Structural Demographic Drivers of Unit Labor Cost

Timothy R. McDade
timothy.mcdade@blackrock.com

BlackRock, Inc.

April 17, 2024

Abstract

It is well known that the price of labor relates to cyclical factors such as the unemployment rate and inflation, but its relationship with structural demographic factors has been less well investigated. Over the coming decades, this relationship could become more important: countries experiencing shrinking and aging populations may have tighter labor markets, meaning that the price of labor may be more responsive to noncyclical demographic characteristics of the labor force than to cyclical factors. I argue that the degree to which economies find other sources of labor – aside from decreasing unemployment – will affect the price of labor. I expect that the Unit Labor Cost (ULC) will increase more slowly when countries more efficiently use their demographic endowment for labor input into production. Moreover, this relationship should be stronger when labor markets are tighter: as an economy runs out of cyclical employment cushion, it must turn to other structural methods of adding hours worked to the economy. To ground the theory in history, I apply this logic to Japan's experience since 1990. Then, in a cross-national time-series analysis, I find evidence supporting both hypotheses. This model helps contextualize the relationship between demographics, growth, and inflation: demographic characteristics of the labor supply serve as a release valve for wage pressures. Moreover, this model connects the two main concepts affected by the price of labor: growth and price level. Equity investors, in particular, might be interested in the model's implications for an economy's free cash flow.

Contents

List of Figures

1 Introduction

It is well known that the price of labor relates to cyclical factors such as the unemployment rate and inflation. But its relationship with structural demographic factors has been less well investigated. Over the coming decades, this relationship could become more important: countries experiencing shrinking and aging populations may have tighter labor markets, meaning that the price of labor may be more responsive to noncyclical demographic characteristics of the labor force than to cyclical factors.

1.1 Argument

The relationship between unemployment and inflation is depicted in the Phillips Curve, which historically was a line with negative slope (Phillips, 1958). The intuition makes sense – as unemployment rises, inflation decreases because households have less disposable income to stimulate the demand side of the economy. Although this relationship held true for the mid-20th century, the Phillips curve has flattened in recent decades: inflation has fluctuated despite unemployment staying low. Despite the disintegration of the Phillips Curve, Unit Labor Cost (ULC) indeed has a relationship with inflation and unemployment in the US. ULC increases faster when unemployment drops and when inflation is higher. This is consistent with the expectations of standard literature.

But why has the Phillips Curve flattened in recent decades? St. Louis Federal Reserve President James Bullard recently chalked it up to more responsive monetary policy; Federal Reserve Chair Jerome Powell has attributed it up to increasingly strong inflation expectations (Engemann, 2020).[1] I offer another explanation: chronically tight labor markets have very little room to trigger cyclical changes in the price of labor. Instead, when unemployment is low, structural demographic characteristics of economies likely play an increasingly important role in determining the price of labor. Chronically tight labor markets contribute to a decay in the information content of cyclical employment measures to predict the price of labor; other measures of labor tightness instead contain more information.

The employed population is not the only group of people in an economy capable of working. It

[1] The latter seems endogenously determined.

is merely a subset of the labor force, which is a subset of a country's working age population. If economies find that their employed population is a near-exhaustive portion of their labor force, they will have to turn to sources outside their labor force to increase labor input (hours worked). I argue that, in tight labor markets, the degree to which economies find other noncyclical sources of labor will affect the price of labor.

1.2 Hypothesis

I expect that the cost of labor will increase more slowly when countries more efficiently use their demographic endowment. Moreover, I expect this relationship to be stronger when labor markets are tighter – as an economy runs out of cyclical room to maneuver,[2] it is left with other structural methods of adding hours worked to the economy: increasing participation rates, increasing immigration, raising effective retirement ages, and increasing hours worked per worker.

1.3 Empirical Setting

To test this hypothesis, I conduct a cross-national analysis for as long a time period as is available. Perhaps unsurprisingly, the countries with the most data availability are developed economies, mostly members of the OECD, and data reporting begins in earnest after World War II. However, the increasingly detailed international reporting conventions mean that measurement of relevant economic phenomena really picks up in the 1990s.

1.4 Empirical Model

This paper makes several empirical findings. First, I validate findings from prior literature: labor is more expensive when inflation is high and unemployment is low. Second, I find that structural demographic characteristics have a statistically significant negative relationship with ULC. As countries get better at extracting more hours from their working age population, the price of labor increases more slowly. Third, I find that this relationship persists when including an interaction term: tightening labor markets (decreasing unemployment) correspond to a stronger effect of structural demographic char-

[2]The "reserve army of labour" (Marx, 1867).

acteristics on ULC, and slackening labor markets (increasing unemployment) moderate their effect. These results support my hypothesis. I find that these results are robust to accounting for separate movements in the numerator and denominator of ULC as well as to including controls for alternate determinants of ULC.

1.5 Contributions

This model helps contextualize the relationship between demographics, growth, and inflation. Adjustments to an economy's labor supply have implications for its growth, but this model shows a specific channel through which they also have implications for the price of labor (and therefore inflation). Demographic characteristics of the labor supply serve as a release valve for wage pressures. Increasing hours worked can stem upwards pressure on wages, but once an economy can no longer add any more hours, it faces higher increases in the price of labor.

Moreover, this model connects the two main concepts affected by the price of labor: growth and price level. Equity investors in particular might be interested in the model's implications for an economy's free cash flow (and therefore its growth potential). Demographic characteristics and tight labor markets might put upward pressure on ULC – and on the labor share of income. Companies have two choices about how to respond. Their first option is to defend their margins, putting upwards pressure on price inflation. Their second option is to take a hit to their margins, decreasing returns to shareholders (i.e., decreasing the capital share of income) – putting downward pressure on free cash flow.

Better forecasts of hours worked (and therefore an economy's effectiveness at turning its demographic endowment into labor) will permit more accurate predictions of future ULC. To best take advantage of the theoretical linkage this model presents, next steps could include constructing granular, high-frequency forecasts of hours worked.

1.6 Roadmap

The next section of this paper will chronicle analytical approaches towards determinants of the cost of labor. The subsequent section will expand on my theoretical argument; the following section will introduce the empirical approach I use to adjudicate my hypothesis. Results will follow, and the final section concludes.

2 Background

It is well known that the price of labor relates to cyclical factors such as the unemployment rate and inflation. But the relationship that the price of labor has with structural demographic factors has been less well investigated. Over the coming decades, this relationship could become more important: countries experiencing shrinking and aging populations may have tighter labor markets, meaning that the price of labor may be more responsive to noncyclical demographic characteristics of the labor force than to cyclical factors.

2.1 An Economy's Production

A common starting point for modeling economic output is a standard Cobb-Douglas production function. This approach, first implemented by Solow (1956), has inspired numerous theoretical evolutions and underlies the reporting and bureaucratic collection of economic data today (Shackleton, 2018). Assuming diminishing marginal returns to capital, non-zero rates of labor growth, and non-zero technological progress, the Solow approach models economic output Q as a function of input factors labor L and capital K as well as the proportion of overall income that goes to each factor of production (α and $1-\alpha$). Factors not included in the model are described by the Solow Residual A, which commonly is referred to as "multi-factor productivity", and is usually defined as how efficiently production turns inputs into outputs.

$$Q = AK^{1-\alpha}L^{\alpha} \tag{1}$$

Since its publication, this model has inspired various evolutions. Among others, Cass (1965) built on the microfoundations of the Solow model to endogenize the household saving rate. Mankiw et al. (1992) and Topel (1999) analyze human capital as an alternative factor input and Romer (1994) endogenizes the processes that contribute to growth.

The most famous example of endogenous growth theory is Romer (1990), who expands on the thought that ideas are nonrival – infinitely usable – and therefore are a key driver of long-run economic growth. Jones (1995) evolved these ideas in semi-endogenous growth theory, which argues that while technological change itself may be endogenous, long-run economic growth is still driven by exogenous population growth.

In these models, the diminshing marginal returns to factor inputs imply "convergence": countries with lower GDP per capita, who are on the steeper part of the growth curve, will eventually catch up with countries with higher GDP per capita, who are seeing far slower increases in output because of diminshing marginal returns.

Other explanations of economic growth focus on the role that institutions play in resource allocation. Robinson and Acemoglu (2012) argue that political institutions affect countries' long-run growth trajectories. Banerjee and Duflo (2005) argue that convergence does not occur (i.e., macro growth theories fall apart at the micro level) because "the assumption of optimal resource allocation fails radically". They argue that in poor countries "government failures, credit constraints, insurance failure, family dynamics, and behavioral issues" make returns on investment lower and more variable, which depresses economic growth there (Banerjee and Duflo, 2005: p. 5).

2.2 Cyclical Drivers of the Cost of Labor

But none of these models question the importance of labor to economic output. Although the endogenous and semi-endogenous growth literatures assume stable or growing populations, they acknowledge that the main drivers of labor costs are how much labor there is and how much employers demand it. Demand for labor is driven by the production function, and supply is driven by processes often taken as exogenous but that can be policy-driven (immigration, productivity, demographics). Attempting to isolate the causes of the cost of labor requires a time-series cross-sectional empirical approach.

The major avenue for pricing the cost of labor has been cyclical. A literature focused on economic cycles of booms and busts has analyzed the relationship between the employment rate and the cost of labor, and found that scarcer labor (lower unemployment) corresponds with higher labor cost and vice versa (Tobin, 1995; Mankiw, 2001; Berentsen et al., 2011).

The relationship between unemployment and inflation is depicted in the Phillips Curve. Phillips found in the late 1950s that the UK had a stable negative relationship between unemployment and wages (Phillips, 1958). The intuition makes sense – as unemployment rises, inflation decreases because households have less disposable income to stimulate the demand side of the economy, and firms produce less supply of goods and services, so the economy reaches a new equilibrium at a lower price point. Although this relationship held true for the United States in the mid-20th century, the Phillips curve has flattened in recent decades: inflation has fluctuated despite unemployment staying low.

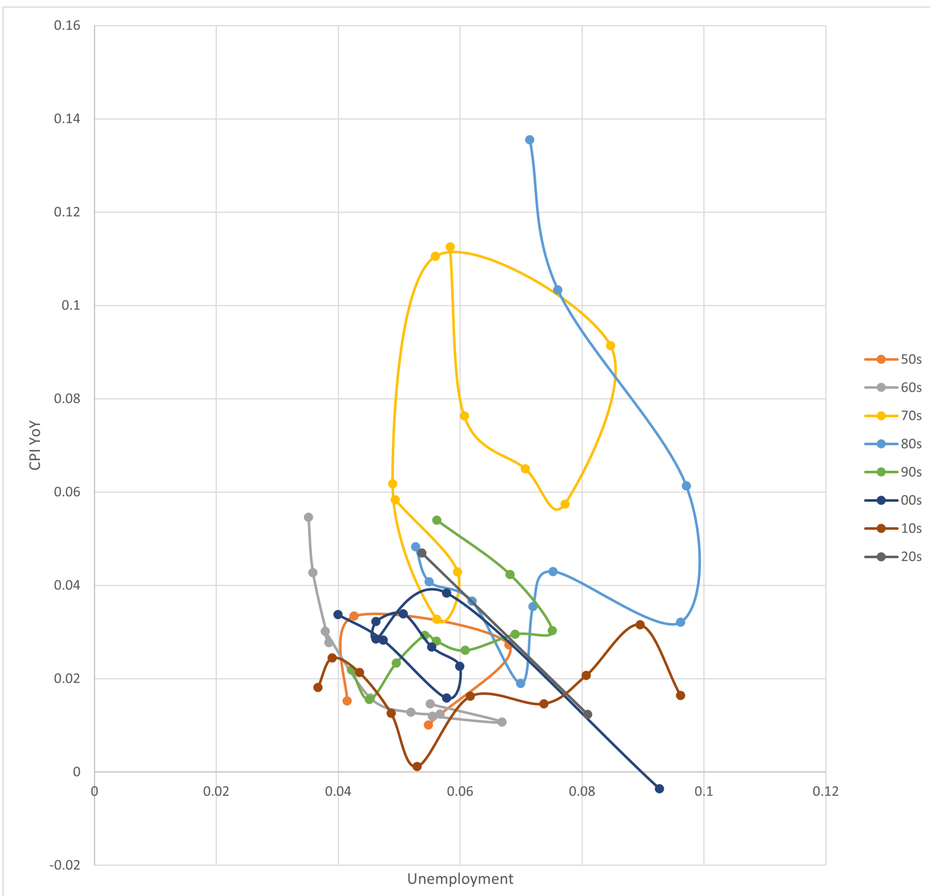

Figure 1: The Phillips Curve in the United States, 1955 to 2022 (Source: OECD, as of 1 August 2023).

Figure 1 shows the Phillips Curve in the US from the 1950s through 2022. Although it followed perfectly its theorized form in the 1960s, the neat relationship has since broken down. The economic crises of the 1970s culminated in stagflation, which persisted through the early 1980s (Friedman, 1977). In the 1990s, the curve took roughly the reverse of its theorized form, and performed various acrobatics through the 2000s. Over the 2010s, it has been roughly flat.

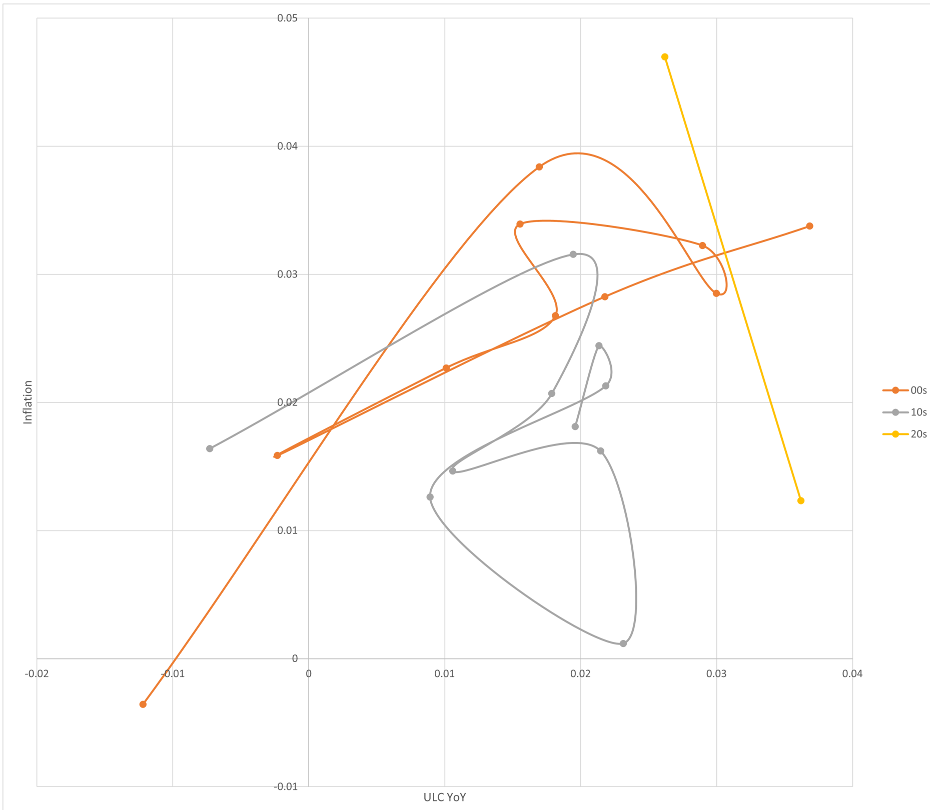

Figure 2: Unit Labor Cost vs. inflation in the USA, 2000-2022 (Source: OECD, as of 1 August 2023).

Why has the Phillips Curve flattened in recent decades? St. Louis Federal Reserve President James Bullard recently chalked it up to more responsive monetary policy; Federal Reserve Chair Jerome Powell has attributed it up to increasingly strong inflation expectations (Engemann, 2020).[3] Despite this wobbly adhesion to the theorized relationship, unemployment and inflation each hold down relatively well-behaved relationships with Unit Labor Cost. Figures 2 and 3 show despite the disintegration of the Phillips Curve, Unit Labor Cost indeed has a relationship with inflation and unemployment in the US. ULC increases faster when unemployment drops and inflation is higher.

[3]The latter seems endogenously determined.

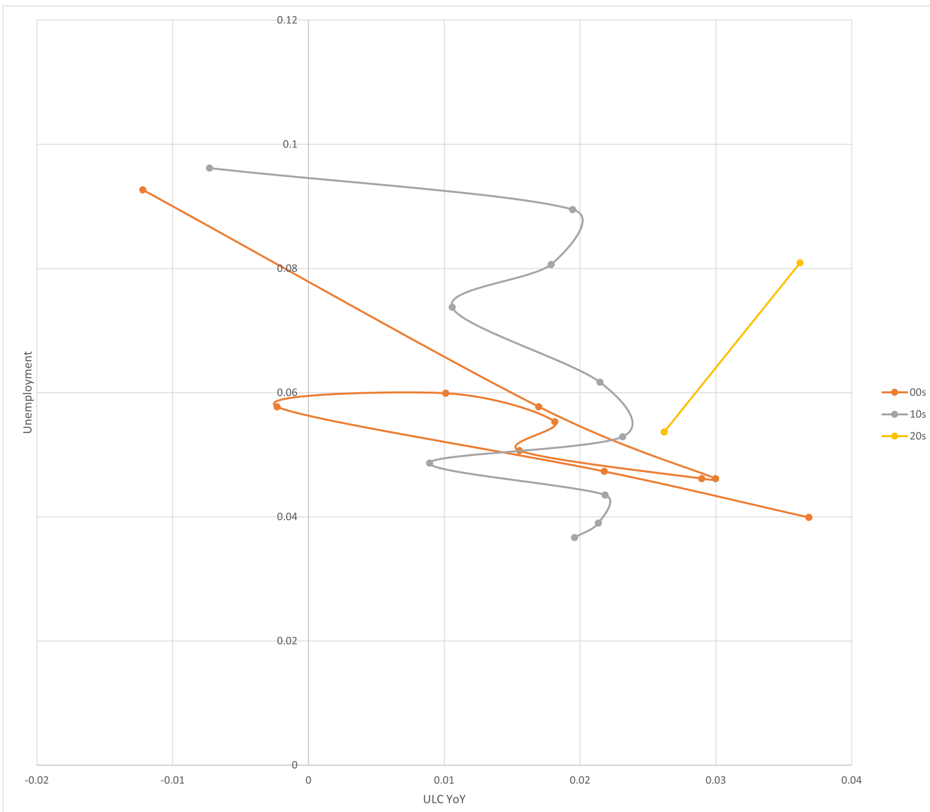

Figure 3: Unit Labor Cost vs. unemployment in the USA, 2000-2022 (Source: OECD, as of 1 August 2023).

Perhaps not surprising at first because of its consistency with the expectations of standard literature, this relationship suggests that ULC could be a channel relating the two variables. Notwithstanding monetary policy, inflation expectations, or other factors, I offer another explanation of the flattening of the Phillips Curve: tight labor markets have very little room to trigger cyclical changes in the price of labor. Instead, structural demographic characteristics of economies and labor markets likely play an increasingly important role.

2.3 Demographic Endowments

Demographics exert pressure on the labor supply in two waves. First, countries naturally face certain demographic trends simply due to their pre-existing population (demographic endowment). As fertility rates drop across the world and life expectancy increases, populations age and shrink. This phenomenon is most prevalent in more advanced economies, but is starting to take place in developing economies as well. Figure 4 shows fertility rates for G20 countries since 1950: all have dramatically decreased, and the vast majority are now well below the replacement rate of 2.1 children per woman.

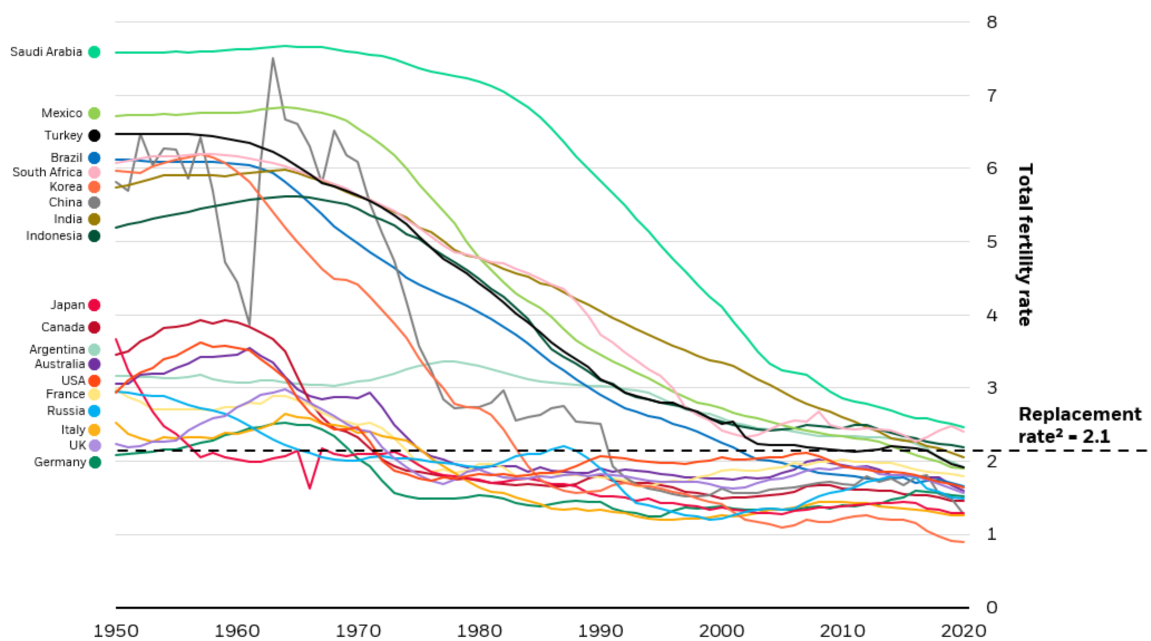

Figure 4: G20 Fertility Rates, 1950 to 2022 (Source: United Nations, as of 1 August 2023).

Figures 5 and 6 show the way these fertility patterns cascade through to working age populations (ages 15-64). Absent migration and working past 65, the working age population is the absolute maximum number of workers an economy could have: the labor force is a subset of working age population, and the employed population is a further slice of the labor force. Figures 5 and 6 have two main takeaways. First, several countries with advanced economies had working age populations that peaked well before 2020 and will continue to shrink in the coming decades. Second, countries whose working age populations grew through 2020 could shrink soon – but even if they keep growing, they will do so far more slowly than they did in the latter half of the 20th century.

These figures use the United Nations ”medium” projection, but serious arguments exist that such a projection is too optimistic, and something resembling the ”low” fertility projection will in fact come to pass.[4] Bricker and Ibbitson call this the “Empty Planet” scenario.

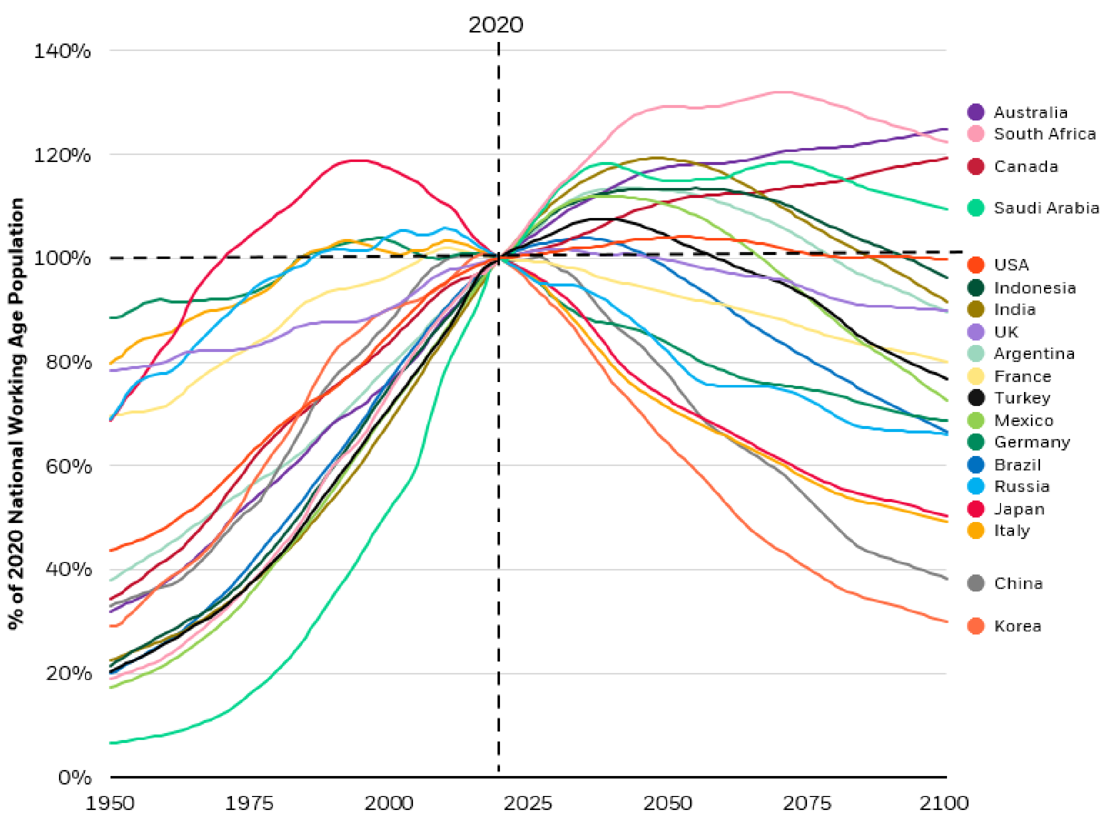

Figure 5: G20 Working Age Populations indexed to 2020, 1955 to 2100 (Source: United Nations, as of 1 August 2023).

[4]The United Nations Medium Projection assumes medium fertility and mortality trends.

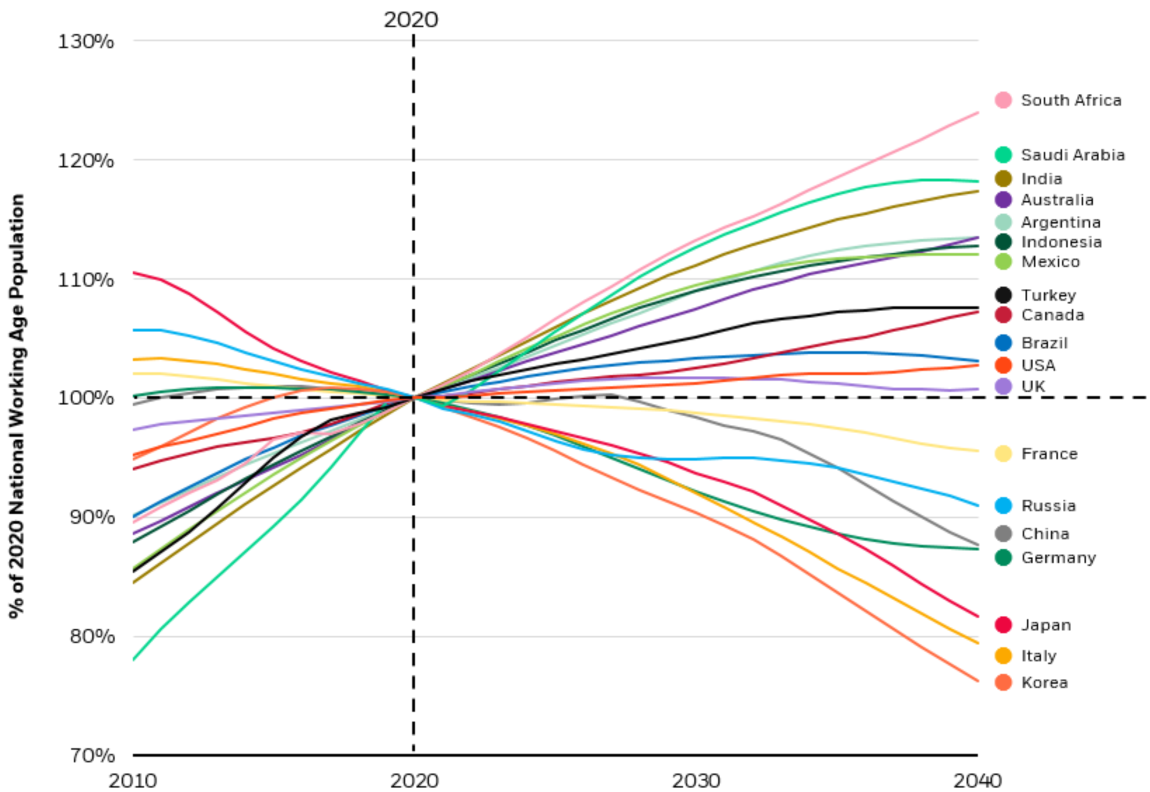

Figure 6: G20 Working Age Populations indexed to 2020, 1980 to 2040 (Source: United Nations, as of 1 August 2023).

Yet most of the models upon which our understanding of economics rests, dating back to Solow (1956), depend on the assumption that populations are rising – or at least stable. The preceding figures show that this assumption clearly is not reflective of the state of today's world, and especially not in advanced economies. Jones (2022) takes the "Empty Planet" scenario seriously and models possible equilibria from an economic growth perspective.[5] This paper takes inspiration from Jones, and instead considers the effect that structural demographic pressures have on the price of labor.

2.4 Country Adaptations to Demographic Pressures

Towards the end of his paper, Jones (2022) briefly touches on the ways that countries can implement policies that help jump towards a favorable long-term growth equilibrium. Countries don't have to sit by and accept their demographic fate; they can adapt to the cards they've been dealt. There are sevearl ways countries can do so.

The employed population is not the only group of people in an economy capable of working. It is merely a subset of the labor force, which is a subset of a country's working age population. If economies find that their employed population is a near-exhaustive portion of their labor force, they will have to turn to other sources to increase the labor input (hours worked).

There are four main structural levers countries have and one cyclical lever. The four structural levers are migration, labor force participation, effective retirement ages, and hours worked per worker. The cyclical lever is unemployment.

- Countries can increase immigration, thereby increasing the number of people in their country – and the number of workers in their economy. They can also offshore existing production to other countries where wages are lower, reducing their need to compete for scarcer local labor.
- The labor force participation rate (LFPR) is the proportion of the labor force that is employed or actively seeking employment.
- The effective retirement age is the average age at which a population exits the work force. Importantly, this is a computed descriptive statistic; it is different from the statutory age at which

[5] Other recent work has considered the same topic from a different point of view – see, among others, Bucci (2023).

workers are eligible to draw upon their social benefits (e.g., Social Security). By encouraging workers to stay working for longer, countries can extract more hours from them, increasing the economy's supply of labor.

- Increasing the number of hours that a typical worker works in a given year could add more hours to the economy, but is politically fraught. Such a measure is also made more complicated by the fact that one way to get older workers to stay in the labor force for longer is to offer part-time work, which drags down the hours worked per worker figure.

The aforementioned levers structurally affect the number of people in the labor force, but some may remain unemployed, still not contributing hours worked to the economy. By decreasing unemployment, a country can ensure its economy is running hot.

3 The Theoretical Model

3.1 Implications for Growth

Crucially, all of these levers affect the quantity of labor in the economy. And adjustments in quantity have implications for price. Because of its position in the economy as a factor input to production, labor has implications for growth – but because of its role as a source of income for households, it has implications for inflation.

Labor is one of the major factor inputs to economic growth. Economic models have long recognized this – official government calculations now quantify labor input using very detailed methodology (Shackleton, 2018). Yet the important takeaway is that labor is measured in the number of hours worked in a country's economy in a given period of time. All else equal, aging and shrinking labor forces mean fewer hours worked, which drags down output. More hours worked, on the other hand, means more output.

3.2 Effect on the Price of Labor

Movements in the quantity of the labor supply also affect the price of labor. Usually, this occurs on a cyclical basis, but I argue that structural demographics will likely affect the price of labor, especially in tight labor markets.

In countries whose working age populations are aging and shrinking, but are not yet adapting, the quantity of labor is shrinking slowly, meaning upwards pressure on wages because workers have more bargaining power. If supply does not increase along with wages, firms face upwards pressure on the cost side of their income statement without correspondingly increasing revenues and attempt to defend their margins by passsing this cost on to the consumer.

But countries that face demographic pressures can implement structural adaptations discussed above, any of which increase the supply of labor, thereby putting upward pressure on wages. Accordingly, each of the aforementioned structural demographic levers has implications for the balance between supply and demand:

1. Adding labor through immigration adds new consumers and new producers to the economy, leaving open the possibility that labor supply and demand for labor stay in equilibrium with no net effect on prices. Offshoring labor releases pressure on domestic wages but does not affect consumption.

2. Increasing labor force participation turns a slice of the population that is already in the economy consuming into producers, meaning a net deflationary impulse. Labor force participation dynamics have deep and complex social and political dependencies.

3. Raising effective retirement ages by enabling workers to work until later in their lives provides a net deflationary impulse, as people who would otherwise stop producing keep doing so for longer.

4. Increasing hours worked per worker maintains the current structure of the labor force but adds more hours, putting downward pressure on inflation.

And then, on top of all these structural factors, cyclical employment patterns affect which of the im-

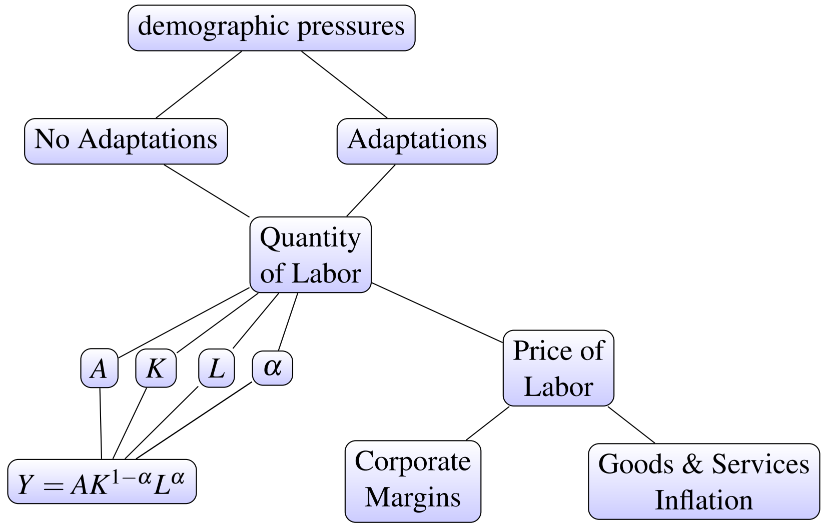

migrants, older workers, and new labor force entrants actually have jobs and work positive hours. Increasing the proportion of the now-bigger labor force that is working by lowering the unemployment rate only serves to further increase the quantity of labor in the economy, putting even further downward pressure on prices.

But each of these levers have a logical or political limit. Politics will only support so many immigrants entering a country. It's hard to make people work full-time past the age of 68 or so. Labor force participation rates can only ever be 100%, and usually don't get much above 85% for men (and slightly lower for women). When it is no longer possible for an economy to add much more to the labor supply, the slowing growth of the quantity of labor results in faster upwards pressure on wages.

3.3 Implications for Equities

The downstream effects of ULC have long been of core interest to central banks with dual mandates to stabilize prices and reduce unemployment. As Chair of the US Federal Reserve, Alan Greenspan commonly argued that corporations defend their margins. When the price of inputs rise, corporations try to pass cost increases along to customers. According to Greenspan, the question wasn't whether they will pass along the costs – the question was whether they could make a price increase stick. If price increases stuck, inflation could take hold. If price increases didn't stick, corporations go back to the drawing board and try to find another way to defend their margins (Bank of International

Setlements, 1999).

Unsurprisingly, this logic has implications for equity investors, who are essentially investing in the profit margin of a company (assuming constant revenues).

3.4 Implications for Price Inflation

The mechanism of interest is labor's function as the major source of income for households. As households earn more labor income, they have more disposable income to consume or save. Assuming overhead expenses and savings patterns remain constant, households can channel this greater disposable income into consuming goods and services, increasing demand and providing upwards pressure on prices.

3.4.1 Generalized Elasticity of Substitution

We define Unit Labor Cost as

$$ULC = \frac{\text{Labor compensation/hour}}{\text{output/hour}}$$

Movements in ULC have two parts: the numerator and denominator. The numerator (wage payments) has a fairly straightforward relationship with structural demographics: wages should decrease when countries make better use of their demographic endowment, increasing the labor supply.

The denominator, labor productivity, is trickier. An economy that follows a Cobb-Douglas production function assumes that the elasticity of substitution between input factors is one. In such a model, when wages go up, cheaper capital is simply substituted for scarcer (and more expensive) labor, leaving output constant. Likewise, Cobb-Douglas implies that when hours worked increases and wages decrease, economies will substitute cheaper labor in for relatively more expensive capital, keeping output level the same but decreasing output per hour worked. Labor productivity should be decreasing in hours worked and increasing in wages.

But in practice, the suitability of the Cobb-Douglas analytical approach depends on an economy's realized elasticity of substitution: the degree to which a substitution of capital for labor affects output. When capital is perfectly substitutable for labor, output may remain constant but lower labor input leads to higher labor productivity (output per hour worked). This is not always what happens in all sectors of an economy. When capital is not perfectly substitutable for labor, output increases more slowly with less labor input: lower marginal productivity of labor. An economy's sector mix will imbue it with its own overall elasticity of substitution.

In general, when wages go up relative to cost of capital (rents), common economic theory expects the labor share of income to increase. This is what happens in the US, and is the case because the US economy has an elasticity of substitution greater than one. But in Japan, the opposite occurs – because elasticity of substitution is less than one, a relative increase in wages leads the labor share of income to drop and the capital share to increase (Hirakata et al., 2018).

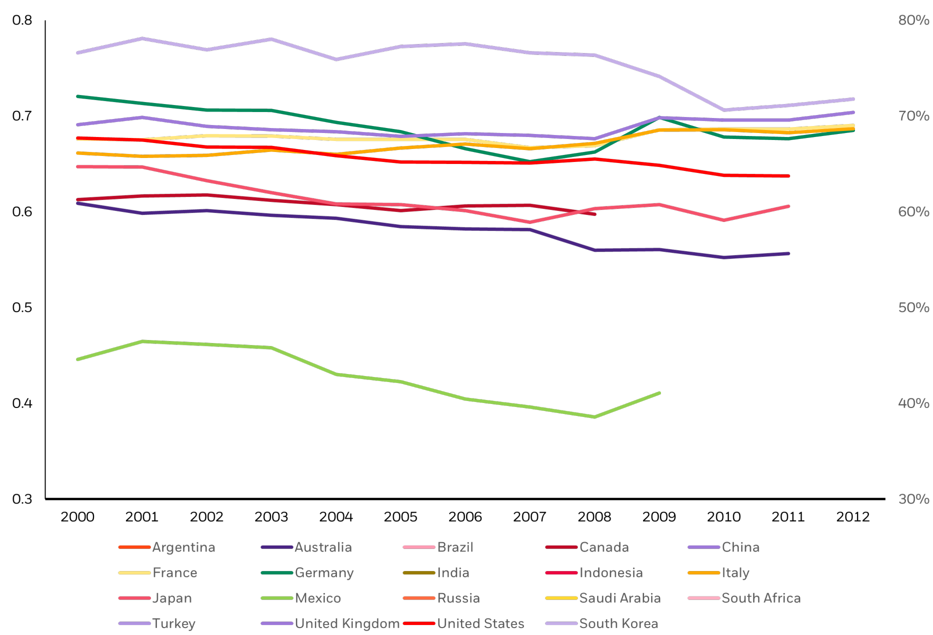

Figure 7: Labor Share of Income, selected G20 countries (Source: Penn World Tables, as of 1 August 2023).

Over the last several decades, economic literature has amassed evidence rejecting the Cobb-Douglas assumption of unity elasticity of substitution. Hirakata et al. (2018) apply this idea to the US and Japan from 1985-2017, arguing that a non-Cobb-Douglas production function (with non-unity Elasticity of Substitution) better reflects the way that substitution occurs in US and Japanese economies. Figure 7 shows further evidence of this: the labor share of income over time has not been constant. Figures 11 and 10 show a third way of telling that economies are not Cobb-Douglas: by looking at ULC.

How does ULC's denominator, labor productivity, respond if elasticity of substitution is not unity? Then the effect of changes in the labor supply on output should vary by economy, and the denominator of ULC should move separately from the numerator. These movements merit considering the numerator and denominator of ULC separately.

3.5 Hypothesis

I expect that the cost of labor will grow more slowly when countries more efficiently convert their structural demographic endowment into hours worked in their economy. This should manifest most directly in downward pressure on wages, unit labor cost's numerator, but I also expect it to have an increasing effect on output and hence labor productivity, ULC's denominator. The combination of the two means that ULC should grow more slowly when countries structurally extract more hours from their populations. I expect this relationship to be stronger when cyclical labor markets are tighter.

4 Empirics

To adjudicate these hypotheses, I take a multi-faceted empirical approach. First, I walk through a case study of how these concepts have manifested in Japan since 1990. Second, I conduct a cross-national time-series regressions to examine broader relationships.

The first part of this cross-national time-series analysis is modeling the effect of structural demographic pressures on ULC itself – a high bar to meet. Not only could structural demographic increases to the labor supply put pressure on the numerator of ULC (wages); they could also plausibly affect the denominator (productivity). Any increased output resulting from more hours worked could increase the denominator of ULC, putting overall downward pressure on ULC. For ULC growth to slow – my hypothesized effect – the wage effect of increasing labor supply would have to be as strong or stronger than its output effect. So any robust relationship between structural demographics and ULC would pass a high bar.

But, as noted above, the degree to which output increases with changes in labor depends on the economy's elasticity of substitution. I take two steps to more explicitly model the possibility that denominator effects are driving ULC results. First, I include various measures of productivity in these empirical specifications to account for movements in ULC's denominator. Second, I run the same empirical models analyzing the same independent variables' effect on labor compensation per hour worked – the numerator of ULC.

The third and fourth empirical tests are designed to address the long and variable lags of these interrelated macroeconomic concepts. I dive deeper into the time-series nature of the data by doing an AutoRegressive Moving Average analysis of ULC and a Vector AutoRegression designed to capture the way that these relationships persist over time.

4.1 Japan: A Case Study

I begin with an examination of this above theoretical framework in the context of stylized facts from Japan. For decades, Japan has been widely known as an aging society. Its working age population peaked in 1994 and has shrunk ever since; its total population peaked in 2010. It is truly an aging society with a shrinking population. In addition to these structural pressures, Japan faces different cyclical constraints than other countries. When Japanese firms are under pressure, they do not fire workers; instead, they find other outlets to relieve pressure on their balance sheets (Goodhart and Pradhan, 2020). This approach to employment results in chronically tight labor markets. Indeed, Japanese unemployment has not risen above 6% since at least 1970, and has spent the majority of the time under 3%.[6] Modern Japan provides prime conditions to examine the effect that structural demographic characteristics have on ULC.

Since Japan's demographic endowment is particularly unfavorable, we might expect that it experienced a scarcity of labor since the peak of its working age population. In fact, Figure 8 shows that

[6] The European Union, on the other hand, is just now experiencing chronically tight labor markets with unemployment just below 7% (Arnola, 2023).

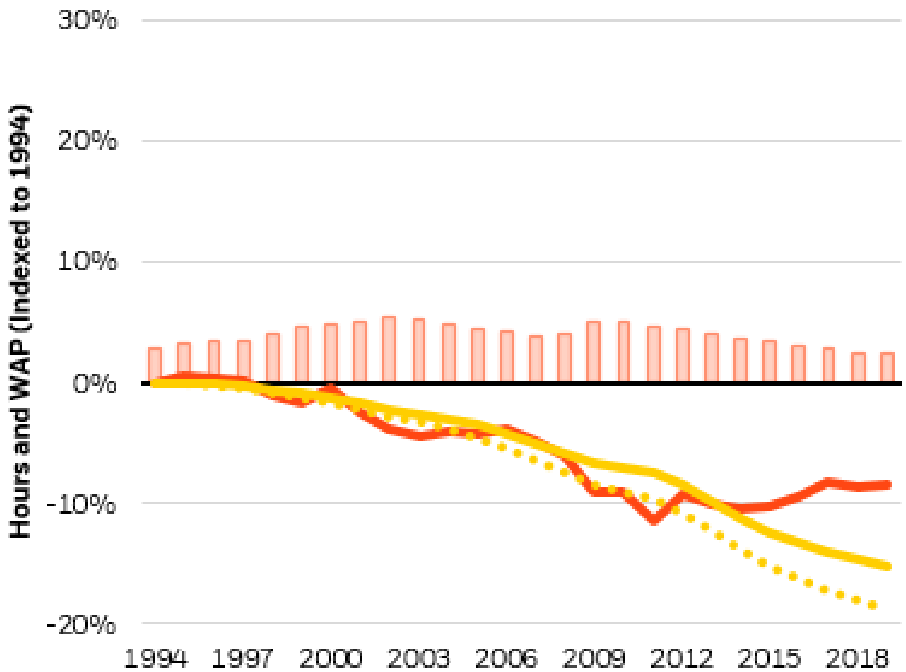

Figure 8: Working Age Population and Hours Worked in Japan (Source: United Nations & OECD, as of 1 August 2023).

until about 2010, the hours worked in its economy (the orange line) shrank alongside its dropping working age population (the solid yellow line). Japan extracted neither more nor less hours from its working age population than trend; it merely accepted the demographic cards it had been dealt.

But in 2010, things changed. The number of hours worked in the Japanese economy stabilized despite the working age population continuing to fall. Figure 9 shows that the main contributors were increases in female LFPR and female effective retirement age (both in orange; male are yellow and overall is black).

The theoretical framework outlined above might suggest that these movements in labor supply would precipitate changes in ULC. Figure 10 shows Japan's ULC movements alongside Germany and the US for comparison. The Bank of Japan attributes the drop in ULC to an economy-wide shift to non-regular work, a term of art which corresponds to the western concept of "salaryman" (Bank of Japan Research and Statistics Department, 2005: p. 31). This fits with our framework: a shift to non-regular work is made possible by companies' ability to move factors of production across

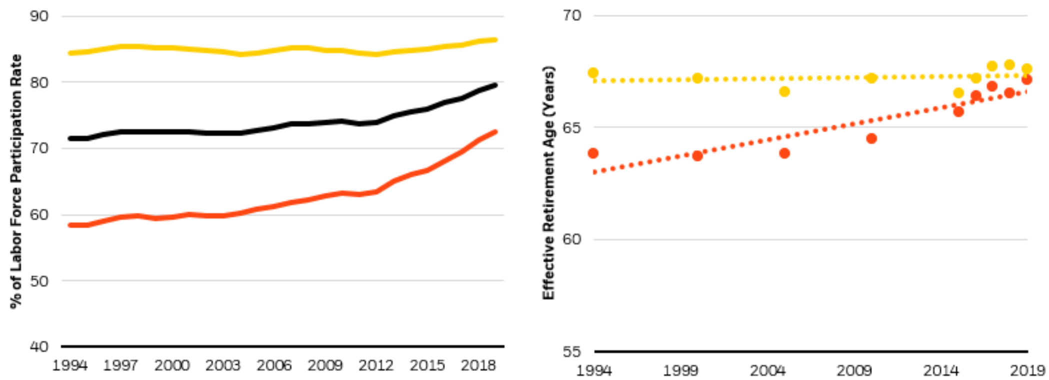

Figure 9: Labor Supply Adaptations in Japan (Source: United Nations & OECD, as of 1 August 2023).

international borders. If an economy has more jobs than people, countries can bring in immigrants, moving labor to where the capital is, or shift jobs overseas, moving capital to where the labor is. Throughout the 1990s and 2000s, Japanese firms offshored jobs to China, where labor was plentiful and cheap (Goodhart and Pradhan, 2020: Ch. 9). Offshoring labor became a release valve for the upwards pressure on the price of labor resulting from demographics that firms faced domestically. Because Japanese firms prefer not to fire employees, they instead offshored jobs as a way to pay local employees less rather than firing them. This approach was made possible because firms did not have to bid domestically for scarcer labor. The model they used to pay domestic workers less was to shift full-time workers to non-regular work.

4.2 Cross-National Empirical Setting

To test this hypothesis on a broader scale, I conduct a cross-national analysis for as long a time period as is available. Perhaps unsurprisingly, the countries with the most data availability are developed economies, mostly members of the OECD, and data reporting begins in earnest after World War II. However, the increasingly detailed international reporting conventions mean that measurement of the economic phenomena that interest us really picks up in the 1990s.[7]

There is theoretical justification for restricting the scope of analysis to only countries with strong

[7]The analysis covers Canada, Switzerland, Chile, Colombia, Germany, Denmark, Spain, France, Italy, Japan, South Korea, the Netherlands, Russia, Sweden, the United Kingdom, and the United States from 1950 to 2022.

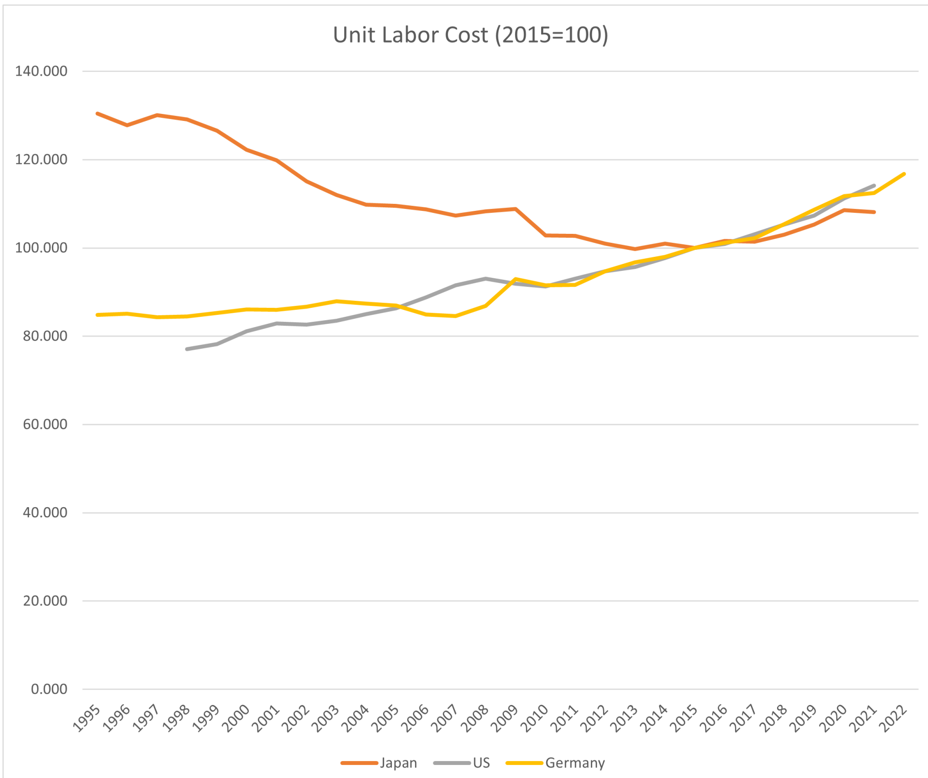

Figure 10: Unit Labor Cost in Japan, Germany, and the US (Source: OECD, as of 1 August 2023).

reporting quality. Because this theory is designed to pertain to the way an economy is using its total labor supply, the existence of a large informal labor sector may pollute the accuracy of results. When workers move from the formal to informal sector – perhaps in tough economic times – they appear to drop out of the labor force entirely, as far as official statistics are concerned. To avoid this data quality issue, I restrict the scope of analysis to countries without large informal sectors.[8]

4.3 Data Description

I compile data using Haver Analytics' API. I source national accounts variables from the OECD. Various measures of labor cost, contributions to GDP growth come from OECD National Accounts data. Effective retirement age, labor force participation, union density, and unemployment data come from the OECD Labor Force Survey. Inflation and exchange rates data come from OECD Main Economic Indicators. I source tariff measures and the labor force's skilled/unskilled mix from the World Bank's World Development Indicators. I pull central bank policy rates from the Bank of International Settlements, and all population data from the United Nations Population database.

Table 1 shows the mean, standard deviation, and inter-quartile range of the data points used in the analysis. Figure 11 shows that ULC mostly increases fairly evenly across countries, with the notable exception of Japan. Because Japan's ULC dynamics are explainable using the above logic, I include it in the empirical analysis as a high bar for the empirical model to clear.

4.4 Empirical Model

I use a simple OLS with HC1 standard errors to predict the annual growth rate in ULC υ in country c in time period t, including a lagged dependent variable to account for autocorrelation. I run four specifications. The base specification includes as predictors variables typically associated with ULC: unemployment u, inflation π, and the the central bank policy rate (CBPR) i.

The second specification adds a structural demographic variable: the gap between a country's annual growth rate of total hours worked and the growth rate of its working-age population (called Δ, presented as a first difference). Δ represents a country's ability to grow the labor input to its economy

[8] I exclude Mexico and Israel.

Table 1: Descriptive Statistics: Cross-National Analysis

	count	mean	std	min	25%	50%	75%	max
A_Contrib	487.0	0.01	0.01	-0.04	-0.00	0.01	0.01	0.07
CBPR	628.0	0.05	0.14	-0.01	0.01	0.03	0.06	2.10
CPI	1010.0	55.46	36.46	0.00	19.33	58.06	89.83	140.13
Delta_FD	792.0	-0.09	2.53	-24.53	-1.02	0.18	1.02	17.88
Labor_Advanced_Educ	401.0	0.79	0.05	0.66	0.76	0.80	0.83	0.90
Tariff_Mean	479.0	0.05	0.03	0.01	0.03	0.04	0.05	0.19
ULC	471.0	0.03	0.04	-0.07	0.01	0.02	0.04	0.37
ULC_lag1	463.0	0.03	0.04	-0.07	0.01	0.02	0.04	0.37
ULC_r5_SD	409.0	0.02	0.01	0.00	0.01	0.01	0.03	0.09
USD_LC	783.0	0.94	0.75	0.00	0.18	0.98	1.34	3.32
Unemp	942.0	0.06	0.04	0.00	0.03	0.05	0.08	0.26
Unions	799.0	0.32	0.20	0.09	0.17	0.27	0.37	0.87
Year	1168.0	1986.00	21.08	1950.00	1968.00	1986.00	2004.00	2022.00

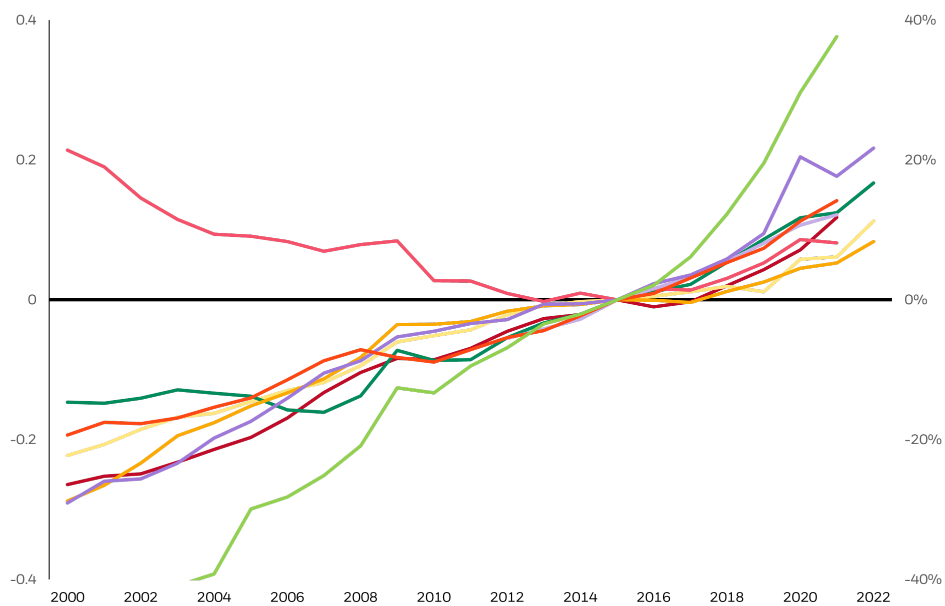

Figure 11: Selected G20 Unit Labor Costs over time (Source: OECD, as of 1 August 2023).

commensurate with its working-age population (its demographic endowment, or the ceiling of its labor force).

The third specification incorporates an interaction term between Δ and unemployment, designed to test the hypothesis that ULC responds more strongly to changes in HW-WAP under tighter labor markets.

The fourth specification includes all the aforementioned variables as well as a vector ζ of measures of other possible explanations of changes in ULC: tariffs, exchange rates, the skilled/unskilled profile of the labor force, growth in TFP contribution to GDP growth, and prevalence of trade unions. The models follow the below form:

$$\upsilon_{ct} = \alpha + \beta_0 u_{ct} + \beta_1 \pi_{ct} + \beta_2 i_{ct} + \beta_2 \Delta_{ct} + \beta_3 \Delta_{ct} u_{ct} + \vec{\beta}_4 \zeta_{ct} + \varepsilon_{ct} \tag{2}$$

Prior literature suggests that β_0 will be negative and β_1 and β_2 will be positive. My theory would find support from a negative value of β_2 and a positive value of β_3.

4.5 OLS Results

Table 2 shows the results. The first specification validates findings from the main school of existing literature: labor is more expensive when unemployment is low and inflation is high. The coefficient for unemployment is much larger than that of inflation, suggesting that most changes in ULC derive from unemployment rather than inflation on its own.

Adding in Δ preserves these relationships, both in size and sign, but reveals a statistically significant negative relationship between Δ and ULC. The negative sign suggests that as countries get better at extracting more hours from their working age population, the price of labor increases more slowly. That its magnitude is smaller than unemployment's coefficient makes sense for a structural, as opposed to cyclical, factor.

This relationship persists when including an interaction term. But the interaction term's positive and statistically significant nature suggests that tightening labor markets (decreasing unemployment) cor-

Table 2: Structural and Cyclical Predictors of ULC

	Dependent variable: Unit Labor Cost AGR			
	Cyclical	Add Structural	Interaction	More Controls
	(1)	(2)	(3)	(4)
Intercept	0.072***	0.073***	0.071***	-0.004
	(0.018)	(0.018)	(0.018)	(0.052)
ULC Lag	0.329***	0.312***	0.307***	-0.071
	(0.061)	(0.059)	(0.058)	(0.051)
ULC 5Y S.D.				0.293**
				(0.115)
Inflation	-0.000***	-0.000***	-0.000**	0.000*
	(0.000)	(0.000)	(0.000)	(0.000)
CBPR	0.106	0.134	0.144	0.475***
	(0.087)	(0.089)	(0.091)	(0.078)
Tariffs				0.017
				(0.092)
USD/LC				-0.001
				(0.005)
Skilled				0.044
				(0.050)
Unions				-0.082*
				(0.044)
TFP Growth				-0.815***
				(0.059)
Unemployment	-0.278***	-0.310***	-0.302***	-0.351***
	(0.057)	(0.062)	(0.060)	(0.043)
Δ		-0.002**	-0.004***	-0.005***
		(0.001)	(0.001)	(0.001)
Interaction			0.020*	0.027***
			(0.011)	(0.008)
Country FE	Yes	Yes	Yes	Yes
Observations	357	357	357	242
R^2	0.697	0.710	0.715	0.684
Adjusted R^2	0.680	0.692	0.697	0.649
Residual Std. Error	0.021	0.021	0.021	0.010
F Statistic	81.287***	86.610***	85.283***	25.324***
Note:			*p<0.1; **p<0.05; ***p<0.01	

respond to a stronger effect of of Δ on ULC, and slackening labor markets (increasing unemployment) moderate its effect. This result supports my hypothesis.

Alongside the above case study, the fixed effect coefficient for Japan is negative and of quite large magnitude, suggesting that the baseline annual growth rate of ULC in Japan – with all other variables set to zero – is negative. Although not shown in this table, such a result supports the intuition that Japan's unique combination of macro and micro factors result in an environment where ULC would experience the most severe downward pressure, as shown in Figure 10.

By its construction, ULC increases more quickly when productivity grows more slowly. To allow for the possibility that movements in productivity (the denominator of ULC) are driving results, I include TFP contribution to GDP growth as an independent variable in the final regression specification. The large, negative, statistically significant coefficient for TFP growth is a noteworthy result, and corresponds with theoretical expectations: with productivity growing more quickly, ULC grows more slowly. Yet despite this relationship, the statistically significant effect of Δ on ULC persists. These results also hold when measuring productivity with year-over-year growth in labor productivity (measured in output per hour worked).

4.6 Robustness

The same empirical approach is robust to using alternate measures of the price of labor as well as including controls for alternate determinants of unit labor cost. However, because the inclusion of such variables does not affect the substance of our results, and requires loss of data observations, I prefer the third specification.

This empirical model finds a positive and statistically significant coefficient for TFP growth – and the same relationship holds when measuring productivity with labor productivity (output per hour worked) instead. When productivity grows more quickly, ULC also increases more quickly (all else equal). This could be initially interpreted as labor receiving a share of productivity gains.

The flipping of productivity's sign in these two models shows that productivity's relationship with ULC has the opposite sign of its relationship with labor cost per hour. Alternatively, it could suggest

Table 3: Structural and Cyclical Predictors of Labor Cost per Hour Worked

	Dependent variable: Unit Labor Cost AGR			
	Cyclical	Add Structural	Interaction	More Controls
	(1)	(2)	(3)	(4)
Intercept	0.095***	0.093***	0.094***	0.026
	(0.019)	(0.014)	(0.013)	(0.060)
LC/HW Lag	0.271**	0.315***	0.322***	0.022
	(0.109)	(0.066)	(0.063)	(0.052)
LC/HW r5 SD				0.237
				(0.158)
Inflation	-0.001***	-0.001***	-0.001***	0.000
	(0.000)	(0.000)	(0.000)	(0.000)
CBPR	0.057	0.110	0.102	0.555***
	(0.080)	(0.078)	(0.078)	(0.078)
Tariffs				-0.009
				(0.096)
USD/LC				-0.003
				(0.005)
Skilled				0.010
				(0.058)
Unions				-0.068
				(0.047)
TFP Growth				0.190***
				(0.067)
Unemployment	-0.188**	-0.291***	-0.300***	-0.315***
	(0.080)	(0.053)	(0.057)	(0.044)
Δ		-0.005***	-0.003**	-0.006***
		(0.001)	(0.001)	(0.001)
Interaction			-0.023	0.019**
			(0.014)	(0.008)
Country FE	Yes	Yes	Yes	Yes
Observations	386	386	386	250
R^2	0.696	0.804	0.809	0.690
Adjusted R^2	0.681	0.794	0.798	0.657
Residual Std. Error	0.025	0.020	0.020	0.011
F Statistic	105.950***	144.858***	145.963***	19.617***

Note: $^{*}p<0.1$; $^{**}p<0.05$; $^{***}p<0.01$

imperfect substitutability of capital for labor: when the labor supply is decreasing and wages are increasing, some capital is substituted for labor, leading to higher productivity. In any case, this relationship is cause for further work based on a model of an economy with generalized elasticity of substitution.

4.7 ARMA Results

To uncover lags and other time-series dependencies, I run an AutoRegressive Moving Average (ARMA) for each country. The emprical model for an $ARMA(p,q)$ process takes the form:

$$\upsilon_{ct} = \alpha + \phi_1 \upsilon_{ct-1} + ... + \phi_p \upsilon_{ct-p} + \theta_1 \varepsilon_{ct-1} + ... + \theta_q \varepsilon_{ct-q} + \varepsilon_{ct}, \tag{3}$$

where p defines the scope of the AutoRegressive part of the model, representing the number of lags over which the effect may persist, and q represents the Moving Average part of the model, allowing the residuals to exert influence for up to q periods. In the empirical model, ϕ_{ct-p} represents the effect that the p^{th} term has on υ_{ct} and θ_{ct-q} represents the effect that the q^{th} term has on υ_{ct}.

For each country, I allow p and q to take values one to ten, and select the model that fits the data best.[9]

Table 12 shows the results of the ARMA analysis. The "ar" coefficients represent ϕ_{cp}, showing the persistence of auto-regressive effects varies across countries. In general, the effects are not auto-regressive to a statistically significant extent. Only in Colombia does any statistically significant auto-correlation persist past one lag. These results support the sufficiency of incorporating a lagged dependent variable in the above panel regressions. The "ma" coefficients represent θ_{cq}, showing the persistence of moving average effects varies across countries. Chile and Colombia have statistically significant moving average effects at one lag, but their effects dissipate after more than one year and no other countries share the same relationship.

[9] With minimum Akaike Information Criterion as the method for choosing best fit.

p	q	Country	Variable	Lag	Coef.	BSE	P Value
1	1	CA	ar	L1	0.66	0.13	0.00
1	1	CA	ma	L1	0.33	0.24	0.17
1	1	CL	ar	L1	-1.00	0.00	0.00
1	1	CL	ma	L1	-1.00	0.00	0.00
4	1	CO	ar	L1	0.03	0.01	0.00
4	1	CO	ar	L2	-0.04	0.00	0.00
4	1	CO	ar	L3	-0.11	0.01	0.00
4	1	CO	ar	L4	0.96	0.00	0.00
4	1	CO	ma	L1	-0.99	0.00	0.00
1	2	DK	ar	L1	0.95	0.07	0.00
1	2	DK	ma	L1	-0.26	0.17	0.14
1	2	DK	ma	L2	-0.30	0.16	0.06
1	1	FR	ar	L1	0.83	0.11	0.00
1	1	FR	ma	L1	0.04	0.19	0.82
1	1	IT	ar	L1	0.95	0.06	0.00
1	1	IT	ma	L1	-0.03	0.22	0.90

Figure 12: ARMA results for each country.

4.8 VAR Results

The ARMA model uncovers relationships that the dependent variable has with its own lags, but does not uncover relationships between the different predictors – all of which may have time-series dependencies. To account for this possibility, I run a Vector AutoRegresion (VAR) for each country with the same predictors as the OLS specifications. The emprical model takes the form:

$$\vec{x}_t = \vec{c}_t + \Phi_1 \vec{x}_{t-1} + \Phi_2 \vec{x}_{t-2} + \cdots + \Phi_p \vec{x}_{t-p} + \vec{\varepsilon}_t, \tag{4}$$

where $\vec{c}_t$ is an nx1 vector of constants, $\vec{x}_t$ is an nx1 vector of dependent variables, Φ_i, $i \in [1, p]$ is an nxn matrix of coefficients, $\vec{x}_{t-i}$ is an nxn matrix of lagged variables, and $\vec{\varepsilon}_t$ is an nx1 vector of error terms (and all variables are first-differenced for stationarity).

Figure 13 shows the forecasted value for ULC and our independent variables for the United States and Japan. I select the number of lags in the VAR by finding the minimum Akaike Information Criterion (AIC) for each country – which was 2 for every country model except Canada, for which it was 6,

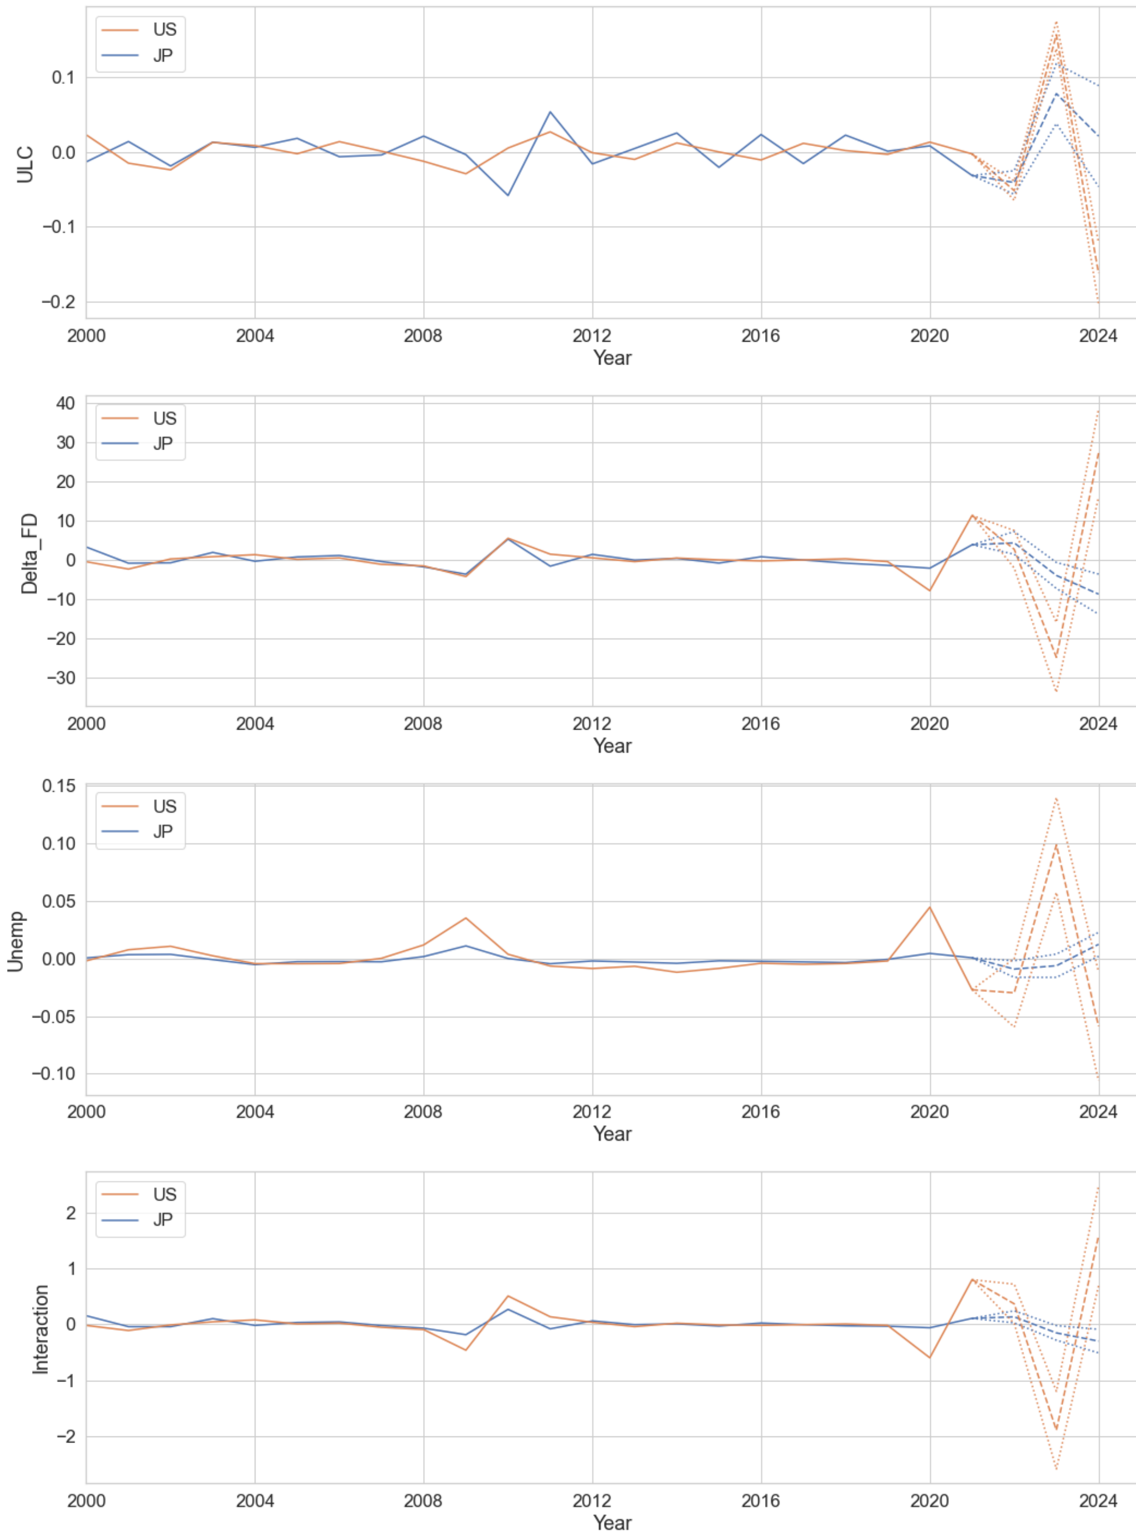

Figure 13: Three-year VAR forecasts for the US and Japan. Actuals in solid lines, projections in dashed lines, projection confidence intervals in dotted lines.

and Russia, where it was 1.[10]

These results support my hypothesis: even when permitting for variable and dependent lags, the degree to which a country adapts to demographic pressures has a marked inverse relationship to ULC, and more so when unemployment is low.

5 Conclusion

This paper has introduced a new way to think about the drivers of the cost of labor in an economy. While academic and industry analysis of the topic have historically depended on cyclical predictors, I show that structural demographic characteristics of an economy predict movements in the price of labor – especially when labor markets are tight. My results are robust to a variety of empirical approaches and persist when adjudicated alongside measures of factors already known to affect the price of labor.

This work will become increasingly important as countries across the world continue to experience aging and even shrinking populations (Arnola, 2023). When such demographic circumstances require economies to run at chronically low unemployment, countries will look outside their existing labor forces to find additional sources of hours worked. Indeed, they are already doing so. The accompanying structural shifts in the quantity of the labor supply will precipitate movements in the price of labor.

This paper helps contextualize the relationship between demographics, economic growth, and inflation. Adjustments to an economy's labor supply have implications for its economic output, but this model shows a specific channel through which they also have implications for the price of labor (and therefore inflation). Demographic characteristics of the labor supply serve as a release valve for wage pressures. Increasing hours worked can stem upwards pressure on wages, but once an economy can no longer add any more hours, it faces increases in the price of labor.

Moreover, this model connects the two main concepts affected by the price of labor: economic output

[10]Models were produced for CA, DE, DK, FR, JP, NL, UK, US, IT, RU, KR, SE, ES, and CH; the models were not defined for CL and CO.

and price level. Equity investors, in particular, might be interested in the model's implications for an economy's free cash flow (and therefore its growth potential). Demographic characteristics and tight labor markets might put upward pressure on ULC – and on the labor share of income. Companies have two choices about how to respond to rising costs on their income statement. Their first option is to defend their margins, putting upwards pressure on price inflation. Their second option is to take a hit to their margins, decreasing returns to shareholders (i.e., decreasing the capital share of income) – putting downward pressure on free cash flow.

This work shows that better forecasts of hours worked will permit more accurate predictions of future ULC. To best take advantage of the theoretical linkage this model presents, next steps could include constructing granular, high-frequency forecasts of hours worked.

References

Martin Arnola. Eurozone unemployment rate hits record low. *Financial Times*, August 2023.

Abhijit V Banerjee and Esther Duflo. Growth theory through the lens of development economics. *Handbook of economic growth*, 1:473–552, 2005.

Bank of International Setlements. Testimony by the chairman of the board of governors of the us federal reserve system, mr alan greenspan, before the committee on banking and financial services of the us house of representatives on 22 july 1999. *BIS Review*, (85), 1999. URL `\url{https://www.bis.org/review/r990726a.pdf}`.

Bank of Japan Research and Statistics Department. The state of the japanese economy: From the perspective of employment and income. *Bank of Japan Research Bulleting*, (85), January 2005. URL `\url{https://www.boj.or.jp/en/research/brp/ron_2005/data/ron0504a.pdf}`.

Aleksander Berentsen, Guido Menzio, and Randall Wright. Inflation and unemployment in the long run. *American Economic Review*, 101(1):371–398, 2011.

Darrell Bricker and John Ibbitson. *Empty planet: the shock of global population decline*. Hachette UK, 2019.

Alberto Bucci. Can a negative population growth rate sustain a positive economic growth rate in the long run? *Mathematical Social Sciences*, 122:17–28, 2023.

David Cass. Optimum growth in an aggregative model of capital accumulation. *The Review of economic studies*, 32(3):233–240, 1965.

Kristie M. Engemann. What is the phillips curve (and why has it flattened?). *Federal Reserve Bank of St. Louis Open Vault Blog*, January 2020. URL `\url{https://www.stlouisfed.org/open-vault/2020/january/what-is-phillips-curve-why-flattened}`.

Milton Friedman. Nobel lecture: inflation and unemployment. *Journal of political economy*, 85(3): 451–472, 1977.

Charles Albert Eric Goodhart and Manoj Pradhan. *The great demographic reversal: Ageing societies, waning inequality, and an inflation revival*. Springer, 2020.

Naohisa Hirakata, Yasutaka Koike, et al. The labor share, capital-labor substitution, and factor augmenting technologies. Technical report, Bank of Japan, 2018.

Charles I Jones. R & d-based models of economic growth. *Journal of political Economy*, 103(4): 759–784, 1995.

Charles I Jones. The end of economic growth? unintended consequences of a declining population. *American Economic Review*, 112(11):3489–3527, 2022.

N Gregory Mankiw. The inexorable and mysterious tradeoff between inflation and unemployment. *The Economic Journal*, 111(471):45–61, 2001.

N Gregory Mankiw, David Romer, and David N Weil. A contribution to the empirics of economic growth. *The quarterly journal of economics*, 107(2):407–437, 1992.

Karl Marx. *Capital: volume I*, volume 1. 1867.

Alban W Phillips. The relation between unemployment and the rate of change of money wage rates in the united kingdom, 1861-1957. *economica*, 25(100):283–299, 1958.

James A Robinson and Daron Acemoglu. *Why nations fail: The origins of power, prosperity and poverty*. Profile London, 2012.

Paul M Romer. Endogenous technological change. *Journal of political Economy*, 98(5, Part 2): S71–S102, 1990.

Paul M Romer. The origins of endogenous growth. *Journal of Economic Perspectives*, 8(1):3–22, 1994.

Robert Shackleton. *Estimating and projecting potential output using CBO's forecasting growth model*. Congressional Budget Office, 2018.

Robert M Solow. A contribution to the theory of economic growth. *The quarterly journal of economics*, 70(1):65–94, 1956.

James Tobin. Inflation and unemployment. *Essential readings in economics*, pages 232–254, 1995.

Robert Topel. Labor markets and economic growth. *Handbook of labor economics*, 3:2943–2984, 1999.