDP in that year, as measured at PPP by version 8.0 of the Penn World Table.[6] (For convenience, I will refer to these normalized time fixed effects as yearly "averages".)

Unlike in the usual presentation, I deal with the *capital* share rather than its complement, the labor share. Of course, since I deal with value added at factor cost, the capital share is always one minus the labor share; I focus on the former because I will emphasize the composition of capital income.

Overall capital shares: net and gross. First I consider average capital shares for the private economy (excluding government, whose net capital share is zero by construction). As discussed in section 2.1, I deal with the problem of self-employment income by following Piketty (2014) and Piketty and Zucman (2014) in the assumption that the net capital share in non-corporate, non-housing sector equals the net capital share in the corporate

dislocations associated with depression and wartime, and mostly postdates the transition from agricultural self-employment to formal employment that bedeviled older analysts like Johnson (1954).

[5]When countries have different trends in $s_{i,t}$, there will be an artifactual discontinuity in α_t when a country enters the sample, which in principle could deliver a misleading impression of the actual year-to-year changes in $s_{i,t}$. In practice, this does not seem to be much of an issue here, and alternative approaches—for instance, averaging the first differences $\Delta s_{i,t}$ across countries in the sample for each year t, then plotting the cumulative average first difference over time—deliver similar results.

[6]See Feenstra, Inklaar and Timmer (2013).

sector.[7]

Figures 1 and 2 report the average net and gross capital shares, respectively. As figure 1 demonstrates, the postwar behavior of the net capital share is characterized not so much by a secular rise as by a precipitous *fall* in the 1970s, which preceded a steady rebound. In this light, it is clear why Karabarbounis and Neiman (2014a,b)—with a sample starting in 1975, the year in which the unweighted estimate for the net capital share hits its minimum—observe such a dramatic and pervasive rise in capital income relative to labor.

Although Piketty (2014) and others have documented an overall U-shaped trend in the capital share, the claims about timing are quite different: for instance, Piketty (2014) observes that capital's aggregate valuation and share of income fell greatly in the first half of the twentieth century, during the depression and two world wars. The postwar period is characterized as a period of recovery from this decline. Yet figure 1 shows that if anything, the first half of the postwar era experienced a fall in the net capital share, and we are only today returning to levels achieved in the immediate aftermath of the war.

Set against figure 1, figure 2 reveals that there is a remarkable difference between the long-run behavior of net and gross shares, echoing the results of Bridgman (2014): since average depreciation as a share of gross value added has risen, the gross capital share displays much more of a long-term upward trend. Crucially, much of this disparity emerges before the mid-1970s, perhaps explaining why Karabarbounis and Neiman (2014a) do not detect such an important role for depreciation in their sample. Given the unreliability of depreciation figures at high frequencies, the sudden rise in depreciation prior to the mid-1970s (which causes the divergence between gross and net) should not be given too much credence. The long-term rise in depreciation, however, appears much more robust. As Koh, Santaeulalia-Llopis and Zheng (2015) discuss, it is due in part to rapidly depreciating intellectual property—especially software—included in the capital stock.

As argued in section 2.1, net shares are likely most relevant for discussions of distribution and inequality. Still, figure 1 paints a perhaps ambiguous picture of the net capital share: the recent rise might be in part just a recovery from the anomalously low levels of the 1970s, but the capital share is now reaching and even surpassing the heights previously achieved in the 1950s and 60s. To what extent, then, is the current high share of capital income a truly novel phenomenon? This question is best addressed by disaggregating further along an important dimension, distinguishing between capital income

[7] There are two exceptions: the Canadian national accounts already provide a decomposition of mixed income into labor and capital, which I use; and the Japanese national accounts do not fully break out the corporate sector, necessitating some additional imputations.

from housing and capital income from the rest of the economy.

Composition of the net capital share: the role of housing. Figure 3 subdivides the aggregate net capital share from figure 1 into two components: net capital income originating in the housing sector, and net capital income from all other sectors of the economy.[8] It reveals that the aggregate net capital share originating in sectors other than housing has seen only a partial recovery since the 1970s; it remains below the levels of the 1950s, and slightly below or at par with the levels of the 1960s. In contrast, housing's contribution to net capital income has expanded enormously, from roughly 3pp in 1950 to nearly 10pp today.

Housing's central role in the long-term behavior of the aggregate net capital share has, to my knowledge, not been emphasized elsewhere. It demands careful scrutiny. Income from housing is unlike most other forms of capital income recorded in the national accounts: in countries where homeownership is dominant, most output in the housing sector is recorded as *imputed* rent paid by homeowners to themselves. It may not be a coincidence that Germany, which table 2 reveals to have by far the lowest housing component of net capital income, also has the lowest homeownership rate in the G7. Indeed, imputed rents from owner-occupied housing should arguably be treated as a form of mixed income akin to self-employment income: in part, they reflect labor by the homeowners themselves. Figure 3 may therefore exaggerate the level of true "capital" income originating in the housing sector.

Nevertheless, even if figure 3 exaggerates the *level* of capital income from the housing sector, this does not necessarily explain the vast *increase* in housing capital income—unless the bias is greater today than in the past. One possible contributor to the trend could be a rise in the rate of homeownership; but this has not been nearly dramatic enough to account for a more than 3x increase in housing capital income.[9] Another distinct source of bias could be rent control: if the rents imputed for homeowners in the national accounts improperly reflect controlled rents in the tenant-occupied sector, then the ebb and flow of rent regulations will have an inflated impact on income in the housing sector as a whole.

These possible biases notwithstanding, the main thrust of figure 3 is that housing has a pivotal role in the modern story of income distribution. Since housing has relatively

[8]For Canada and Japan, the "housing" sector is actually the owner-occupied housing sector due to data limitations. Importantly, Canada and Japan do not drive the trend here: to the contrary, from 1960 (when Canada enters the sample) to 2010, the average contribution of housing to net capital income in Canada and Japan increases by 3pp, while in France, the UK, and the US it increases by 4.5pp.

[9]See, e.g., Andrews and Sanchez (2011) for some discussion of trends in homeownership.

broad ownership, it does not conform to the traditional story of labor versus capital, nor can its growth be easily explained with many of the stories commonly proposed for the income split elsewhere in the economy—the bargaining power of labor, the growing role of technology, and so on.

The divergence between housing and other forms of capital is also hard to reconcile with the accumulation view: in the Piketty (2014) narrative, for instance, it is not clear why a rise in the aggregate wealth-to-income ratio should be channeled entirely into a rise in the housing component of the net capital share, while the non-housing component stagnates.[10]

Net capital share within the corporate sector. For additional clarity, figure 4 plots the average net capital share within the *corporate* sector. Restricting attention to the corporate sector is a common way to deal with perceived conceptual and measurement difficulties elsewhere in the economy—including ambiguity in the labor/capital split of mixed income, as well as the crucial role of housing. Figure 4 echoes the behavior of the non-housing component in figure 3, with a sustained fall until the 1970s and a partial recovery in the decades since. Indeed, this resemblance is no coincidence: as discussed above, figure 3 imputes the net capital share in the non-housing, non-corporate sector to be the same as in the corporate sector, so that movement in all non-housing capital income is fundamentally driven by the corporate capital share visible in figure 4.[11]

Although it does not show any decisive, long-term trend, figure 4 does clash with the Kaldor (1957) view of stable income shares. It also contrasts with the relatively steady upward creep of housing capital income in figure 3. Fluctuations in the average corporate capital share have been rapid and macroeconomically significant—dropping from a high around 26% in 1950 to a trough around 18% in the 1970s and 80s, then rebounding to a peak of 24% in the 2000s. Indeed, the fall in the unweighted average share from 26.4% in 1950 to 17.7% in 1975, all else equal, contributed nearly half a percentage point annually to growth in corporate labor compensation during that interval. In contrast, the rapid rise from 17.7% in 1975 to 23.6% in 1988 *subtracted* slightly more than half a percentage point of annual compensation growth.[12]

Yet over the long term, the role of fluctuating corporate income shares is compara-

[10] In fact, computations using the multisector model in section 5.2 will show that this is backward: the fall in r induced by savings should lead to a concentrated *decline* in the housing component of the net capital share.

[11] As described in footnote 7, different imputations are used for Canada and Japan. Furthermore, since separate data for the corporate sector is not available in Japan, figure 4 displays the overall capital share for Japan instead.

[12] Explicitly, $((1-.177)/(1-.264))^{(1/25)}-1 \approx .45\%$ and $((1-.236)/(1-.177))^{(1/13)}-1 \approx -.57\%$.

tively quite mild. For both the weighted and unweighted averages, the impact on annual compensation growth from the 1948-2010 change in net shares is roughly *three-hundredths* of a percentage point.[13] The overall message is clear, and arguably consistent with the Kaldor (1957) perspective on long-run growth: changes in the distribution of corporate income—even systematic ones spread across several countries—can have a marked effect on the short-to-medium-run growth of paychecks. The impact on long-run labor compensation, however, appears to be little more than a rounding error when set against trend growth.[14]

There is also a pronounced cyclical pattern in figure 4. This is has long been recognized: the labor share tends to rise late in expansions and fall late in recessions. The economic explanation for this pattern, however, is somewhat harder to discern. Conventional wisdom is that low unemployment puts upward pressure on real wages and hence the labor share, while high unemployment keeps real wage growth subdued. This story, however, implicitly involves variation in markups: as Mitchell (1941) observes, "a problem still remains: Why cannot businessmen defend their profit margins against the threatened encroachment of costs by marking up their selling prices?" Answering this challenge, the business cycle literature offers an abundance of proposed explanations for the cyclical pattern of markups, of which Rotemberg and Woodford (1999) provides an excellent summary.

3 Decomposing the capital share

3.1 Bringing in the value of capital

Section 2 provided some preliminary insights into the structure of the net capital share, by distinguishing it into housing and non-housing components. It found that the housing component has seen a steady increase, while the non-housing component has experienced a dramatic fall and then rise. To better understand these movements, it is important to look at another piece of evidence: the value of the capital stock itself.

[13] Explicitly, for unweighted: $((1-.214)/(1-.229))^{(1/62)}-1 \approx .03\%$. For weighted: $((1-.231)/(1-.245))^{(1/62)}-1 \approx .03\%$.

[14] To be clear, the long-run impact in individual countries can be larger. Perhaps the most extreme example is Japan, which table 2 shows to have experienced a decline in the average non-housing share of aggregate capital income from 31% in the 1960s to 20% in the 2000s, implying an annualized contribution to wage growth of roughly three-tenths of a percentage point. But Table 3 does not suggest any long-run tendency for corporate capital shares in different countries to diverge from each other; the distinct paths across countries are therefore probably best interpreted as mean-reverting variations around an apparently trendless average.

Both the Piketty (2014) and Karabarbounis and Neiman (2014b) versions of the accumulation view, for instance, explain the recent rise in the capital share through a rise in the value of reproducible capital relative to aggregate income. In fact, this is a central feature of virtually any narrative that stresses capital accumulation: if capital is earning a larger share because we are building more of it, then data on the value of capital should reveal that it has indeed grown relative to income.

Furthermore, this should be true *within* sectors: for instance, if accumulation explains the rise in the non-housing capital share over the last few decades, then we should see a rising value of capital within the non-housing sector, relative to sectoral value added. This is a simple but crucial check. Elaborating upon it, we can try to disentangle the roles of different influences on the capital share—the observed value of capital itself; the net user cost of that capital; and firms' markups over cost that lead to additional capital income, not attributable to the user cost of the measured capital stock.

Theory. Formally, let $K_1, \ldots, K_n$ be different types of capital, and let $Y = F(N, K_1, \ldots, K_n)$ be a constant-returns-to-scale production function that takes labor N and capital $K_1, \ldots, K_n$ as inputs. Suppose that output Y is sold at a price P that represents a markup of $\mu \geq 1$ over the cost of production[15], such that the share of what I will call "pure profits"—capital income above and beyond the user cost of capital $K_1, \ldots, K_n$—in gross income is $\pi \equiv 1 - \mu^{-1}$. I allow for a potentially time-varying markup μ in part because of the discussion of the corporate capital share in section 2.2, which notes a pronounced cyclical pattern that has been explained in the literature through markup variation.

Letting W_N denote the wage paid to labor and $W_{K_1}, \ldots, W_{K_n}$ denote the user costs of capital, we have

$$(1 - \pi)PY = W_N N + \sum_{i=1}^{n} W_{K_i} K_i \tag{1}$$

Suppose further that the model is cast in continuous time (suppressing time subscripts for convenience), and that the flow real cost of funds is r. Capital K_i has real price P_i, with expected real growth rate g_{P_i}, as well as a flow depreciation rate of δ_i. The user cost W_{K_i} is then

$$W_{K_i} = P_i(r + \delta_i - g_{P_i}) \tag{2}$$

reflecting the real cost $P_i r$ of financing each unit of capital and the expected combined effect $P_i(\delta_i - g_{P_i})$ of depreciation and price growth on the value of capital held.

Combining (1) with (2), we see that we can divide net output into labor income $W_N N$

[15]Since F is constant-returns-to-scale, marginal and average costs are equal, so I will refer to them both as "cost".

and net capital income; the latter can further be divided into a share πPY of profits and a component $(r - g_{P_i})P_iK_i$ corresponding to each type of capital i.

$$\underbrace{PY - \sum_i^n \delta_i P_i K_i}_{\text{net output}} = W_N N + \underbrace{\pi PY + \sum_{i=1}^n (r - g_{P_i}) P_i K_i}_{\text{net capital income}} \tag{3}$$

Letting Y^{net} denote net output on the left of (3), we can divide through by Y^{net} to write (3) in terms of shares:

$$1 = \underbrace{W_N N / Y^{net}}_{\text{net labor share}} + \underbrace{\pi PY / Y^{net} + \sum_{i=1}^n (r - g_{P_i})(P_i K_i / Y^{net})}_{\text{net capital share}} \tag{4}$$

(4) illustrates formally how we can divide the net capital share into components that reflect the ratio P_iK_i/Y^{net} of the value of capital of each type i to net income. As discussed earlier, this allows us to evaluate a central element of the accumulation view—namely, that changes in P_iK_i/Y^{net} have played a key role in the evolution of the net capital share.

Discussion of implementation. Suppose that in practice we have disaggregated capital into n types, for which we have data on the value P_iK_i, and we want to divide the observed capital share of net income Y^{net} into the components identified in (4). First, expected price growth g_{p_i} is needed; this is very difficult to obtain in principle, since we rarely observe agents' individual expectations of price growth, but it can be roughly approximated by assuming that g_{P_i} matches the trend rate of growth over some interval.

The most difficult parts of (4) are π and r: with knowledge of one, we can infer the other, but neither is readily available in the data. In principle, r could be obtained from financial markets, perhaps as some function of bond and equity prices. But this is a notoriously hard problem: it is challenging to know how exactly the costs of borrowing or equity finance map onto the effective cost of funds faced by an enterprise. Furthermore, since this r is pretax while returns on bonds or equity are after corporate taxes, a time-varying tax adjustment would be needed to infer r directly from market returns.

3.2 Implementation: decomposing the net corporate capital share in the US, 1948–2013

I first attempt the disaggregation in (4) for the net capital share in the US corporate sector, at an annual frequency for the years 1948 through 2013.[16] I disaggregate fixed capital into its three most important components: structures, equipment, and land (denoted by $i = s, e, l$), and I obtain the values $P_i K_i$ for the corporate sector from the flow of funds.[17] I assume that the expected price growth g_{P_i} of each form of capital is its actual average real price change from the end of 1947 to the end of 2013. I then try several approaches to resolving the difficulties identified at the end of section 3.1.

Evaluating the accumulation view: assuming constant r. One way to implement the decomposition in (4) is to simply impose constant r. Taken literally, this is probably not a viable assumption, but it *is* a straightforward approach to testing the accumulation view: if we rule out variation in r as a source of change in (4), how much of the time series can $P_i K_i / Y^{net}$ itself explain? How well do movements in $P_i K_i / Y^{net}$ correlate with changes in the net capital share, and what role can they play quantitatively when r is chosen to be of reasonable size?

Note that in this exercise, the "pure profit" term $\pi PY / Y^{net}$ is effectively just a residual. The goal, for now, is not to provide a complete and convincing decomposition of the net capital share into changes in π, r, and $P_i K_i / Y^{net}$; but instead, to see what role $P_i K_i / Y^{net}$ alone can play. This exercise, though similar, is more informative than mere inspection of the paths of $P_i K_i / Y^{net}$ relative to the path of the net capital share, because it provides some indication of *magnitude*. For instance, if $P_i K_i / Y^{net}$ moves together with the net capital share for most i, this pattern would appear consistent with the accumulation view; but to see whether this support is quantitatively viable, it is necessary to map the changes in $P_i K_i / Y^{net}$ onto their contributions to the net capital share. This is the role of (4), together with some choice of constant r.

First, I assume that r takes a constant value over the sample period 1948–2013 such that the average profit share π of corporate revenue over the sample is zero. This implies $r \approx 11\%$.[18] Effectively, the assumption here is that in the long run, there are no pure

[16] Ideally, this exercise would extend to all seven of the G7 countries covered in section 2, but the additional data required makes this difficult.

[17] Since the flow of funds provides end-of-year values for capital, I average the adjacent end-of-year values to obtain the effective capital stock used in production during each year.

[18] Although this seems high for a real return, note that it is a pretax return: the return before taxes are applied either to corporate profits or distributions of interest or dividends. Interestingly, it is slightly lower than the constant return in figure 7 estimated using my alternative approach, which is roughly 12.8%. As

profits in the corporate sector—on average, net capital income reflects a return on equipment, structure, or land. This is consistent with Chamberlinian monopolistic competition, where entry drives monopoly profits to zero, on average, in the long run.

Figure 6 shows how the net capital income for the US corporate sector in figure 5 breaks down into the four components in (3) under this assumption. Though there are some fluctuations in each component's contribution, both the U-shaped pattern and the cyclical fluctuations in the corporate capital share in figure 5 appear dominated by the residual component of "pure profits" π. In other words, contrary to the accumulation view, time series shifts in the capital share in the corporate sector cannot be explained by parallel shifts in the measured value of capital.

Consequences for the falling investment prices hypothesis. As figure 5 further reveals, the contribution from equipment in particular is if anything the *inverse* of the U-shaped pattern in the corporate net capital share in figure 5: it rises in the 1970s and 1980s, and then later trends downward. Since equipment is the component of fixed capital that has experienced a decline[19] in real price, this is hard to reconcile with a central role for falling investment prices in the dynamics of capital's share, the hypothesis emphasized by Karabarbounis and Neiman (2014b). Without a structural model, of course, this exercise is not decisive: falling investment prices might contribute to a rising capital share via some more indirect causal channel, and indeed Karabarbounis and Neiman (2014a) suggest one such possibility. I address these concerns with a multisector model in section 5.2, where I generally do not find a major role for such indirect mechanisms.

Surprisingly, my finding here is consistent with the result of a closely related exercise in section IV.B of Karabarbounis and Neiman (2014b), who also decompose non-labor income into a component reflecting the return on accumulated capital and a component reflecting markups, under the assumption of a constant real interest rate. Although it is not their focus, Karabarbounis and Neiman (2014b) remark that they generally do not find increases in the share of the former component. This implies that the fall in labor share comes in the aggregate from the rise in markups, rather than returns on measured capital.

At face value, this contradicts the emphasis on capital accumulation as a source of the falling labor share; but Karabarbounis and Neiman (2014b) point out that if their elastic-

explained below, this return is higher because according to the flow of funds, the total market value of the corporate sector in the US has actually been lower than the book value, on average, in the postwar era—suggesting that pure profits are, if anything, negative, and that the assumption that pure profits are zero on average is not misattributing these profits to an exaggerated return r on capital.

[19] The equipment investment deflator has relative to the GDP deflator at an annualized rate of 1.5% during the sample period, as opposed to a 1.1% average rise in the deflator for nonresidential structures.

ity estimate is valid, it remains correct to say that *counterfactually*, the labor share would be higher if not for the role of falling investment prices in encouraging investment. Of course, if this is true, it follows that there must be two other unidentified forces influencing the factor distribution of income: (1) some force of similar magnitude that offsets their mechanism in the aggregate by pushing investment downward, and (2) another force leading to the rise in markups, which accounts for the entire aggregate fall in the labor share. With these forces in play, the accumulation view only plays a secondary role regardless.

Smaller r. Figure 5 can also be constructed assuming a smaller r, under the assumption that pure profits in the corporate sector are not all dissipated in the long run. This does not materially change the conclusion that the measured value of capital is unable to account for the major shifts in the net capital share. (Indeed, a smaller r in (4) directly leads to a lower weight on $P_i K_i / Y^{net}$.)

Structural approach: identify time path for r from market minus book value. In an attempt to more convincingly disentangle the roles of r and π, I turn to a more elaborate approach for estimating r. The basic idea is that the difference between the market value of corporations and the value of their fixed assets should reflect the expected stream of future pure profits πPY (up, possibly, to some stochastic pricing error). We can use this observation as a strategy to estimate the implied r: for instance, if market value is much higher than the value of the firm's assets, the expected stream of pure profits πPY is high, and r in the future must be low enough that there are pure profits left over in (3) after the direct return from capital $\sum_i (r - g_{P_i}) P_i K_i$ is subtracted.[20]

Description of the method. Appendix C provides the technical details, along with the specific theoretical assumptions in a continuous-time model that are needed to make the procedure valid. The core equation implied by the theory is (see (32) and (33)):

$$\mathbb{E}\left[\frac{\phi(t)}{\text{output between } t-1 \text{ and } t} \times \left(OMV(t) - \text{discount} \times OMV(t+1)\right)\right]$$
$$= \mathbb{E}\left[\frac{\phi(t)}{\text{output between } t-1 \text{ and } t} \times \text{pure profits between } t \text{ and } t+1\right] \quad (5)$$

[20]For simplicity, I will call the total value of the firm's fixed assets its "book value", even though this is not necessarily book value in the usual sense: I will define it to exclude financial assets—these are instead subtracted from the market value, which includes net financial liabilities—and it uses values from the flow of funds for real estate and equipment, which are updated to reflect changes in price.

where $OMV(t)$ denotes the difference between the market value and book value of corporations recorded at time t, and $\phi(t)$ is an arbitrary time-dependent function. Implicit in (5) is a (nonstochastic) time path $r(t)$ for the real interest rate, which is needed to calculate profits $\pi(t)P(t)Y(t)$ as a residual in (3) and to calculate the proper discount factors.

The interpretation of (5) is straightforward: it states that the expected difference between the present value of next year's excess market value $OMV(t+1)$ and this year's excess market value $OMV(t)$ reflects expected pure profits between t and $t+1$. This relation continues to hold, in expectation, when both sides are normalized by the previous year's recorded output, which I do to render values comparable across time; and it also holds when both sides are multiplied by any choice of the time-dependent function $\phi(t)$.

Technically speaking, equation (5) can be used as a moment condition to estimate $r(t)$. If we have n functions $\{\phi_1(t), \ldots, \phi_n(t)\}$, we obtain n distinct moment conditions (5), and can enforce these conditions in the sample to solve for an n-parameter functional form for $r(t)$. I choose $\phi_1(t) = 1$, $\phi_2(t) = t$, and $\phi_3(t) = t^2$, and estimate three specifications for $r(t)$: a constant value $r(t) = \bar{r}$, a linear trend $r(t) = a_0 + a_1 t$, and a quadratic trend $r(t) = a_0 + a_1 t + a_2 t^2$, using the moment conditions implied by $\{\phi_1(t)\}$, $\{\phi_1(t), \phi_2(t)\}$, and $\{\phi_1(t), \phi_2(t), \phi_3(t)\}$, respectively.

Effectively, I am solving for the constant $\bar{r}$ such that the expression

$$\frac{OMV(t) - \text{discount} \times OMV(t+1) - \text{pure profits between } t \text{ and } t+1}{\text{output between } t-1 \text{ and } t} \tag{6}$$

equals zero on average throughout the sample; and I am also solving for the linear $r(t) = a_0 + a_1 t$ and the quadratic $r(t) = a_0 + a_1 t + a_2 t^2$ such that (6) does not have any linear or quadratic trends over time, respectively.

When calculating OMV, the difference between the market value of the corporate sector and the book value of its fixed capital, I interpret the "market value" to be the total value of all financial claims on a corporation—both its equity market capitalization and its net financial liabilities—in order to be consistent with the computation of capital income in the national accounts, which includes income that ultimately goes to both shareholders and bondholders.[21] Both market and book value are taken from the flow of funds.

[21] This causes some anomalies in the early postwar years, when the corporate sector was left with large cash balances and relatively little debt, making net liabilities negative while equity valuations were already quite low, and leading to an extremely low market relative to book value. To avoid undue influence from this period, I exclude data from prior to 1955 in the benchmark results displayed here; otherwise, there is an even more dramatic estimated downward trend in $r(t)$.

Estimated paths for r. Figure 7 shows the estimated constant, linear, and quadratic time trends for the corporate rate of return $r(t)$ following the procedure above. The most striking feature of these plots is the general downward trend in $r(t)$: according to this procedure, the required return on capital for the US corporate sector has fallen over the postwar era. This reflects the fact that the market value of corporations has grown relative to book value over this period, albeit unevenly, as can be seen in figure 9. The estimation infers from this that pure profits are trending upward, so that the required return on capital $r(t)$ itself must be declining.

Another interesting feature of figure 7 is that estimated constant $\bar{r}$, at roughly 12.8%, is actually *higher* than the r chosen in my benchmark decomposition to set the average share of pure profits to zero. This reflects the fact that according to the flow of funds, on average, the aggregate market value of corporations has actually been slightly below the book value during the sample period, as depicted in figure 9. This suggests that the assumption of zero average pure profits for the benchmark decomposition was not too far out of line: corporations, on average, have not been worth more than the underlying value of their assets.

Since I am only estimating parametric trends for $r(t)$ here, I am not allowing $r(t)$ to vary at high frequencies with the business cycle; market prices at high frequencies are too noisy and volatile to permit credible estimation of $r(t)$ using the method above. This means that I still cannot address, for instance, the role played by cyclical fluctuations in $r(t)$ in driving cyclical fluctuations in the capital share. But by allowing for a long-term trend in $r(t)$, I can disentangle the long-term effects of r from the effects of changing capital-income ratios P_iK_i/Y^{net}, and obtain a better assessment of the role of pure profits $\pi PY/Y^{net}$.

Since long-term trend in the corporate net capital share is U-shaped, with a large fall and recovery, I will emphasize the results from the *quadratic* estimated trend $r(t)$. To the extent that varying r is partly responsible for the U-shaped trend, quadratic $r(t)$ can capture much of its impact.

Implications of quadratic trend in r. Redoing the decomposition in figure 6, using the quadratic trend for $r(t)$ rather than a constant, produces figure 8. The impact of the change in $r(t)$ is unsurprising. Relative to figure 6, figure 8 initially attributes a larger share of returns to fixed capital, offset by substantial negative pure profits; over time, the return on fixed capital falls, and the role of pure profits grows substantially. As in figure 6, pure profits play a central role in the U-shaped path for the overall corporate net capital share—but these movements come in addition to broad offsetting trends, in which pure profits have replaced income from fixed assets in (3) .

It is difficult to say how literally these trends should be interpreted. Given the methodology for identifying $r(t)$, they are ultimately the consequence of the long-term rise in the ratio of market value to book value in the US corporate sector, as seen in figure 9; and this, in turn, may be the result of other, unmodeled changes in financial markets, not a rise in π. Nevertheless, figure 8 is certainly suggestive, and it casts additional doubt on the accumulation view, since it indicates that (contrary to the assumption of relatively stable returns per dollar of capital) $r(t)$ has, if anything, experienced a sizable decline.

3.3 Extending the decomposition: the net capital share for the private economy

I now extend the decomposition in section 3.2 to the net capital share for the private domestic economy as a whole—excluding the non-housing government and NPISH (non-profit institutions serving households) sectors, which have zero net capital share by construction in the national accounts.

Due to the inherent difficulties in apportioning mixed income between labor and capital, as discussed in section 2.1, this requires some imputations. I will assume that both the rate of return r and the pure profit share π are the same in the non-housing, non-corporate sector and the corporate sector, and use the estimated quadratic path for r from the previous section.[22] For the housing sector, I will assume that there is no pure profit, and allow r to vary over time in (3) such that net housing capital income always equals $(r - g_{P_{s2}})P_{s2}K_{s2} + (r - g_{P_{l2}})P_{l2}L_2$, where $P_{s2}K_{s2}$ is the value of residential structures and $P_{l2}L_2$ is the value of residential land.

The results are displayed in figure 11, which decomposes the net capital share displayed in figure 10. Figure 11 is noisy, and for the most part it combines the lessons from sections 2.2 and 3.2: there is a strong, long-term upward trend in net capital income from housing, and the volatile capital share elsewhere in the economy is driven principally by pure profits.

There are, however, some additional insights in the figure 11 decomposition. For instance, the rise in net income for the housing sector has come both from residential structures and land, but figure 11 attributes a larger portion of the increase (and of the level) to structures.

[22] Note that this imputation, which uses data on the value of fixed assets in the non-corporate sector, is different from the imputation in section 2.2, where this data was not available for the full sample and the net capital *share* of income—rather than the return r—in the non-housing, non-corporate sector was assumed to be the same as in the corporate sector.

This may come as a surprise, since one plausible hypothesis for the growth of net housing income is the rising scarcity of land. In part, the secondary role of residential land here comes from its more rapid price appreciation. Since I assume that the net rate of return *including* expected capital gains is equalized between residential structures and land, the net rate of return *excluding* expected capital gains—which is used in the decomposition, because income in the national accounts also excludes capital gains—is significantly lower for land. In a sense, then, the lesser role of land is due to the idiosyncrasies of national accounting; and an alternative definition of net capital income that included some form of expected capital gains would show a larger impact from land. (With this in mind, it is remarkable that housing plays such a large aggregate role in section 2.2 already: if the G7 national accounts data were modified to include capital gains, housing's centrality would only increase.)

Another interesting feature of figure 11 is that there has been a sizable decline in the role of capital income from non-residential land over time, from roughly 10% of net private value added in the first half of the sample to an (erratic) average of roughly 2.5% today. In other words, there has been a shift in net capital income from non-residential to residential land—but the decline in the former has been far larger than the growth in the latter, suggesting that the direct contribution of land to net capital income in the US has actually fallen.

4 Capital share theory: one-sector model

4.1 One-sector, one-good model.

I now take a step back from the decomposition in section 3—with its multiple capital goods—to recount the simplest, traditional model of income shares, with a single production sector and a single good. This offers a first-pass test of the *theoretical* viability of the accumulation view: all else equal, should we expect a larger capital-income ratio to cause an increase or decrease in capital's share?

Let $F(K, N)$ be a constant returns to scale production function, with capital K and labor N as factor inputs, and positive but diminishing returns in each factor. Assume that this is a one-good model, where the relative price of capital and output is fixed at one. The **elasticity of substitution** σ between K and N is defined as

$$\sigma \equiv -\left(\frac{d(\log(F_K/F_N))}{d(\log(K/N))}\right)^{-1} \tag{7}$$

This gives us the (inverse) elasticity of the ratio F_K/F_N of marginal products to the ratio K/N of capital. Equivalently, σ tells us the extent to which a cost-minimizing producer's relative demand for K/N will change if there is a change in the relative cost R/W of using capital and labor as inputs.

From the definition (7), one can show that σ also gives the inverse elasticity of F_K with respect to a change in the capital-output ratio K/F:

$$\sigma = -\left(\frac{d(\log F_K)}{d(\log(K/F))}\right)^{-1} \tag{8}$$

which implies that the elasticity of the capital income *share* $F_K K/F$ with respect to the capital-output ratio K/F is

$$\frac{d(\log(F_K K/F))}{d(\log(K/F))} = 1 - \frac{1}{\sigma} \tag{9}$$

This indicates the critical importance of the threshold $\sigma = 1$. If $\sigma > 1$, the elasticity is positive, so that the capital income share will increase as K/F rises. Inversely, if $\sigma < 1$, the capital income share will fall as K/F rises. In the important special case $\sigma = 1$, diminishing returns exactly offset the increased quantity of capital, and the share remains constant.

Indeed, one of the original motivations behind Cobb and Douglas (1928)'s eponymous production function was the apparent constancy of capital and labor shares in the data; this is guaranteed by the Cobb-Douglas production function $F(K, N) = K^\alpha N^{1-\alpha}$, which has a constant elasticity of substitution $\sigma = 1$.

Net versus gross. Thus far, I have been ambiguous about whether the function F gives *gross* production, or production *net* of depreciation. In principle, either interpretation is legitimate—especially since this is a one-good model, where the relative price of capital and output is fixed at one, and losses from capital depreciation can reasonably be included as part of the production function.

If F is *gross* production, then $1 - 1/\sigma$ is the elasticity of *gross* capital income with respect to the ratio of capital to *gross* output. If F is *net* production, then $1 - 1/\sigma$ is the elasticity of *net* capital income with respect to the ratio of capital to *net* output. As discussed in section 2.1, both measures are useful, but net concepts are probably more meaningful when studying income distribution.

It is important to recognize that σ depends greatly on which measure is used—a subtlety that is often overlooked. Suppose $F(K, N)$ is the gross production function, with an elasticity of substitution of σ. Then the net production function is $\tilde{F}(K, N) =$

$F(K, N) - \delta K$, and from (8) the elasticity of substitution for $\tilde{F}$ is

$$\begin{aligned}
\tilde{\sigma} &= \frac{d(\log(\tilde{F}/K))}{d(\log \tilde{F}_K)} \\
&= \frac{d(F/K - \delta)/(F/K - \delta)}{d(F_K - \delta)/(F_K - \delta)} \\
&= \frac{d(F/K)/(F/K)}{d(F_K)/F_K} \cdot \frac{(F_K - \delta)/(F - \delta K)}{F_K/F} = \sigma \cdot \frac{(F_K K - \delta K)/(F - \delta K)}{F_K K/F} \qquad (10)
\end{aligned}$$

Hence the elasticity of substitution $\tilde{\sigma}$ for the net production function ("net elasticity") equals the elasticity of substitution σ for the gross production function ("gross elasticity") times the ratio of the net capital share $(F_K K - \delta K)/(F - \delta K)$ and the gross capital share $F_K K/F$. Since the net capital share is always less than the gross capital share, it follows that the net elasticity is always below the gross elasticity.

Why, intuitively, is the net elasticity always lower? The net return on capital $\tilde{F}_K$ is less than the gross return F_K by a constant—the depreciation rate δ—meaning that a given change in F_K translates into an equal absolute, and a larger *relative*, change in $\tilde{F}_K$. For instance, if $\delta = 5\%$, and F_K declines from 10% to 8%, $\tilde{F}_K$ will decline from 5% to 3%. A 20% decline in the gross return becomes a 40% decline in the net return, and the ratio of the two is (A). As we increase capital relative to labor, the net marginal product of capital declines more rapidly than the gross—in short, capital is *less substitutable* for labor from a net perspective.

Calibrating (10). To obtain an illustrative calibration, I take the data from section 3.3, where pure profits are estimated using the quadratic path for $r(t)$. I exclude pure profits and land from the capital share, since they are not reproducible forms of capital—and relevant question for the Piketty (2014) hypothesis is whether adding more reproducible capital through investment increases or decreases capital's share of income.

In the most recent year in the sample, 2013, the resulting US private net capital share (excluding pure profits and land) was 25.6%, while the US private gross capital share (excluding pure profits and land) was 34.5%. This results in a ratio of approximately 0.74, and (10) implies

$$\tilde{\sigma} \approx 0.74 \times \sigma \qquad (11)$$

so that the net elasticity is slightly less than three-quarters the gross elasticity.

If the decomposition in section 3.3 is performed assuming a lower rate of return r, such that a more significant share of net capital income is attributed to pure profits rather than returns on measured capital, then the ratio can be appreciably lower than in (11). For

Table 1: Distribution of elasticity estimates compiled by Chirinko (2008): in gross terms (as originally stated) and in net terms (converted using (11)).

	$[0, 0.5)$	$[0.5, 1)$	$[1, 1.5)$	$[1.5, 2)$	$[2, 4)$
Frequency of gross σ	14	12	3	1	1
Frequency of net $\tilde{\sigma}$	21	8	1	0	1

instance, in an alternative estimate where r is chosen to be roughly 5.5% for the corporate sector—implying that half of long-run net capital income is attributable to pure profits—the private net and gross capital shares (excluding pure profits and land) become 15.5% and 25.6%, respectively, resulting in a ratio of approximately 0.60.

Empirical implications. Ever since Arrow, Chenery, Minhas and Solow (1961) first proposed the constant elasticity of substitution (CES) production function, researchers have attempted to estimate the key elasticity parameter. These studies have virtually always looked at the elasticity of substitution in the *gross* production function.

The literature is vast and its conclusions muddled, but one consistent theme has been the rarity of high elasticity estimates. Chirinko (2008) provides an excellent summary of the empirical literature, listing estimates from many different sources and empirical strategies; table 1 displays the estimates compiled there, both in their original gross terms and converted to net terms, where the conversion factor of 0.74 from (11) is used.

Of the 31 sources[23] listed, table 1 reveals that only 5 show a gross elasticity above 1, and only 2 imply a net elasticity above 1. From (9), it follows that a rise in the capital-income ratio, holding the production function constant, most likely will cause a decline in the net share of capital income. This is inconsistent with the Piketty (2014) and Piketty and Zucman (2014) version of the accumulation view, which holds that a rise in the capital-income ratio has led—and will lead going forward—to a rise in capital's net share.

Implications for r-g. A closely related theme in Piketty (2014) is the gap $r - g$ between the real return r on capital and the real growth rate g of the economy. This gap, for instance, gives the rate at which a wealthy dynasty can withdraw capital income for consumption purposes without decreasing its wealth relative to the size of the economy. More generally, when $r - g$ is higher, "old" accumulations of wealth become more important relative to "new" ones. Higher $r - g$ generally implies that the power law tail of the wealth distribution has a smaller exponent—so that there is more inequality of wealth at

[23]For a few sources that list a range of elasticities, I take the midpoint. This has minimal effect on the distribution.

the top, and extreme levels are more likely. Many readers take the dynamics of $r - g$ to be *the* central theme of the book.

Both Piketty (2014) and Piketty and Zucman (2013) make heavy use of the identity

$$\frac{K}{Y^{net}} = \frac{s}{g} \tag{12}$$

where s is the net savings rate and g is the growth rate, which is dubbed the "Second Fundamental Law of Capitalism". This identity only holds asymptotically—if s or g changes, convergence to the new value of K/Y^{net} does not happen instantaneously—and it is unlikely that s is exogenous and invariant to changes in g. Nevertheless, Piketty (2014) argues that it is useful to explore the implications of this identity given exogenous s, particularly the fact that K/Y^{net} rises as g falls, which is central to the projection that the capital-income ratio will rise in the future.[24]

If $r = F_K$, then (8) shows that the elasticity of r with respect to K/Y^{net} is simply $-\tilde{\sigma}^{-1}$, where $\tilde{\sigma}$ is the net elasticity of substitution. For exogenous s, (12) indicates that the elasticity of K/Y^{net} with respect to g is -1, implying that the elasticity of r with respect to g is $\tilde{\sigma}^{-1}$. It follows that

$$\frac{\partial(r-g)}{\partial g} = \frac{r}{g}\tilde{\sigma}^{-1} - 1 \tag{13}$$

This expression is positive if $r/g > \tilde{\sigma}$. Again taking data from section 3.3, in 2013 the average return on measured capital was 7.5%. Taking this to be the r in (13), and taking g to be 2.5% (approximate trend real GDP growth in the US in the last 25 years), we have $r/g = 3$, in which case the derivative in (13) is positive as long as the net elasticity $\tilde{\sigma}$ is less than 3.

The evidence in table 1 indicates that this is overwhelmingly likely. Indeed, converting via (11), a net elasticity of 3 corresponds to a gross elasticity of $\sigma = \tilde{\sigma}/0.74 = 4.05$, which is above every estimate listed in Chirinko (2008) and virtually every estimate in the wider literature. Hence a decline in g will result in a *decline* in $r - g$: the decline in g itself is less than the decline in r that it induces through capital accumulation and diminishing returns. Given the assumption (12) on capital accumulation, the prediction in Piketty (2014) that $r - g$ will rise as g falls is especially hard to reconcile with empirically plausible degrees of substitutability.

[24]There is some conflict between the assumption of exogenous s for all income and the emphasis on $r - g$. If only this fraction s of capital income r is saved, then existing fortunes will grow at the rate $s \cdot r - g$, not $r - g$; and for plausible values of s as a share of all income, $s \cdot r - g$ is likely to be quite negative, implying the rapid erosion of existing wealth.

4.2 One-sector, two-good model.

The canonical model in section 4.1 can be enriched slightly by allowing the price P_K of capital relative to the output good to vary. This modification is central to the account in Karabarbounis and Neiman (2014b), who attribute the rise in the gross capital share to high capital demand induced by a fall in P_K.

To be more explicit, take the net required return r on capital as given. Ignoring expected capital gains, demand for capital is pinned down by the condition

$$F_K(K,N) = P_K(r+\delta) \tag{14}$$

The elasticity of the gross capital/output ratio K/F with respect to P_K is then

$$\frac{\partial(\log(K/F))}{\partial(\log P_K)} = \frac{d(\log(K/F))}{d(\log F_K)} \cdot \frac{\partial(\log F_K)}{\partial(\log P_K)} = -\sigma \tag{15}$$

where $\partial(\log F_K)/\partial(\log P_K) = 1$ follows directly from (14), and $d(\log(K/F))/d(\log F_K) = -\sigma$ follows from (8). The elasticity of the gross capital share with respect to P_K becomes

$$\frac{\partial(\log(F_K K/F))}{\partial(\log P_K)} = \left(1 + \frac{d(\log(K/F))}{d(\log F_K)}\right) \cdot \frac{\partial(\log F_K)}{\partial(\log P_K)} = 1-\sigma \tag{16}$$

implying that a decline in the relative price P_K of capital will increase the gross capital share if $\sigma > 1$.

Meanwhile, the elasticity of the *net* capital share with respect to P_K is can be obtained through a somewhat more involved computation.[25] The result, first derived byKarabarbounis and Neiman (2014a), is

$$\frac{\partial(\log((F_K - \delta P_K)K/(F-\delta P_K K))}{\partial(\log P_K)} = (1-\sigma)\cdot\frac{F}{F-\delta P_K K} \tag{17}$$

Note that $\sigma = 1$ is still the critical threshold: a decline in the relative price P_K of capital increases both the net and gross capital shares if $\sigma > 1$. This consistency is a noteworthy contrast with the distinction (10) between gross and net elasticities of substitution, where

[25] For instance, one can write

$$\begin{aligned}\frac{\partial(\log((F_K-\delta P_K)K/(F-\delta P_K K))}{\partial(\log P_K)} &= \frac{\partial(\log(F_K-\delta P_K))}{\partial(\log P_K)} + \frac{\partial(\log(K/F))}{\partial(\log P_K)} - \frac{\partial(\log(1-\delta P_K K/F))}{\partial(\log P_K)}\\ &= 1-\sigma + \frac{\partial(\log(\delta P_K K/F))}{\delta(\log P_K)}\frac{\delta P_K K/F}{1-\delta P_K K/F} = (1-\sigma) + (1-\sigma)\frac{\delta P_K K}{F-\delta P_K K} = (1-\sigma)\frac{F}{F-\delta P_K K}\end{aligned}$$

a rise in the capital-output ratio could produce an increase in the gross capital share and a decrease in the net capital share. From an intuitive standpoint, this is unsurprising: since we are holding r constant, the ratio $r/(r+\delta)$ of net to gross capital income is fixed, and the two move in the same direction in response to a change in P_K.

Karabarbounis and Neiman (2014a) stress the role of (17), which shows that their focus on the role of changes in P_K can potentially account for simultaneous changes in *both* the gross and net capital shares, assuming that the gross elasticity of substitution σ is greater than 1. In light of the estimates compiled in table 1 (26 out of 31 of which find $\sigma < 1$), $\sigma > 1$ still appears unlikely, but it is somewhat more plausible than $\tilde{\sigma} > 1$.

5 Capital share theory: a multisector model

5.1 Design of the multisector model

The theory in section 4 enables a first-pass analysis of how the distribution of income is affected by various forces. It shows that accumulation of capital—all else equal—will likely result in a *decline* in the net capital share, since the net elasticity of substitution is almost certainly below one. This counters the central hypothesis of Piketty (2014). It also shows that a decline in the relative price P_K of capital, holding the required return r constant, will result in an increase in the net capital share if the gross elasticity of substitution is above 1—a claim that is still hard to reconcile with the bulk of empirical evidence, but for which Karabarbounis and Neiman (2014b) mount a spirited case.

Nevertheless, the one-sector model in section 4 is in many ways unsatisfactory as a model of the distribution between capital and labor. For instance, sections 2 and 3 demonstrated the decisive role of the housing sector in the long-term trajectory of the net capital share—but a one-sector model is by construction unable to account for a shift toward housing. Indeed, Piketty (2015) has recently voiced discomfort with the one-sector interpretation of the rising capital share, arguing that "the right model to think about rising capital-income ratios and capital shares in recent decades is a multi-sector model of capital accumulation." In this section I will construct a tentative version of such a model.

Nested framework. Given the central role of housing in sections 2 and 3, it is first important to distinguish between non-housing and housing output. If household preferences are homothetic in these two types of output, the household objective can be written as a monotonic transformation of a constant returns to scale aggregator $Z(Y_{nh}, Y_h)$ that takes non-housing and housing services as inputs. We can view Z as the "top-level" production function for the economy.

For the non-housing sector, it will be useful to model the production process in a way that reflects the different types of capital studied in section 3 (equipment, structures, and land), so that the results from that disaggregation exercise can be used to inform the model. One natural approach is to assume that structures and land together provide "real estate" services that serve as an input to production, while labor and equipment together provide all other services. This approach enables me to draw upon several empirical literatures, which estimate the relevant elasticities of substitution—for instance, the elasticity of substitution between structures and land in the production of real estate services, or the elasticity of substitution between housing and non-housing in consumer preferences.

Concretely, let $H(N, K_e)$ be a constant returns to scale aggregator combining labor N and equipment K_e, and let $G_1(K_{s1}, L_1)$ be another constant returns to scale aggregator combining nonresidential structures K_{s1} and land L_1. Finally, let F be another constant returns to scale aggregator that combines H and G_1, so that the consolidated production function for the non-housing sector takes the form

$$Y_{nh} = F\left(H(N, K_e), G_1(K_{s1}, L_1)\right) \tag{18}$$

Following section 3, I assume that gross output in the non-housing sector is sold at some markup μ over marginal cost.

Similarly, suppose that residential structures K_{s2} and land L_2 are combined by an aggregate $G_2(K_{s2}, L_2)$ to provide housing services, so that the production function for the housing sector takes the form

$$Y_h = G_2(K_{s2}, L_2) \tag{19}$$

Finally, as already mentioned, Z combines Y_{nh} and Y_h into an aggregate that reflects household preferences:

$$Y = Z(Y_{nh}, Y_h) \tag{20}$$

This multisector economy captures the distinction between the non-housing and housing sectors, as well as all five forms of capital analyzed in section 3: equipment (K_e), nonresidential structures (K_{s1}), nonresidential land (L_1), residential structures (K_{s2}), and residential land (L_2).

The aggregate, nested structure of production in the economy is depicted in the tree below.

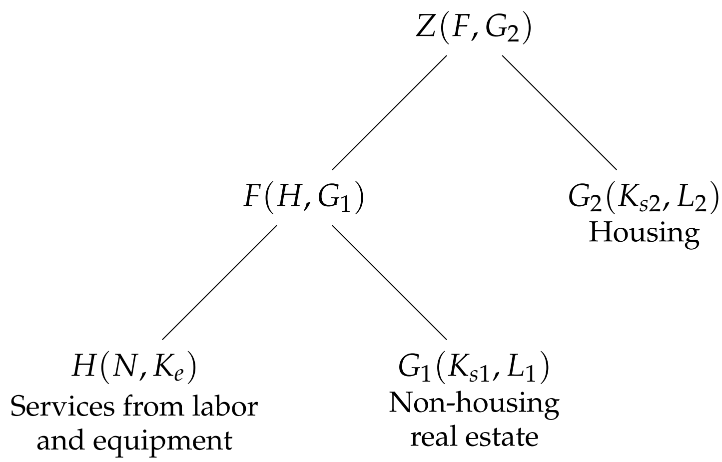

Elasticities of substitution. The response of the multisector model to various shocks is influenced by the local (gross) elasticities of substitution $(\sigma_Z, \sigma_F, \sigma_{G_1}, \sigma_{G_2}, \sigma_H)$ for each of the five constant-returns-to-scale production functions (Z, F, G_1, G_2, H) in the model above.

Although there are extensive empirical literatures that study many of these elasticities, a convincing research design is often elusive, and there is rarely strong consensus around a single point estimate. In the absence of such consensus, I will draw upon each literature to obtain plausible *ranges* for each elasticity, and study the implications of choosing different values within each range. The objective is to see which, if any, conclusions emerge robustly from the multisector model despite allowing for some uncertainty about the σs. Another goal is to investigate which σs matter most to aggregate outcomes, both to clarify thinking and to direct future research toward the most crucial targets.

Surveying the relevant literatures, I find:

- σ_Z equals the elasticity of demand for housing services (as a share of total output) with respect to its price (relative to the aggregate price index for Z). Closely related elasticities of demand for housing have been studied in the literature, which has generally obtained relatively low values. For instance, in a review of the literature Ermisch, Findlay and Gibb (1996) state that "price elasticity estimates are less dispersed than the income elasticity measures, yielding results between 0.5 and 0.8"; and themselves provide an estimate of 0.4.[26] I set a range of $\sigma_Z \in [0.4, 0.8]$.

[26] $\sigma_Z < 1$ is strongly supported by casual observation as well. For instance, as the real price of housing services has risen in the US over the last several decades, its share of consumption has increased slightly; there is also a well-known tendency for consumers to spend a larger share of their budgets on housing in areas where housing is expensive.

- σ_F, the elasticity of substitution between real estate and other services in the non-housing sector, does not map closely onto any empirically studied elasticity. In the absence of direct evidence, I set a wide range of $\sigma_F \in [0.5, 1.5]$.

- σ_{G_1} and σ_{G_2} are the elasticities of substitution between structures and land in the non-housing and housing sectors, respectively. These elasticities play an important role in the urban economics literature, where substitutability between structures and land in the provision of real estate services is of great practical and theoretical interest.

  - The more voluminous literature is for housing, σ_{G_2}, with a widely cited early entry by Muth (1971), who estimates $\sigma_{G_2} = 0.5$ using several approaches. More recently, Thorsnes (1997) surveys the literature and finds that recent estimates have generally been below 1, in the range $[0.5, 1]$; but he also argues that some of these estimates may be biased downward due to measurement error, and that the true elasticity may not be much below 1. This claim is seconded by Ahlfeldt and McMillen (2014). In light of these findings, I set a range of $\sigma_{G_2} \in [0.5, 1]$.
  - The literature for non-housing real estate, σ_{G_1}, is more scattered, with a range of elasticity estimates similar to that for housing—generally below one, but with concerns about bias from measurement error. For instance, Clapp (1979) obtains elasticities from high-rise office data mostly in the range $[0.5, 0.75]$, but in a tentative attempt to correct for measurement error finds that elasticities closer to 1 may be appropriate. Interpretation is complicated by the fact that non-housing real estate is much more heterogenous than housing real estate, spanning everything from high-rise office towers to farmland. Amid this uncertainty, I also set the range $\sigma_{G_1} \in [0.5, 1]$.

- σ_H is the elasticity of substitution between equipment and labor. This is of great speculative interest—there are frequent discussions about the extent to which automation, for instance, can replace existing workers, and σ_H governs the extent to which the decline in equipment prices documented by Karabarbounis and Neiman (2014b) will lead to substitution away from labor. In his survey, Chirinko (2008) reports a wide range of relevant estimates; the majority are still below one, but several are above one as well, and he suggests that the elasticity for equipment may be higher than the aggregate elasticity. Cummins and Hassett (1992), for instance, obtain implied elasticities of 0.93 for equipment but only 0.28 for structures, and estimates listed by Chirinko (2008) that use computer investment obtain values as high as 1.58. I therefore set a range $\sigma_H \in [0.5, 1.5]$.

5.2 Response of the net capital share to exogenous shocks

General methodology. I now study the elasticity of the net capital share with respect to various shocks, in the multisector model whose structure is described in the previous section.

I assume that the quantities (N, L_1, L_2) of labor and both types of land are exogenous. I take final output from Z to be the numeraire, and assume that the prices P_e, P_{s1}, and P_{s2} of reproducible capital in terms of this numeraire are exogenously fixed by technology.[27] As in (2), the user cost of reproducible capital for $i \in \{e, s_1, s_2\}$ are

$$W_{K_i} = P_i(r + \delta - g_{P_i})$$

where the required return r, the depreciation rate δ, and the expected real change in prices g_{P_i} are also all assumed to be exogenous. As in section 3, r may differ between the non-housing and housing sectors. The quantities (K_e, K_{s1}, K_{s2}) of reproducible capital are then given endogenously by demand at this user cost.

I will consider exogenous shocks to either the quantities (N, L_1, L_2) or to either the prices (P_e, P_{s1}, P_{s2}) or r, which jointly determine the user costs $(W_{K_e}, W_{K_{s1}}, W_{K_{s2}})$. The elasticity of factor shares in the model with respect to either these shocks depends only on the initial gross and net shares and the local elasticities $(\sigma_Z, \sigma_F, \sigma_{G_1}, \sigma_{G_2}, \sigma_H)$ of substitution at each level of production; with these in hand, it can be obtained numerically. (Unfortunately, unlike in Oberfield and Raval (2014), elasticities here cannot be expressed in closed form as a weighted average of the individual elasticities $(\sigma_Z, \sigma_F, \sigma_{G_1}, \sigma_{G_2}, \sigma_H)$. Analytically, this is due to the fact that I assume more than one exogenous quantity.)

I calibrate the initial shares to match the decomposition of the US economy in section 3.3 for the final year in the sample, 2013. Table 4 displays the resulting gross and net shares of each factor as a fraction of total income, while table 5 shows the gross shares of each factor as a fraction of the parent aggregate.

Implementation and results. I focus on the elasticity of the net capital share with respect to four specific exogenous shocks:

[27] Since housing is probably not an input to the production of equipment or structures, it would be slightly more natural to assume that these prices are fixed relative to the price of non-housing output F; I assume they are fixed relative to Z for convenience, and in general the relative prices of F and Z do not change enough that this has a sizable impact on the results.

- A shock to the required rate of return r.
- A shock to the price of equipment investment P_e.
- A shock to the price of residential structures investment P_{s2}.
- A shock to the quantity of residential land L_2.

As discussed in greater detail below, the first and second correspond to the Piketty (2014) and Karabarbounis and Neiman (2014b) versions of the accumulation view, respectively. The third and fourth shocks, which relate to residential housing, correspond to my proposed alternative of a "scarcity view".

The core results are summarized in tables 6, 7, and 8. For table 6, I calculate the elasticity of the net capital share with respect to each shock over the full range of σ_i deemed plausible in the previous section

$$(0.4, 0.5, 0.5, 0.5, 0.5) \leq (\sigma_Z, \sigma_F, \sigma_{G_1}, \sigma_{G_2}, \sigma_H) \leq (0.8, 1.5, 1.0, 1.0, 1.5) \tag{21}$$

and report the *minimum* and *maximum* elasticities of the net capital share obtained for any combination of σ_i in this range. I also calculate the elasticity of the net capital share at a set of "benchmark" σ_i, which I choose to be the midpoint of the range: $(\sigma_Z, \sigma_F, \sigma_{G_1}, \sigma_{G_2}, \sigma_H) = (0.6, 1.0, 0.75, 0.75, 1.0)$.

Table 7 provides additional insight into how different assumptions on σ_i combine to produce an aggregate response to shocks. For each shock, the table shows the *sensitivity* (partial derivative) of the net capital share elasticity to changes in each of the underlying σ_i, starting from the benchmark values. Essentially, table 7 shows the gradient of the values in the "benchmark" column of table 6 with respect to perturbations in the σ_i.

For instance, in the case of a shock to P_e, the second row of table 7 shows small sensitivities to all σ_i except σ_H, for which the sensitivity is -0.29. This means that if σ_H is increased slightly from its benchmark value—say, from $\sigma_H = 1.0$ to $\sigma_H = 1.1$—the elasticity of the net capital share with respect to P_e will decline by 0.029. The intuition in this case is straightforward: when σ_H is higher, it is easier to replace equipment with labor in response to higher equipment prices, meaning that a rise in P_e will result in a smaller increase in (or greater decline in) net capital income.

Finally, table 8 decomposes the elasticity of the net capital share, at the benchmark σ_i, into contributing changes in each source of capital income. Each row of table 8 sums to the elasticity for the corresponding shock in the "benchmark" column of table 6, with one exception: an extra row is included for a shock to P_e, showing the decomposition in the "high elasticity" case where each elasticity σ_i is chosen to be at the maximum of the range. (This is because there is virtually no effect from the shock to P_e at the benchmark σ_i.) For instance, for a shock to the price P_{s2} of residential structures investment, the contribution of residential structures K_{s2} is 0.09, out of a total elasticity (from table 6) of 0.07; this means that when the cost of residential investment rises, more than 100% of the resulting increase in the net capital share is due to a rise in income from residential structures themselves.

I now discuss and interpret the results for each shock.

Shock to the required rate of return r. This case tests the Piketty (2014) hypothesis that a rise in savings will push up the net capital share. In general equilibrium, increased savings influences capital income by pushing down the real interest rate; hence, to learn the sign of the effect of savings on the net capital share, it suffices to study the partial equilibrium effect of a change in the real interest rate.

Since the decomposition of the US economy in section 3.3 allows r in the non-housing and housing sectors to be different, I define a "shock to r" to be a parallel shift dr in these two rates of return. I then define the elasticity of the net capital share with respect to this shock to be

$$\frac{\partial(\text{net capital share})\,/\,(\text{net capital share})}{\partial r\,/\,r^{ave}}$$

where r^{ave} is the *average* return on capital across the economy as a whole, including both the non-housing and housing sectors.

Table 6 shows that for all σ_i within range (21), the response of the net capital share to r is positive: barely so at minimum (0.04) and strongly so at maximum (0.54). This is inconsistent with the Piketty (2014) hypothesis that a decline in r can produce an increase in the net capital share, and it corroborates the findings from the single-sector model in section 4.

Table 7 reveals that the response of the net capital share to r depends primarily on three elasticities, all negatively: σ_Z, σ_F, and σ_H, each with a sensitivity of about -0.20. Each of these elasticities governs the extent to which an aggregate that includes labor (which is unaffected by r) can be substituted for an aggregate that does not include labor. But even when these elasticities are chosen at the maximum level in the range ($\sigma_Z = 0.8$, $\sigma_F = 1.5$, $\sigma_H = 1.5$), the response of the net capital share to r remains slightly positive.

Table 8 shows that the vast majority of the response to r comes from residential structures: at benchmark σ_i, a contribution of 0.23 out of an overall elasticity of 0.26. This is for two reasons. First, since both housing G_2 and aggregate consumer demand Z have σs below 1, the direct positive impact of rising r on income from residential structures outweighs the negative effect of substitution—much more so than for nonresidential structures or equipment. Second, since section 3.3 finds a lower r for the housing sector than the non-housing sector, a parallel shift in these rates has a disproportionate effect on housing. This reinforces the centrality of housing to any assessment of the Piketty (2014) narrative.

Finally, table 8 indicates the importance of a crucial distinction—namely, the distinction between (A) the ratio of housing capital to aggregate income and (B) the share of housing capital *income* in aggregate income. In response to rising r, (A) falls: higher r pushes down the demand for residential structures relative to aggregate income[28], and since residential land's share of income remains roughly constant in table 8, higher r will push down the valuation of this land relative to aggregate income. At the same time, as already discussed, (B) rises dramatically. Hence a shock to r pushes (A) and (B) in different directions, making it important to document (A) and (B) separately.

Shock to the price of equipment investment P_e. This case tests the Karabarbounis and Neiman (2014b) hypothesis that declining investment prices—which have been concentrated in equipment—will push up the net capital share. As table 6 shows, this remains ambiguous for the range of σ_i specified in (21), which are consistent with either a positive or negative relationship between P_e and the net capital share.

Table 7 makes clear the source of this ambiguity: the response of the net capital share to P_e depends almost entirely on the elasticity of substitution σ_H between labor and equipment. When σ_H is near the top of the $[0.5, 1.5]$ range, falling P_e leads to a rise in the net capital share, consistent with Karabarbounis and Neiman (2014b); when σ_H is near the bottom, the opposite is true.

The "high elasticity" row in table 8, however, provides cause for skepticism of the Karabarbounis and Neiman (2014b) channel. Here, P_e has a substantial negative effect on the net capital share. But this effect comes almost exclusively from the net capital income of equipment itself—which, in this partial equilibrium exercise, moves in parallel with the value of the equipment stock—rather than through some less direct channel. Section 3 found that the value of equipment (which has recently fallen) has followed a path quite distinct from the path of the net capital share (which has recently risen). This is not consistent with a major role for P_e.

For the P_e hypothesis to be consistent with the data, it would be necessary for declining P_e to push up the net capital share through some channel other than a rise in the value of the equipment stock. Karabarbounis and Neiman (2014a) sketch one such possibility,

[28] This occurs in the model, but not directly visible in tables 6 through 8 .

where falling P_e can lead to an increase in the net capital share despite an actual decline in the aggregate value of equipment, but the multisector model here does not corroborate their mechanism.[29]

Shock to the price of residential structures investment P_{s2}. In table 6, a rise in P_{s2} leads to a rise in the net capital share for benchmark σ_i; for other choices of σ_i within range (21), there is at worst roughly no effect. According to table 7, the effect is most sensitive to the elasticity σ_Z of substitution between housing and non-housing output; and according to table 8 it works almost entirely through the net capital income of residential structures themselves. The mechanism here is relatively simple: when the ability to substitute away from housing is limited, costlier residential investment leads to a higher-value housing stock and a larger share of income accruing to housing.

Shock to the quantity of residential land L_2. This is similar to the previous case. In table 6, a decline in the quantity of residential land L_2 leads to a rise in the net capital share for benchmark σ_i; for other choices of σ_i within range (21), there is at worst roughly no effect. Again, according to table 7, the effect is most sensitive to σ_Z; now, however, the effect is smaller and works mainly through the net capital income earned by residential land.

Summary of results and conclusion. I have examined the response of the multisector model to four exogenous shocks. The first two shocks correspond to versions of the accumulation view: a shock to r captures the general equilibrium channel through which the rise in savings postulated by Piketty (2014) affects factor shares, while a shock to P_e is central to the Karabarbounis and Neiman (2014b) narrative.

In both cases, the results do not support the proposed mechanism. For all choices of $\{\sigma_i\}$ within the range considered, a fall in r leads to a fall in the net capital share, in contrast with Piketty (2014). Meanwhile, although a fall in P_e can produce a rise in the net capital share, it only does so by pushing up the net capital income from equipment itself, which is at odds with the evidence from section 3.

[29] Karabarbounis and Neiman (2014a) devise a model where two types of capital, high-depreciation (which can be interpreted as equipment) and low-depreciation (which can be interpreted as structures) combine to form a capital aggregate; the elasticity of substitution between these types of capital is less than 1, while the elasticity of substitution between the capital aggregate and labor is greater than 1. A decline in the price of equipment lowers the price of the capital aggregate, which induces substitution from labor to the capital aggregate; but since the elasticity of substitution between equipment and structures is less than 1, this also causes a decline in equipment relative to structures. With the right parameters, it is possible for a decline in the price of equipment to increase the net capital share while net capital income from equipment itself actually declines.

The latter two shocks both embody some form of the *scarcity view*, which is more successful in the multisector model. Either a rise in the price P_{s2} of residential investment or a fall in the quantity L_{s2} of residential land leads to a rise in the net capital share, for the vast majority of $\{\sigma_i\}$ in the range (21). In both cases, the mechanism works through increasing the net capital income earned by housing, consistent with the dramatic rise in the contribution of housing documented in section 2.

5.3 Counterfactual exercise

Building upon the promise of the scarcity view in the previous section, I now use the multisector model to perform a counterfactual exercise, exploring the implications of alternative paths for P_{s2} and L_{s2}.

The real price P_{s2} of residential investment has risen in the last several decades in the US; furthermore, real output has grown substantially, putting pressure on the supply of residential land. I consider a counterfactual where these two forces are not present: where the real price of P_{s2} is instead constant from the beginning of the sample period (1948) onward, and where the quantity of residential land L_{s2} grows in tandem with real output from the beginning of the sample period onward.[30]

In contrast to the exercises in section 5.2, which consider only local shocks to exogenous variables, this counterfactual involves large *global* changes. It requires additional, global assumptions to compute; for this purpose, I will assume that the production functions (Z, F, G_1, G_2, H) each have a globally constant elasticity of substitution. I consider two choices of $\{\sigma_i\}$: first, the benchmark $(\sigma_Z, \sigma_F, \sigma_{G_1}, \sigma_{G_2}, \sigma_H) = (0.6, 1.0, 0.75, 0.75, 1.0)$; and second, an alternative $(\sigma_Z, \sigma_F, \sigma_{G_1}, \sigma_{G_2}, \sigma_H) = (0.4, 0.5, 0.75, 0.75, 0.5)$ that sets σ_Z, σ_F, and σ_H (the σs that govern the response to P_{s2} and L_2, according to table 7) to the minimum values in the range (21).

Figure 12 displays the results of this exercise, distinguishing between the housing and non-housing components of the net capital share. Consistent with table 8, there is little effect working through the non-housing component. Furthermore, the large initial increase and then decline in the housing component, through 1980, is left untouched by the counterfactuals. Much of the subsequent increase in the housing component, however, is eroded.

This is consistent with a role for rising residential investment costs, along with growing scarcity of residential land, in driving up housing's contribution to the net capital

[30] To make this modification, I assume that the quantity of land L_{s2} was in reality constant, and then expand it in each year by a fraction equal to cumulative real GDP growth since 1948. Depending on the interpretation of L_{s2}, the assumption that it has been constant may or may not be appropriate.

share: when these forces are reversed in a counterfactual, we see less of a rise. At the same time, figure 12 makes clear the limitations of this account. It does not explain the fall and rise in the non-housing component, nor can it explain all aspects of the housing time series. The scarcity view, therefore, is only a partial replacement for the accumulation view: it achieves better consistency with data and theory, but does not purport to explain more than a fragment of the evolving factor income distribution.

6 Conclusion

The aging Kaldor facts have retreated in the face of experience. Today, macroeconomists no longer claim that factor shares are constant—but what should replace the old consensus?

It is increasingly commonplace to believe that labor is ceding ground to capital. But a closer look at postwar experience reveals a murkier story, in which steady increase is limited to the gross capital share. The net share, by contrast, has fallen and then recovered; it consists of a large long-term increase in net capital income from housing, and a more volatile contribution from the rest of the economy with little cumulative movement in either direction.

Even more elusive than these facts is a cohesive explanation of them. The accumulation view, in both its Piketty (2014) and Karabarbounis and Neiman (2014b) variants, falters in multiple respects. It cannot explain the dominant role of housing, nor can it be readily reconciled with the evidence on elasticities of substitution. Outside of housing, there appears to be little correlation between the capital-income ratio and the net capital share.

The rise in housing's contribution to the capital share, by contrast, can be explained in part as the result of scarcity. The rising real cost of residential investment and the limited quantity of residential land have conspired to make housing more expensive, and given low elasticities of substitution this has meant a rise in housing's share of income.

With these trends in mind, policymakers concerned about the distribution of income should keep an eye on housing costs—many urban economists, including Glaeser, Gyourko and Saks (2005) and Quigley and Raphael (2005), have documented explicitly how restrictions on land use and residential construction inflate the cost of housing. Outside of housing, however, this paper raises more questions than it answers about the evolution of the net capital share: once the accumulation view has been discarded, there is no master narrative at hand that can explain the postwar fall and rise.

If anything, these results suggest that concern about inequality should be shifted away from the overall split between capital and labor, and toward other aspects of distribution, such as the within-labor distribution of income. Although the net capital share has at times seen dramatic shifts both up and down, away from housing its long-term movement has been quite small, and there is no compelling reason to suggest that this pattern will change going forward.

No doubt, however, the distribution between capital and labor will continue to be a salient issue: we surely have not seen the last of Ricardo (1821)'s principal problem of Political Economy.

References

Ahlfeldt, Gabriel M. and Daniel P. McMillen, "New Estimates of the Elasticity of Substitution of Land for Capital," in "ERSA conference papers" European Regional Science Association 2014.

Andrews, Dan and Aida Caldera Sanchez, "The evolution of homeownership rates in selected OECD countries: Demographic and public policy influences," *OECD Journal: Economic Studies*, 2011, *2011* (1), 1–37.

Arpaia, Alfonso, Esther Perez, and Karl Pichelmann, "Understanding labour income share dynamics in Europe," Working Paper May 2009.

Arrow, K. J., H. B. Chenery, B. S. Minhas, and R. M. Solow, "Capital-Labor Substitution and Economic Efficiency," *The Review of Economics and Statistics*, August 1961, *43* (3), 225–250.

Azmat, Ghazala, Alan Manning, and John Van Reenen, "Privatization and the Decline of Labour's Share: International Evidence from Network Industries," *Economica*, July 2012, *79* (315), 470–492.

Bonnet, Odran, Pierre-Henri Bono, Guillaume Chapelle, and Etienne Wasmer, "Does housing capital contribute to inequality? A comment on Thomas Piketty's Capital in the 21st Century," *Sciences Po Economics Discussion Papers*, 2014, *7.*

Bridgman, Benjamin, "Is Labor's Loss Capital's Gain? Gross versus Net Labor Shares," *Bureau of Economic Analysis (June 2014)*, 2014.

Chirinko, Robert S., "sigma: The long and short of it," *Journal of Macroeconomics*, June 2008, *30* (2), 671–686.

Clapp, John M., "The substitution of urban land for other inputs," *Journal of Urban Economics*, January 1979, *6* (1), 122–134.

Cobb, Charles W. and Paul H. Douglas, "A Theory of Production," *The American Economic Review*, March 1928, *18* (1), 139–165.

Cummins, Jason G. and Kevin A. Hassett, "The effects of taxation on investment: new evidence from firm level panel data," *National Tax Journal*, 1992, pp. 243–251.

Elsby, Michael W. L., Bart Hobijn, and Aysegul Sahin, "The Decline of the U.S. Labor Share," *Brookings Papers on Economic Activity*, 2013, *2013* (2), 1–63.

Ermisch, J. F., J. Findlay, and K. Gibb, "The Price Elasticity of Housing Demand in Britain: Issues of Sample Selection," *Journal of Housing Economics*, March 1996, *5* (1), 64–86.

Feenstra, Robert C., Robert Inklaar, and Marcel Timmer, "The Next Generation of the Penn World Table," Working Paper 19255, National Bureau of Economic Research July 2013.

Glaeser, EdwardÂ L., Joseph Gyourko, and Raven Saks, "Why Is Manhattan So Expensive? Regulation and the Rise in Housing Prices," *Journal of Law and Economics*, October 2005, *48* (2), 331–369.

Gollin, Douglas, "Getting Income Shares Right," *Journal of Political Economy*, April 2002, *110* (2), 458–474.

Johnson, D. Gale, "The Functional Distribution of Income in the United States, 1850-1952," *The Review of Economics and Statistics*, May 1954, *36* (2), 175–182.

Kaldor, Nicholas, "A Model of Economic Growth," *The Economic Journal*, December 1957, *67* (268), 591–624.

Kalecki, Michal, "The Determinants of Distribution of the National Income," *Econometrica*, April 1938, *6* (2), 97–112.

Karabarbounis, Loukas and Brent Neiman, "Capital Depreciation and Labor Shares Around the World: Measurement and Implications," Technical Report, National Bureau of Economic Research 2014.

_ **and** _ , "The Global Decline of the Labor Share," *The Quarterly Journal of Economics*, February 2014, *129* (1), 61–103.

Keynes, John Maynard, "Relative Movements of Real Wages and Output," *The Economic Journal*, March 1939, *49* (193), 34–51.

Koh, Dongya, Raul Santaeulalia-Llopis, and Yu Zheng, "Labor Share Decline and the Capitalization of Intellectual Property Products," SSRN Scholarly Paper ID 2546974, Social Science Research Network, Rochester, NY April 2015.

Mitchell, Wesley C., *Business Cycles*, Berkeley: University of California Press, 1913.

_ , *Business cycles and their causes*, Berkeley and Los Angeles: University of California press, 1941.

Modigliani, Franco and Richard A. Cohn, "Inflation, Rational Valuation and the Market," *Financial Analysts Journal*, March 1979, *35* (2), 24–44.

Muth, Richard F., "The derived demand for urban residential land," *Urban studies*, 1971, *8* (3), 243–254.

Oberfield, Ezra and Devesh Raval, "Micro Data and Macro Technology," Working Paper 20452, National Bureau of Economic Research September 2014.

Piketty, Thomas, *Capital in the Twenty-First Century*, Cambridge Massachusetts: Belknap Press, March 2014.

_ , "Putting Distribution Back at the Center of Economics: Reflections on Capital in the Twenty-First Century," *Journal of Economic Perspectives*, February 2015, *29* (1), 67–88.

_ **and Gabriel Zucman**, "Capital is Back: Wealth-Income Ratios in Rich Countries 1700-2010," Working Paper 2013.

_ **and** _ , "Capital is back: Wealth-income ratios in rich countries, 1700-2010," *The Quarterly Journal of Economics*, 2014, p. qju018.

Quigley, John M. and Steven Raphael, "Regulation and the High Cost of Housing in California," *The American Economic Review*, May 2005, *95* (2), 323–328.

Ricardo, David, *On the Principles of Political Economy, and Taxation*, John Murray, 1821.

Rotemberg, Julio J. and Michael Woodford, "Chapter 16 The cyclical behavior of prices and costs," in John B. Taylor and Michael Woodford, ed., *Handbook of Macroeconomics*, Vol. Volume 1, Part B, Elsevier, 1999, pp. 1051–1135.

Solow, Robert M., "A Skeptical Note on the Constancy of Relative Shares," *The American Economic Review*, September 1958, *48* (4), 618–631.

Thorsnes, Paul, "Consistent Estimates of the Elasticity of Substitution between Land and Non-Land Inputs in the Production of Housing," *Journal of Urban Economics*, July 1997, *42* (1), 98–108.

Weitzman, Martin L., "On the Welfare Significance of National Product in a Dynamic Economy," *The Quarterly Journal of Economics*, February 1976, *90* (1), 156–162.

A Figures

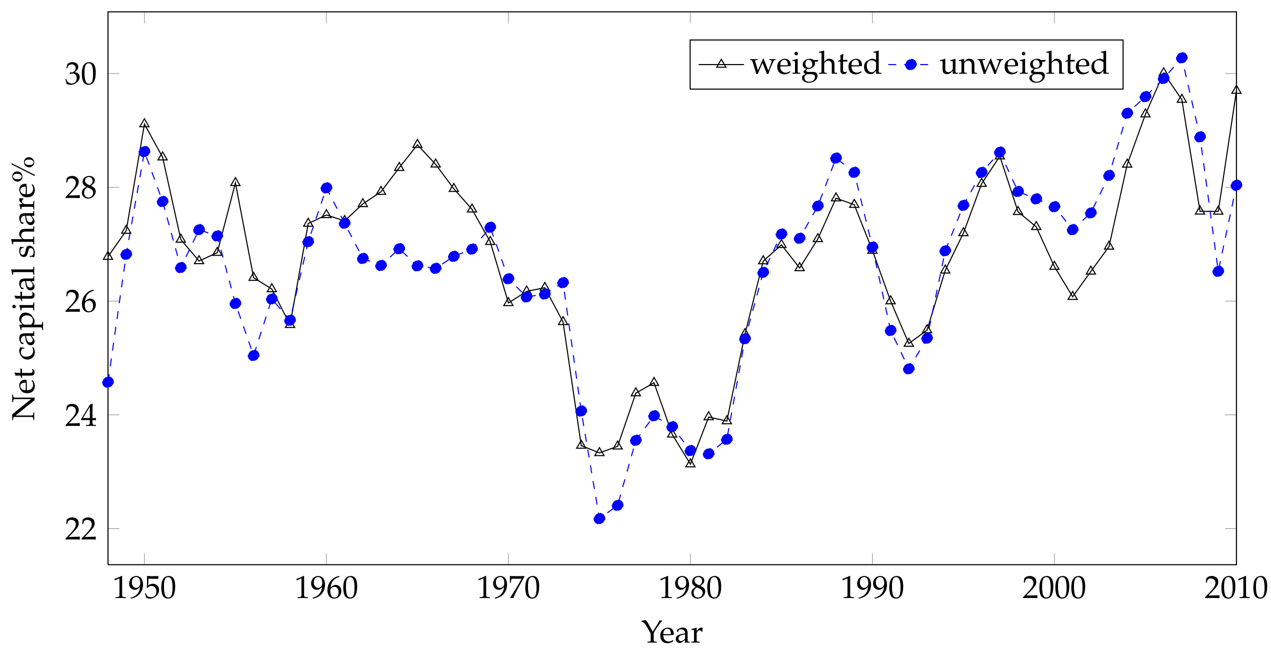

Figure 1: Average net capital share of private domestic value added for G7 countries.

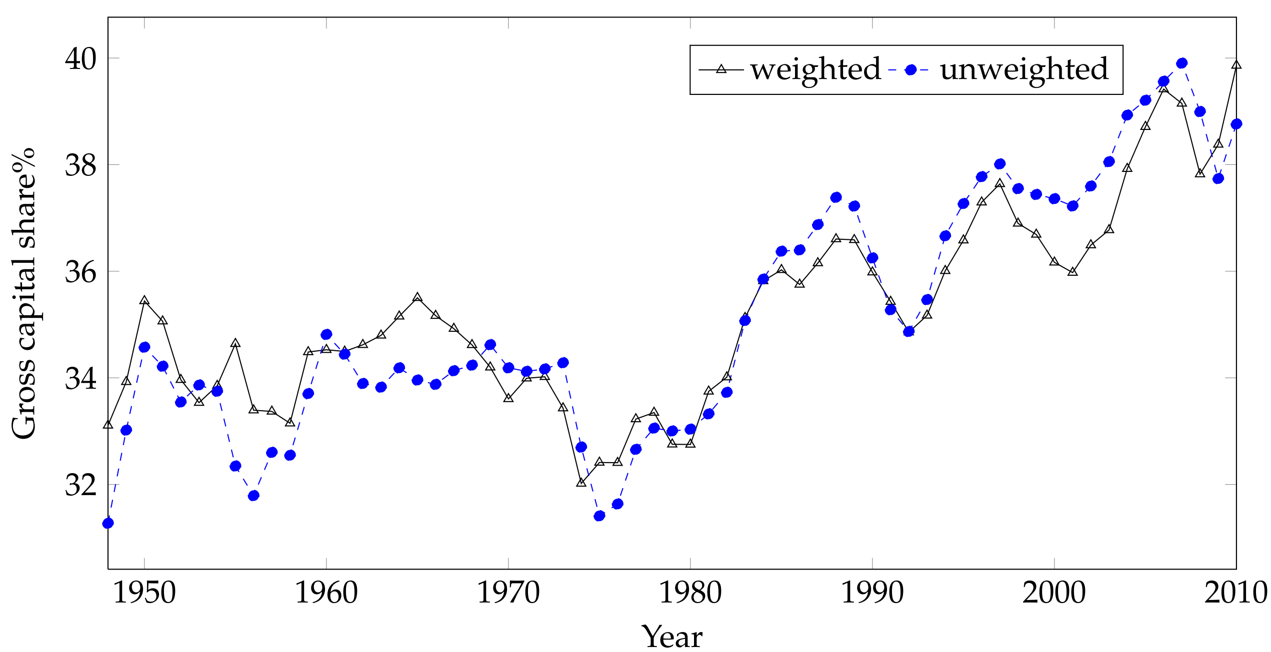

Figure 2: Average gross capital share of private domestic value added for G7 countries.

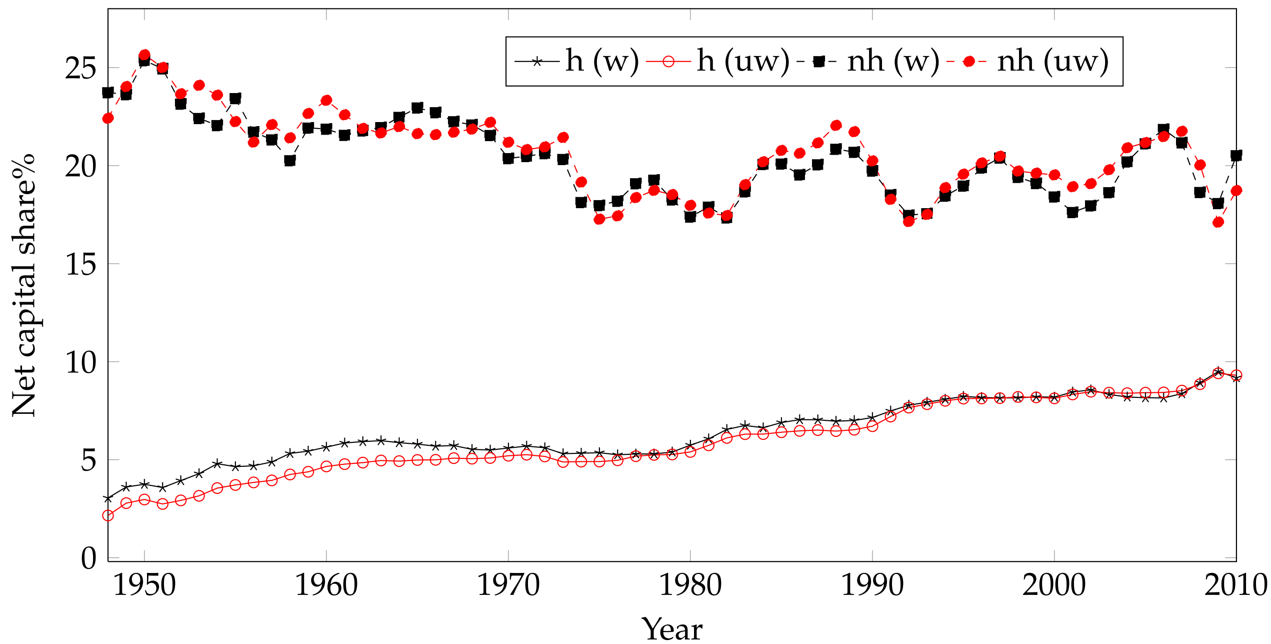

Figure 3: Components of average net capital share of private domestic value added for G7 countries: housing (h) versus other (nh) sectors, weighted (w) and unweighted (uw).

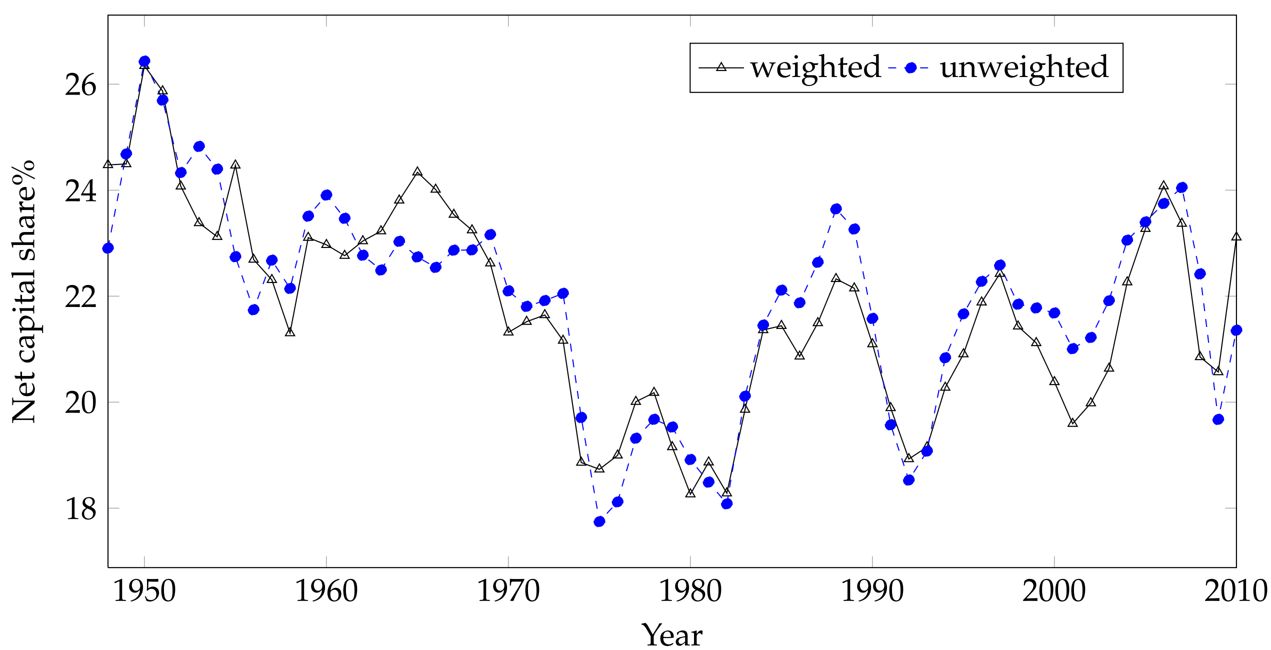

Figure 4: Average net capital shares of corporate sector value added for G7 countries.

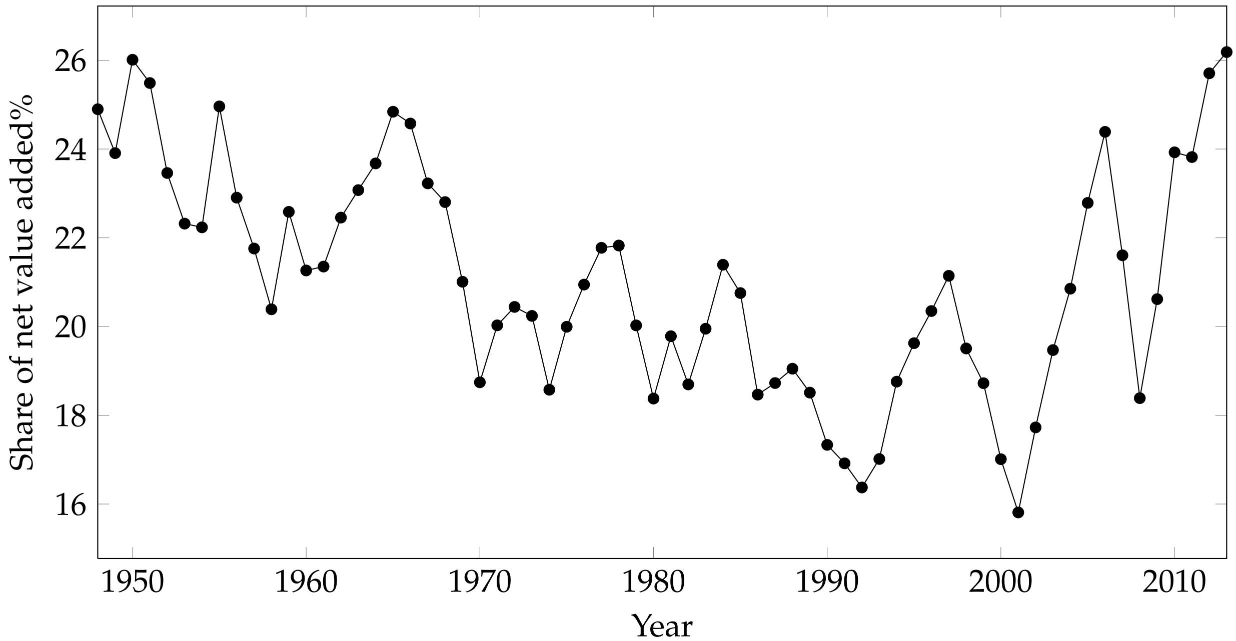

Figure 5: Net capital share of corporate sector value added in the US.

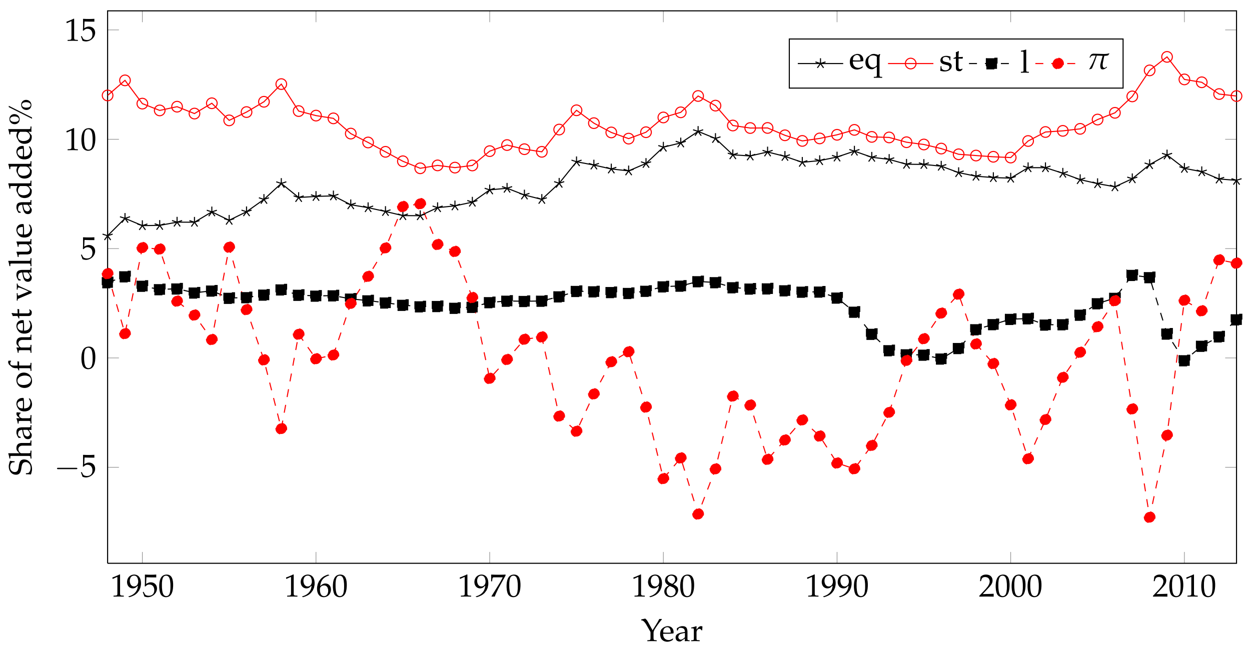

Figure 6: Decomposition of net capital share of corporate sector value added in the US: return on equipment, structures, land, and pure profits π.

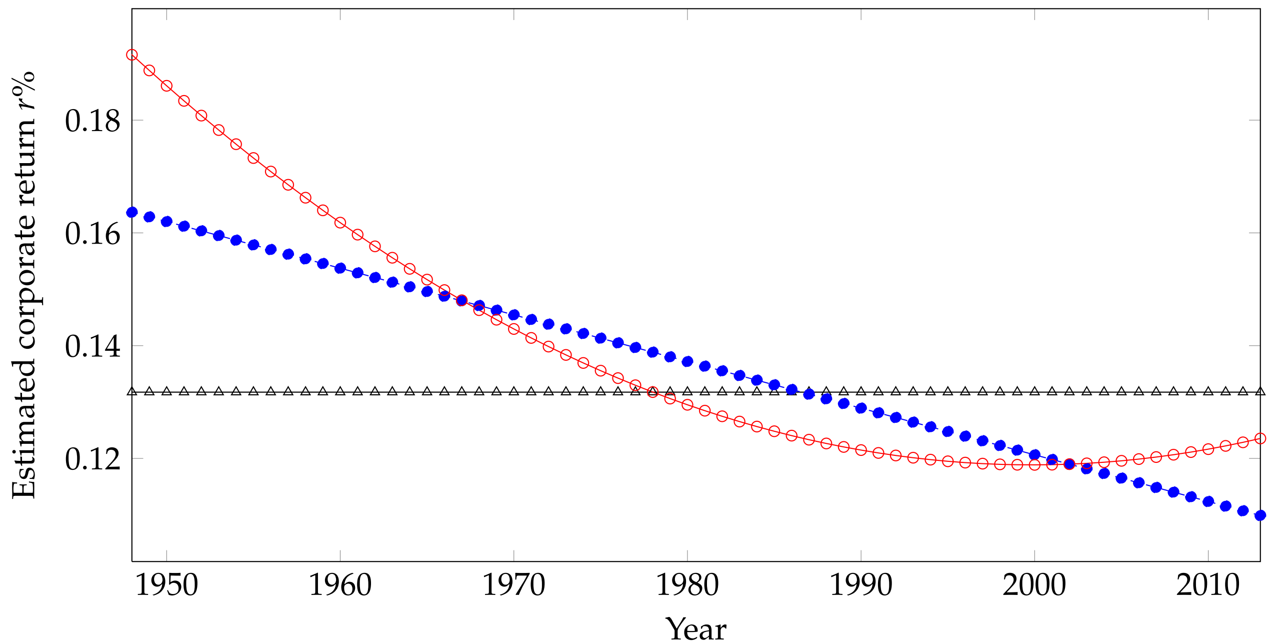

Figure 7: Estimated constant, linear, and quadratic time trends for the corporate rate of return $r(t)$.

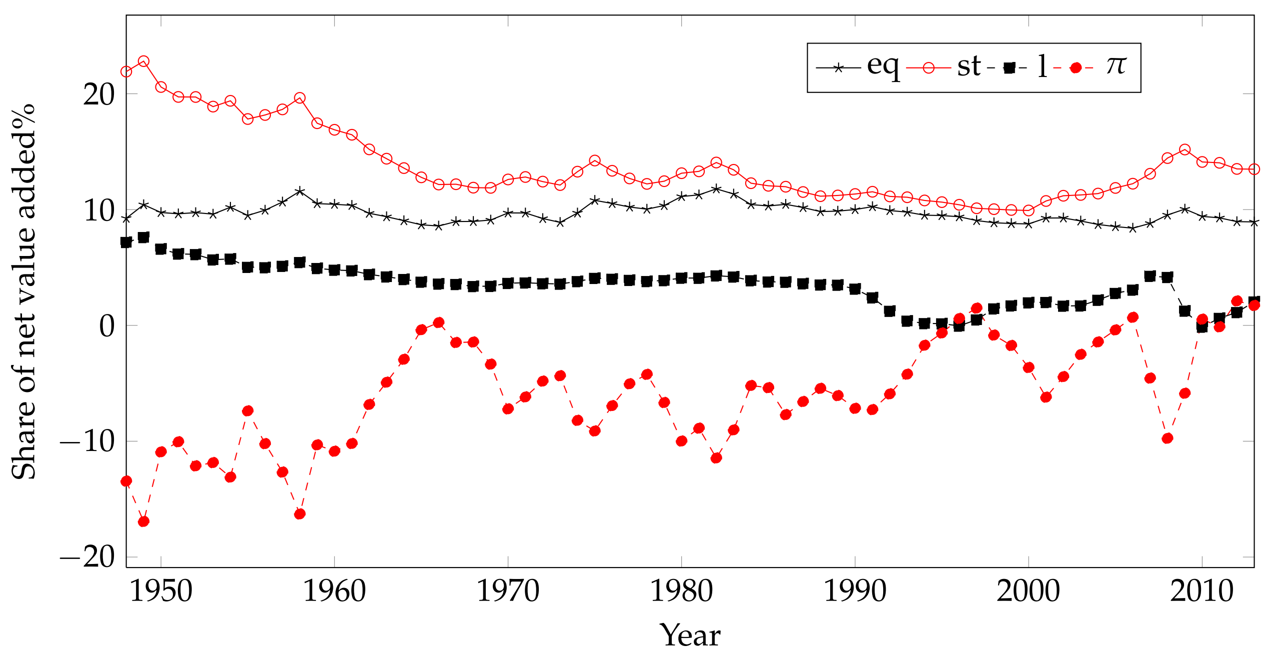

Figure 8: Decomposition of net capital share of corporate sector value added in the US: return on equipment, structures, land, and pure profits π, using quadratic trend for $r(t)$.

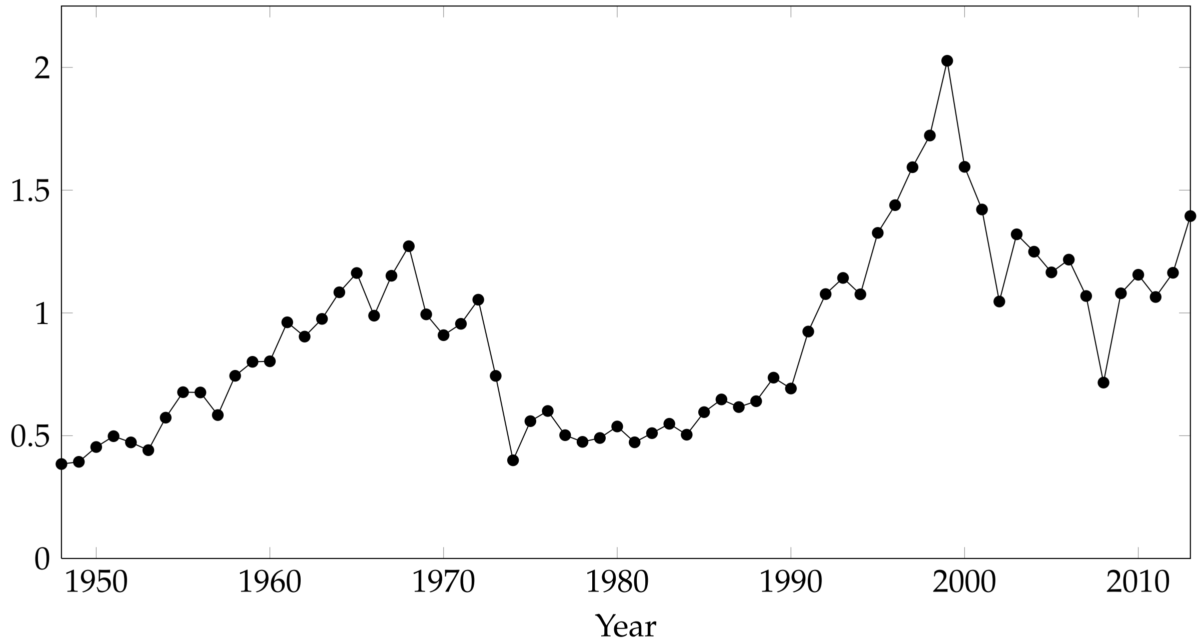

Figure 9: Ratio of total market value to the recorded value of equipment, structures, and land ("book value"), US corporate sector.

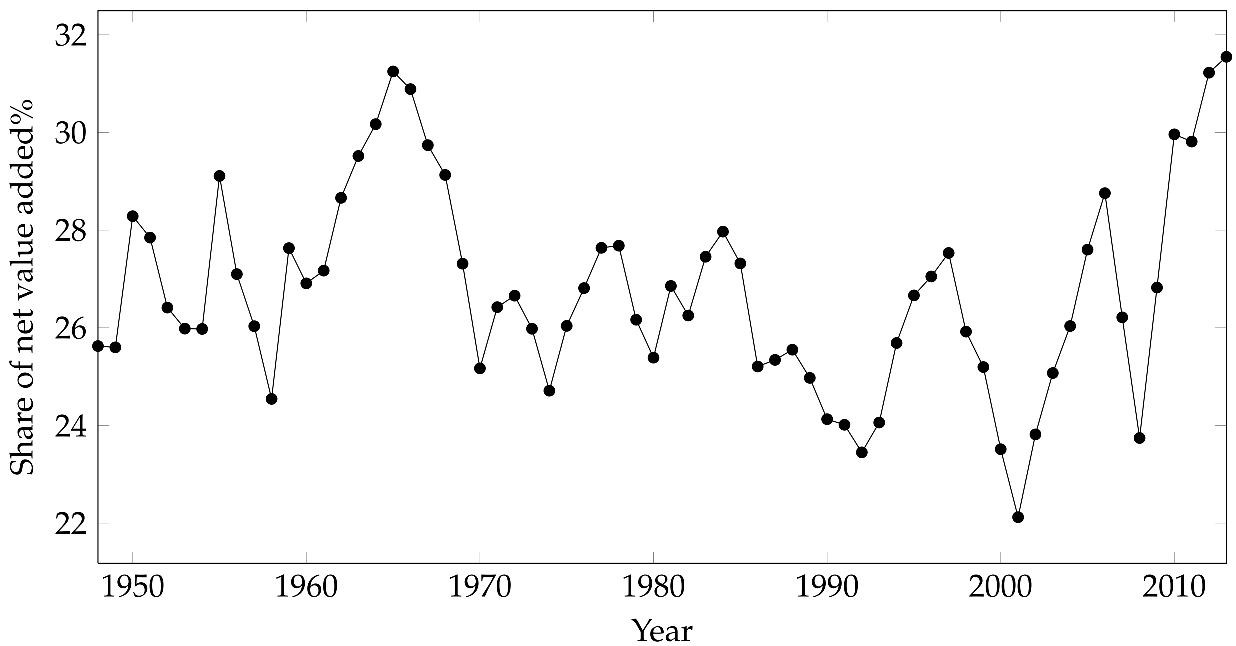

Figure 10: Net capital share of private value added in the US.

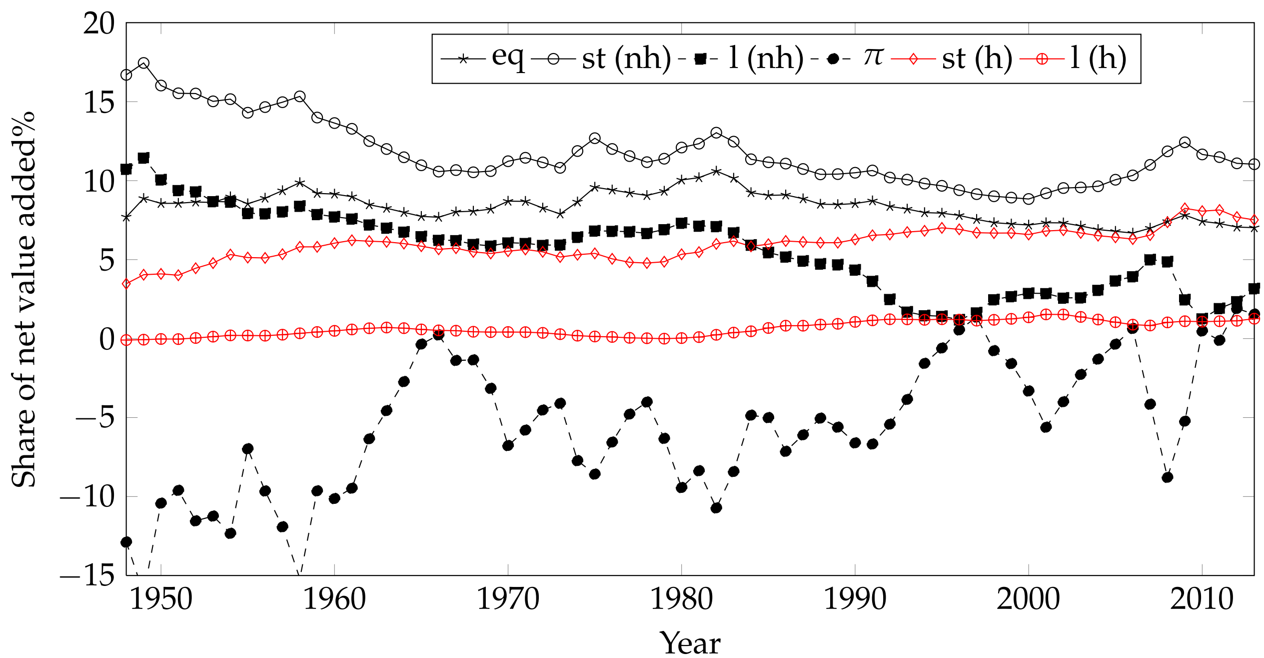

Figure 11: Decomposition of net capital share of private domestic value added in the US: return on equipment (eq), non-residential structures (st-nh), non-residential land (l-nh), pure profits π, residential structures (st-h), and residential land (l-h).

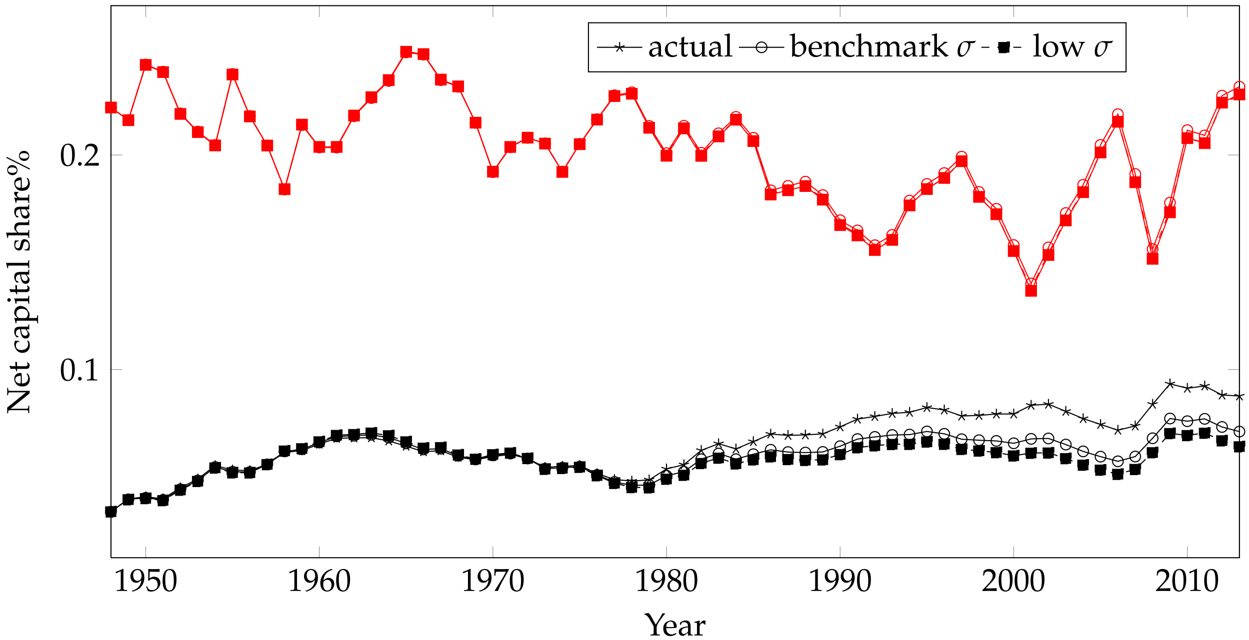

Figure 12: Counterfactual paths—assuming no change in the real price of residential structures investment, and a constant ratio of the quantity of residential land to the quantity of real output—for the housing (black) and non-housing (red) components of the net capital share.

B Tables

Table 2: Decadal averages for the net capital share of private domestic value added, broken into housing and non-housing ("other") components.

		1950s	1960s	1970s	1980s	1990s	2000s
USA	Housing	5.3%	6.5%	5.7%	7.2%	8.4%	8.2%
	Other	22.0%	21.7%	18.6%	18.4%	19.2%	19.4%
	Total	27.3%	28.2%	24.2%	25.6%	27.5%	27.6%
Japan	Housing		4.2%	3.6%	4.1%	5.2%	7.0%
	Other		31.2%	26.9%	25.7%	21.6%	20.1%
	Total		35.4%	30.5%	29.8%	26.9%	27.1%
Germany	Housing					2.9%	3.4%
	Other					23.5%	28.0%
	Total					26.4%	31.4%
France	Housing	3.6%	5.1%	5.9%	7.1%	9.8%	10.8%
	Other	21.3%	19.8%	17.9%	16.6%	19.9%	18.0%
	Total	24.9%	24.9%	23.8%	23.7%	29.7%	28.8%
UK	Housing	1.2%	2.1%	3.8%	4.6%	5.8%	7.3%
	Other	27.2%	23.9%	18.3%	21.6%	23.2%	23.4%
	Total	28.4%	26.0%	22.1%	26.2%	29.0%	30.7%
Italy	Housing					4.3%	6.4%
	Other					33.9%	32.5%
	Total					38.2%	38.9%
Canada	Housing		6.6%	6.6%	8.1%	10.4%	8.6%
	Other		22.5%	24.0%	25.8%	21.2%	27.2%
	Total		29.1%	30.6%	33.8%	31.6%	35.8%

Table 3: Decadal averages for the net capital share of value added in the domestic corporate sector.

	1950s	1960s	1970s	1980s	1990s	2000s
USA	23.2%	23.2%	19.7%	19.8%	20.9%	21.1%
Germany					24.2%	29.0%
France	22.1%	20.9%	19.0%	17.9%	22.1%	20.1%
UK	27.6%	24.4%	19.0%	22.7%	24.7%	25.3%
Italy					35.4%	34.6%
Canada		24.5%	26.1%	28.5%	24.3%	30.1%

Table 4: Gross and net shares of factors and higher-level aggregates taken from 2013 decomposition in section 3.3, used to calibrate the multisector model.

	Gross aggregate share	Net aggregate share
N	60%	68%
K_e	12%	7%
K_{s1}	12%	11%
L_1	3%	3%
K_{s2}	10%	8%
L_2	1%	1%
π	1%	2%
H	72%	76%
G_1	15%	14%
G_2	11%	9%
F	88%	90%

Table 5: Gross shares of production within each higher-level aggregate in calibrated the multisector model, based on shares in table 4.

		Gross share
H	N	83%
	K_e	17%
G_1	K_{s1}	82%
	L_1	18%
G_2	K_{s2}	90%
	L_2	10%
F	H	83%
	G_1	17%
Z	F	89%
	G_2	11%

Table 6: Minimum and maximum elasticities (for choices of σ_i within range) of net capital share with respect to shocks, in addition to elasticity for benchmark σ_i.

Shock	Min	Max	Benchmark
r	0.04	0.54	0.26
P_e	-0.18	0.15	0.00
P_{s2}	-0.00	0.16	0.07
L_2	-0.04	0.00	-0.01

Table 7: Sensitivity of the elasticity of net capital share to changes in each σ_i, starting at benchmark values.

	Sensitivity				
Shock	σ_Z	σ_F	σ_{G_1}	σ_{G_2}	σ_H
r	**-0.21**	**-0.21**	0.03	-0.01	**-0.19**
P_e	0.02	0.01	0.00	0.01	**-0.29**
P_{s2}	**-0.17**	-0.03	0.00	0.01	-0.04
L_2	0.03	0.00	0.00	0.01	0.00

Table 8: Contribution to the aggregate elasticity of the net capital share, for benchmark σ_i.

	Contribution					
Shock	K_e	K_{s1}	L_1	K_{s2}	L_2	Profit
r	0.04	0.03	-0.04	**0.23**	0.00	-0.01
P_e	0.00	0.00	0.00	0.00	0.00	0.00
P_e (high elasticity case)	**-0.12**	0.00	0.00	-0.01	0.00	0.00
P_{s2}	-0.01	-0.01	-0.01	**0.09**	0.01	0.00
L_2	0.00	0.00	0.00	0.00	-0.01	0.00

C Description of alternative procedure in section 3.2 for estimating path of r.

I sketch here the procedure in section 3.2 for estimating the effective required return $r(t)$ on capital for the US corporate sector. I specify the model in continuous time, and use superscripts to denote the time t for economy of notation. I am also more explicit here about how the underlying continuous time flows are aggregated into the measured flow for a given time period.

Assumed stochastic processes

I assume two stochastic processes beyond what is already visible in the data:

- π^t, which is a stationary, ergodic process for the share of gross output that goes to profits.
- ζ^t, which reflects stochastic pricing error for the total market value of corporations, with mean 1 (where 1 corresponds to no error).

Core relations

First relation: profit share of flows. We know that all non-profit income will be allocated between depreciation, labor, and the various types of capital. This is a flow relation

$$(1-\pi^t)Y^t = w^t L^t + \sum_i (\delta_i + r^t - g_{P_i}) P_i^t K_i^t$$

which can be rewritten as

$$\pi^t = 1 - \frac{w^t L^t}{Y^t} - \sum_i \frac{\delta_i P_i^t K_i^t}{Y^t} - \sum_i (r^t - g_{P_i}) \frac{P_i^t K_i^t}{Y^t} \tag{22}$$

or, if we don't want to divide by Y^t, as

$$\pi^t Y^t = Y^t - w^t L^t - \sum_i \delta_i P_i^t K_i^t - \sum_i (r^t - g_{P_i}) P_i^t K_i^t$$

Consolidating into an accumulated flow. Suppose that we write

$$\int_t^{t+\Delta t} \pi^s Y^s \, ds = \int_t^{t+\Delta t} (Y^s - w^s L^s - \sum_i \delta_i P_i^s K_i^s) \, ds - \sum_i \int_t^{t+\Delta t} (r^s - g_{P_i}) P_i^s K_i^s \, ds \tag{23}$$

We can identify the first part as simply real net capital income during the period, while for the second term we must write

$$\int_t^{t+\Delta t} (r^s - g_{P_i}) P_i^s K_i^s \, ds \approx \left((r^{t+\Delta t} - g_{P_i}) P_i^{t+\Delta t} K_i^{t+\Delta t} + (r^t - g_{P_i}) P_i^t K_i^t \right) \frac{\Delta t}{2} \tag{24}$$

Second relation: asset pricing. The expected discounted value of the profit stream from time t onward is (in real terms)

$$Y^t \cdot \int_0^\infty e^{(g_Y - \delta_\pi)s - \int_t^{t+s} r^u} \mathbb{E}_t[\pi^{t+s}]$$

where δ_π is the rate at which pure profits decay. (We can think of it as the rate at which a given company, for instance, on average loses the ability to make pure profits. There is no clear basis for picking δ_π, and I will choose $\delta_\pi = .015$, which implies a half-life of just below 50 years—within reason given the typical lifetimes of American corporations. Fortunately, the precise choice of δ_π does not matter much for the results.)

This expected discounted value plus the value of capital itself is (again in real terms)

$$Y^t \cdot \int_0^\infty e^{(g_Y - \delta_\pi)s - \int_t^{t+s} r^u du} E_t[\pi^{t+s}] \, ds + \sum_i P_i^t K_i^t$$

I assume that the market value of the corporate sector equals this overall value times ζ^t, the multiplicative stochastic pricing error that has mean 1, follows a stationary, ergodic process, and is drawn independently of π, $\{P_i^t\}$, $\{K_i^t\}$, and Y^t:

$$MV^t = \zeta^t \left(Y^t \cdot \int_0^\infty e^{(g_Y - \delta_\pi)s - \int_t^{t+s} r^u du} E_t[\pi^{t+s}] \, ds + \sum_i P_i^t K_i^t \right) \tag{25}$$

Define $OMV^t \equiv MV^t - \sum_i P_i^t K_i^t$, and rewrite (25) as

$$OMV^t = \zeta^t \left(Y^t \cdot \int_0^\infty e^{(g_Y - \delta_\pi)s - \int_t^{t+s} r^u du} E_t[\pi^{t+s}] \, ds \right) + (\zeta^t - 1) \sum_i P_i^t K_i^t \tag{26}$$

Second relation, part two: taking first differences. Now use (26) to compute

$$\frac{\phi(t)}{Y^{t-1,t}} \left(OMV^t - e^{-\delta_\pi \Delta t - \int_t^{t+\Delta t} r^u du} \cdot OMV^{t+\Delta t} \right) \tag{27}$$

for some small Δt, dividing by $Y^{t-1,t}$ (to be defined later, but known at time t) and multiplying by any deterministic function $\phi(t)$ of t. Expanding the term inside the parentheses in (27), we obtain

$$
\begin{aligned}
&OMV^t - e^{-\delta_\pi \Delta t - \int_t^{t+\Delta t} r^u du} OMV^{t+\Delta t} \\
&= \zeta^t \left(Y^t \int_0^\infty e^{(g_Y - \delta_\pi)s - \int_t^{t+s} r^u du} \mathbb{E}_t[\pi^{t+s}]\, ds \right) - \zeta^{t+\Delta t} \left(e^{-g_Y \Delta t} Y^{t+\Delta t} \int_{\Delta t}^\infty e^{(g_Y - \delta_\pi)s - \int_t^{t+s} r^u du} \mathbb{E}_{t+\Delta t}[\pi^{t+s}]\, ds \right) \\
&\qquad + (\zeta^t - 1) Y^t \sum_i P_i^t K_i^t + e^{-\delta_\pi \Delta t - \int_t^{t+\Delta t} r^u du} (\zeta^{t+\Delta t} - 1) Y^{t+\Delta t} \sum P_i^{t+\Delta t} K_i^{t+\Delta t}
\end{aligned}
\tag{28}
$$

Suppose now that we take the unconditional expectation of (27). Given the assumed independence of ζ^t, Y^t, and π^t, (28) simplifies dramatically and we are left with

$$
\mathbb{E}\left[\frac{\phi(t)}{Y^{t-1,t}} \left(OMV^t - e^{-\delta_\pi \Delta t - \int_t^{t+\Delta t} r^u du} \cdot OMV^{t+\Delta t} \right) \right] = \mathbb{E}\left[\frac{\phi(t)}{Y^{t-1,t}} Y^t \int_0^{\Delta t} e^{(g_Y - \delta_\pi)s - \int_t^{t+s} r^u du} \pi^{t+s}\, ds \right]
\tag{29}
$$

We can further manipulate (29), using the law of iterated expectations to obtain

$$
\begin{aligned}
\mathbb{E}\left[\frac{\phi(t)}{Y^{t-1,t}} Y^t \int_0^{\Delta t} e^{(g_Y - \delta_\pi)s - \int_t^{t+s} r^u du} \pi^{t+s}\, ds \right] &= \mathbb{E}\left[\frac{\phi(t)}{Y^{t-1,t}} \mathbb{E}_t \left[\int_0^{\Delta t} (e^{g_Y s} Y^t) e^{-\delta_\pi s - \int_t^{t+s} r^u du} \pi^{t+s}\, ds \right] \right] \\
&= \mathbb{E}\left[\frac{\phi(t)}{Y^{t-1,t}} \mathbb{E}_t \left[\int_0^{\Delta t} e^{-\delta_\pi s - \int_t^{t+s} r^u du} \pi^{t+s} Y^{t+s}\, ds \right] \right] \\
&= \mathbb{E}\left[\frac{\phi(t)}{Y^{t-1,t}} \int_0^{\Delta t} e^{-\delta_\pi s - \int_t^{t+s} r^u du} \pi^{t+s} Y^{t+s}\, ds \right]
\end{aligned}
\tag{30}
$$

Assuming that Δt is small enough and $\pi^{t+s} Y^{t+s}$ is sufficiently close to being continuous, we can approximate the integral inside (30) by

$$
\int_0^{\Delta t} e^{-\delta_\pi s - \int_t^{t+s} r^u du} \pi^{t+s} Y^{t+s} ds \approx \frac{1 + e^{-\delta_\pi \Delta t - (r^t + r^{t+\Delta t})/2}}{2} \int_0^{\Delta t} \pi^{t+s} Y^{t+s} ds
\tag{31}
$$

Full estimation strategy. We have shown that the unconditional expectation of (27), which we can approximate by

$$
\mathbb{E}\left[\frac{\phi(t)}{Y^{t-1,t}} \left(OMV^t - e^{-\delta_\pi \Delta t - (r^t + r^{t+\Delta t})/2} OMV^{t+\Delta t} \right) \right]
\tag{32}
$$

has unconditional expectation approximately equal to

$$
\mathbb{E}\left[\frac{\phi(t)}{Y^{t-1,t}} \cdot \frac{1 + e^{-\delta_\pi \Delta t - (r^t + r^{t+\Delta t})/2}}{2} \cdot \int_0^{\Delta t} \pi^{t+s} Y^{t+s} ds \right]
\tag{33}
$$

where according to (23) and (24), we can obtain the approximate flow of pure profits $\int_0^{\Delta t} \pi^{t+s} Y^{t+s} ds$ in (33) as

$$\begin{aligned} &\int_0^{\Delta t} \pi^{t+s} Y^{t+s}\, ds \\ \approx &\int_t^{t+\Delta t} (Y^s - w^s L^s - \sum_i \delta_i P_i^s K_i^s)\, ds - \sum_i \left((r^{t+\Delta t} - g_{P_i}) P_i^{t+\Delta t} K_i^{t+\Delta t} + (r^t - g_{P_i}) P_i^t K_i^t \right) \frac{\Delta t}{2} \end{aligned} \tag{34}$$

where the first term is just the recorded net return on capital in the period $[t, t + \Delta t]$ as measured in the national accounts, while the second term can be derived from the nominal quantities $P_i K_i$ of each type of capital.

The moments (32) and (33) are equal, and we can set the corresponding sample moments equal to each other. Generally I will look at an annual frequency, such that $\Delta t = 1$. Given a functional form for r^t with n free parameters to be pinned down, we can choose n functions for $\phi(t)$ to give us n sample moment conditions that determine those parameters.