Okay Boomer... Excess Money Growth, Inflation and Population Aging

Joseph Kopecky*

22nd May 2022

Abstract

What determines the strength of the relationship between money growth rate and inflation? A large literature suggests that it has weakened since the 1980s, without a definitive explanation of the cause. In this paper, I explore the extent that population age structure explains changes in the pass through of money growth rates to inflation. In a long run annual panel of countries, I show that the quantity theory of money holds over long time horizons, with sizable estimates of the impact of money growth on inflation in the short-to-medium term. I then estimate state dependent local projections at five year horizons, showing that various measures of population age structure have significant impact on the strength of this relationship. These demographics account for a substantial increase in the transmission of money growth to prices in the 1970s, as well as a weakening of this relationship throughout the great moderation. I find that the baby boomer cohort, now in the age group around retirement, may exert upward pressure on this money transmission to prices at present, with ambiguous implications in the future as low fertility and rising longevity continue to play a role.

*Trinity College, Dublin (jkopecky@tcd.ie)

1. Introduction

In this paper I explore the role that population age structure plays on the pass through of money growth rates to inflation. I first establish evidence in support of the long run relationship implied by the quantity theory of money. Estimating this in a historical panel of 16 developed countries since 1870, I find strong money growth correlations with prices over short-to-medium horizons and convergence to a one-to-one relationship after twelve years. My primary contribution is to estimate the state dependent impact of population age structure on the transmission of changes in money growth to prices. I show that progression of the baby boomer cohort through the life-cycle, in addition to long run trends in old and young dependents, can explain a strengthening of the money growth-inflation relationship during the 1970s, as well as a subsequent weakening throughout the great moderation.

A recent shift in central bank policy is the diminishing importance of money. While monetary aggregates are still nominally used as one of the "two pillars" of European Central Bank analysis, their importance have substantially waned. Early critics of money growth pointed to an apparent instability in the money-inflation relationship during the turbulent period of the 1970s and early 1980s.[1] The great moderation has further minimized money's role, with low and stable inflation seeming to react weakly to changes in monetary aggregates. The empirical evidence is fairly clear: growth in traditional monetary aggregates has had a weaker correlation with inflation. Less clear is *why* this weakening has taken place. This paper presents evidence that aging may be an important channel for consideration. The aggressive actions of central banks around Covid-19 resulted in a large shock to monetary aggregates, with official measures of M2 in the United States growing by 19% in 2020 relative to the prior year, and again by 16.24% in 2021. If demographics have some role to play in the transmission of money growth to inflation, then the policy environment may be quite different today than in the tranquil period of the great moderation.

Why might population age structure matter for money and prices? Population aging has gained a great deal of attention as a potential source of future inflation. This is one of the primary theses of Goodhart and Pradhan (2020), and is supported by a large and growing literature.[2] Separately, many have identified demographics as potentially important for secular trends in equilibrium interest rates and asset prices.[3] These papers show that life-cycle demand for assets drive secular trends in interest rates as households accumulate savings during working life and spend them down in retirement, with relative asset prices

[1]For a thorough treatment of this policy discussion at the time see Goodhart (1989).

[2]See for example: Bobeica, Nickel, Lis, and Sun (2017), de Albuquerque, Caiado, and Pereira (2020), Aksoy, Basso, Smith, and Grasl (2019), Edo and Melitz (2019), and Juselius and Takáts (2021)

[3]See fore example: Gagnon, Johannsen, and Lopez-Salido (2021), Eggertsson, Mehrotra, and Robbins (2019), Kopecky and Taylor (2022)

Figure 1: *U.S. Survey of Consumer Finances: Money Holdings by Age*

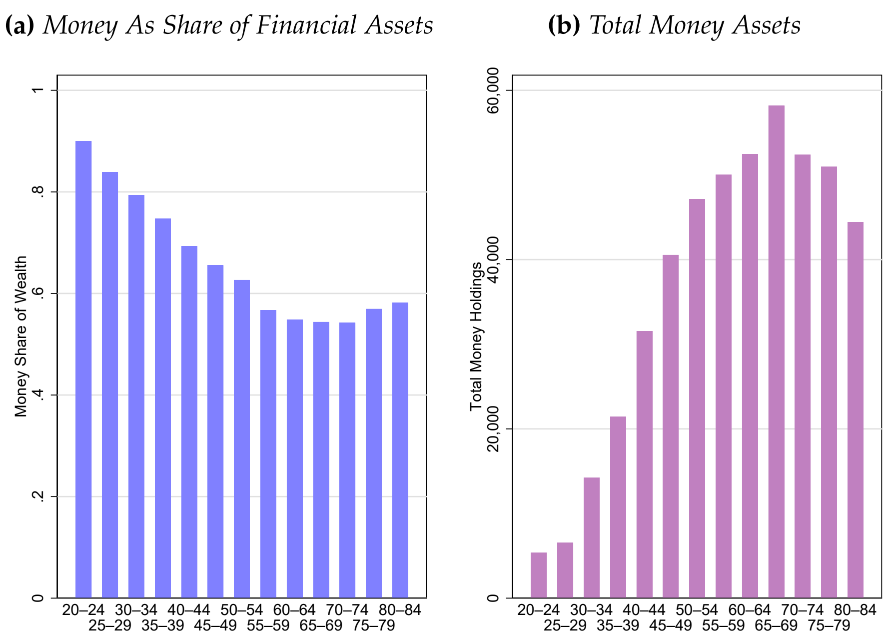

Notes: Survey of Consumer Finance Data pooled from 2001, 2004, and 2007, 2010, 2013, and 2016 surveys. Money here is defined as sum of checking, savings, and money market mutual fund accounts. Financial assets defined as sum of all non-housing financial assets excluding trusts and , life insurance. Averages are calculated by five year age group for all respondents with less than $2,000,000 in total assets.

affected by evolution of portfolio allocation across the life-cycle. Aoki, Michaelides, and Nikolov (2019) estimate a quantitative model of this type. Unlike other work, they include money as one of three assets held by households. Their model shows that money holdings (as a share of financial wealth) are generally high early in working career, before falling dramatically throughout mid-to-late career, and rebounding in retirement. This is broadly consistent with Survey of Consumer Finance data on money holdings, which I show in Panel A of Figure 1.

Aoki et al. (2019) are interested in money holdings as a share of assets, as their work focuses on life-cycle portfolio decisions. More relevant for aggregate money demand is the life-cycle nature of total money holdings, which I plot in Panel B of Figure 1. This suggests that one might expect economies with relatively more households in the years surrounding retirement to have significant upward pressure on money demand. Such a relationship is explored in Wang and Zhu (2021), who show that the share of old age (> 65) populations have a negative impact on the velocity of money. They interpret this result as population aging leading to increases in money demand, and suggest that aging may be a channel for structural changes in money demand throughout the great moderation. I seek to contribute by providing an empirical estimate of how much of the change in money growth's pass through to inflation may be accounted for by population change, while also providing a deeper empirical understanding of this relationship across the entire age distribution.

My average estimates of money growth on inflation (those independent of the age distribution) lend support to existing work, which shows a robust medium-to-long term correlation. This is in line with a recent literature[4] supporting the long run Friedman and Schwartz (1963) relationship. Meanwhile, my heterogeneous estimates (those conditional on age), provide an explanation for weakening of this relationship suggested in work such as: Teles, Uhlig, and Valle e Azevedo (2016), Benati (2005), and Gertler and Hofmann (2018). This literature finds a significant change in the quantity theory relationship since the 1980s. My conclusions are not mutually exclusive of other structural explanations. It is entirely possible that in addition to aging, changing central bank policies have caused structural breaks in money-inflation relationships by grounding expectations.[5] Here I suggest that aging may be an additional channel and should be taken into account when trying to understand the low frequency relationship between money and inflation. Establishing a strong inflation anchor[6] likely plays a role in the decline of money growth's importance, but demography may provide currents to pull on that anchor. Recent criticism of the importance of inflation expectations in Rudd (2022) has gained substantial attention. While in my view this critique undervalues evidence in favor of expectations, I agree that this is an area where causal evidence is scarce, and that we should be careful when enshrining conventional wisdom as unassailable truth. This paper provides a different, and potentially complementary perspective.

Methodologically, I use a tool developed in Cloyne, Jordà, and Taylor (2020), who study the effect of monetary accommodation of fiscal multipliers. They suggest a parsimonious way of dealing with decomposition of heterogeneous policy estimates in the local projections estimations of Jordà (2005) by implementing a Kitagawa (1955)-Blinder (1973)-Oaxaca (1973) style decomposition.

The paper proceeds as follows. In Section 2, I discuss the data and my methodological approach. Section 3 establishes a baseline relationship between money growth and inflation over this sample period, showing substantial short and medium run correlations with long run estimates approaching unity. Section 4 performs the Kitagawa-Blinder-Oaxaca estima-

[4]See Assenmacher-Wesche and Gerlach (2007) and Fratianni, Gallegati, and Giri (2021). As well as Belongia and Ireland (2016) who show that much (though not all) of the weakening relationship is due to mismeasurement and fixed using a Barnett (1980) Divisia measure of money.

[5]For two papers who discuss regime shifts to "anti-inflationary" policy as changing this relationship see: Sargent and Surico (2011) and Benati (2005)

[6]See Reis (2021) for a very recent discussion of historical anchoring as well as a present view on movements in expectations.

tion and decomposition proposed by Cloyne et al. (2020), and demonstrates quantitatively that aging may have a large impact on the transmission of money growth to prices. I show that these results remain when studying a smaller sample with more appropriate Divisa aggregations of money in Section 5. Section 6 concludes.

2. Data and Empirical Strategy

2.1. Data

Population data come primarily from the Human Mortality Database (2019), (HMD). I use their data on population by age to construct various population shares across the age distribution. In addition, I use US Census National Intercensal Tables to extend the United States demographic series to 1900. I consider two methods for controlling for the population age structure. The first is to use the working age population, defined as the number of individuals aged 20-64 as a percentage of the overall population. The second relies on demographic variables that cover population shares over the full age distribution as in Fair and Dominguez (1991).

In Figure 2a, I plot the working-age population for the full panel of countries used in my estimations. This highlights an advantage of such a long sample. While a great deal of work is limited to either a single country, or to panels that include only the period since the baby boomer generation entered the workforce, I can take advantage of a large amount of variation both across countries, as well as in the period before the boomer cohort. While all of the countries in this sample have some sort of baby boom, the severity and timing are different.

I merge this population data with the macrohistory dataset of Jordà, Schularick, and Taylor (2017) (JST) which provides information on price levels and a wide range of macroeconomic variables of interest. With the addition of Jordà, Knoll, Kuvshinov, Schularick, and Taylor (2019) this also contains useful information on asset returns and credit market conditions. It will be crucial to control for a broad set of macroeconomic and financial controls from these data, which have been shown to be important. While the dataset contains 18 developed countries, my preferred specification drops Ireland and Canada as they lack some of the financial variables used as controls. For the remaining 16 countries[7] data go back as far as 1870 with information on population age, macroeconomic controls, and financial variables. The shortest panel is 32 years (Belgium) due to lack of data on broad money, but all other series are more than 50 years, with the average in the panel

[7] For ease of comparability I present all results only on this truncated sample. Including them has no impact on overall results for specifications where data is available.

Figure 2: *Working Age Populations and Excess Money Growth*

(a) *Working Age Population* **(b)** *Excess Money Growth*

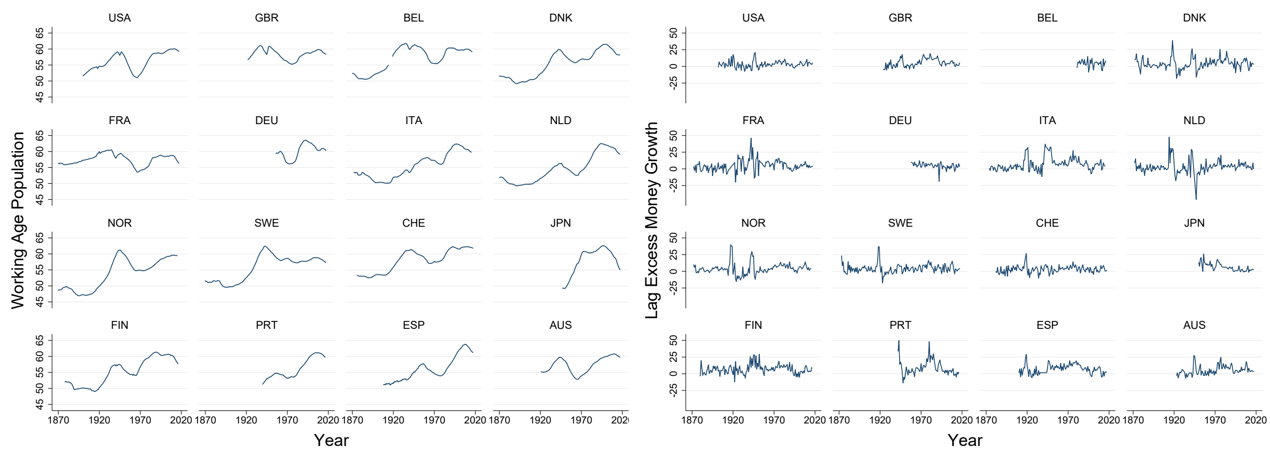

spanning 85 years of data.

A simple quantity theory of money states the growth rate of inflation should be the growth rate of money less growth in output and velocity. I take as my measure of *excess money growth* the difference between broad money growth rates and growth of output. Assuming a constant velocity the quantity theory then suggests one-to-one relationship. Figure 2b shows my measure of excess money growth. It is important to note that part of the weakening of money is due to mismeasurement issues around official aggregates. Belongia and Ireland (2016) show that by using the Divisia monetary aggregates of Barnett (1980) that there has been a lengthening of the lag time for money to affect output and prices since the early 1980s, but that the relationship remains strong. This is in line with Barnett, Offenbacher, and Spindt (1984), and many others who show that these Divisia aggregates generally outperform official measures. While these measures are not available for the long run panel I study, it seems plausible that they might reflect a faster convergence to quantity theory convergence than I find using official measures of broad money from the JST data. I will show for a subsample of my data (the United States from 1967) that my age dependent estimates are similar when using these more appropriate measures money, though consistent with this literature such money growth measures will have a much stronger estimated impact on inflation.

2.2. Estimation Strategy

I wish to study the impact of excess money growth rates on cumulative inflation conditional on the state of demographics. As such I estimate cumulative impulse response functions

(IRFs) of the change in price level to the excess money growth rate. I use the reduced form local projections method of Jordà (2005), which has been shown in Plagborg-Møller and Wolf (2021) to estimate the same IRFs as an equivalently specified VAR, with flexibility regarding a wide range of structural assumptions. In my case, the local projection is one with a structural recursive identification with money growth in prior period affecting contemporaneous values of my outcome and controls. This standard impulse response, not taking into account state dependence, is shown in Equation 1:

$$\pi_{i,t+h} - \pi_{i,t-1} = \mu_i^h + m_{i,t-1}\beta^h + (x_{i,t} - \bar{x})\gamma^h + D_{i,t}\gamma_d^h + \epsilon_{i,t+h}\ ; \quad \textit{for } h = 0, \dots, H \tag{1}$$

The outcome, $\pi_{i,t+h} - \pi_{i,t-1}$, is the cumulative growth rate of price levels from the current period to the h annual horizon. The key independent variable, $m_{i,t-1}$ is excess money growth, defined as growth in M2 money less growth in output. Estimates of horizon $\beta^{h=0}$ are therefore one-year relationship between changes in excess money growth rate the year prior and inflation in the current year. Estimates at this horizon might be considered the short run impacts of these policies, while later horizons reflect a medium-to-long term relationship. These are correlations with money and inflation, similar to those estimated in Belongia and Ireland (2016), who study the inflation and output response to growth of Divisia money aggregates using a recursive VAR identification. This is far from a causally identified shock, and a narrative instrument would be preferred and easily incorporated into this framework. However, such an instrument is not readily available for the panel of countries and time period studied. Although these results suggest a strong empirical relationship, I caution against any causal interpretation.

I include in the vector $x_{i,t}$ for a broad set of macroeconomic and financial controls. In order to have comparability with estimates using the Kitagawa-Blinder-Oaxaca (KBO) decomposition described below, all controls are taken as deviations from their sample mean. In the baseline specification, I populate $x_{i,t}$ with: growth rate of real output, growth of real consumption per-capita, the change in the ratio of investment-to-gdp, change in the current account-to-gdp ratio, change in government debt-to-gdp ratio, real equity prices, growth in total loans to the household sector, population growth rates, and the change in the short term bill rate. Two lags, along with contemporaneous values are included in the regression specification for each control. Additionally, three lags of current inflation are included. I include dummy variables for both world wars, which are interacted with demographics to avoid clouding estimates with large swings in some age groups related to wartime periods. I also include a measure of global GDP growth, here taken as the aggregate GDP growth of the countries in my sample, as a means of controlling for global time varying trends. This captures much of the variation of time fixed effects in the same way as similar long-run studies[8] on this data, without the need to add an additional 143 controls as time fixed effects.

Inclusion of time fixed effects give qualitatively identical results for estimated demographic changes to policy transmission. They have larger error bands, and a smaller, though still sizable, response of inflation to money growth over reported horizons. Time fixed effects can bias estimates in the presence of unobserved time trends that are not uniform across countries. Such trends are easy to imagine in this context. For example structural changes to monetary policy framework likely occur across this sample, with differential implementation and timing. I opt to consider my baseline estimates without such effects. They are included in Appendix D for reference. I include country fixed effects with μ_i^h. The term $D_{i,t}$ represents the set of controls for population age structure. This is either the working age population (WAP), or constructed as in Fair and Dominguez (1991) as polynomial controls for the entire population age distribution. I describe construction of the latter below.

2.3. Kitagawa-Blinder-Oxaca Heterogeneous Estimates

In Cloyne et al. (2020), the standard local projection framework is extended to incorporate an estimation strategy well known in the microeconomic literature, and first widely used in work by Kitagawa (1955), Blinder (1973), and Oaxaca (1973). They propose this method as a way of parsimoniously dealing with heterogeneity of treatment impacts of policy implementation. Introducing all control variables as demeaned relative to the sample population, and including them in the local projection directly (as above) and also as interactions with the policy variable allows for decomposition of estimates into: average estimates for the policy on the outcome, the average estimates for controls on the outcome, and the change in the policy estimates when controls move away from their sample means.

I note, as they emphasize in their paper, that without a clear identification of independent variation along various dimensions of heterogeneity it not possible give a causal interpretation of this inherently partial equilibrium approach. The estimates coming out of the interaction between a given control and the policy shock requires identification of not only the policy, but also the control. However, such concerns are widely present in the large existing literature on state dependent policy, which typically split samples across discrete states for estimation without clearly dealing with the additional identification issues arising from potential selection across states. Further, the identification challenge outlined for the policy shock above remains, and while estimates are comparable to similar recursively

[8]See Jordà, Schularick, and Taylor (2020), for example.

identified VARs, the KBO approach used here does not provide any further *identification*, but rather a *decomposition* of the empirical relationship between the policy and inflation across all states of the control set.

The KBO estimates of IRFs by local projections are shown in Equation 2. I split out my vector of demographic variables $D_{i,t}$ from other controls $x_{i,t}$ for ease of reference, but they are treated identically.

$$\begin{aligned}\pi_{i,t+h} - \pi_{i,t-1} = \mu_i^h + m_{i,t-1}\beta^h + (x_{i,t} - \bar{x})\gamma^h + m_{i,t-1}(x_{i,t} - \bar{x})\theta^h \\ + D_{i,t}\gamma_D^h + m_{i,t-1}D_{i,t}\theta_D^h + \epsilon_{i,t+h}\ ; \quad \textit{for } h = 0, \ldots, H\end{aligned} \tag{2}$$

Cloyne et al. (2020) show that under fairly standard assumptions the local projections estimate from Equation 2 can be estimated by any usual method. Here I use OLS, as is conventional for local projections. All residuals from the estimations below have covariance stationary error processes. While all additional controls used are stationary I(0) processes, the demographic variables (and cumulative inflation itself) are I(1). I discuss in subsection 2.5 that population age variables are cointegrated with inflation, consistent with prior evidence in the literature. This is important for my results as it implies that my estimates of γ^D and θ^D should be consistent, and differencing demographics to create a stationary I(0) process would be inappropriate in this context.

Equation 2 decomposes estimates into three types of results. A *direct* relationship between the policy variable and inflation, an *indirect* relationship between policy (conditional on the state of underlying controls), and *composition* effects of the changing control set on the outcome itself. Estimates from changing composition of the control set on outcomes directly are given by γ^h. For instance, older populations may have an impact on the level of inflation directly, as has been shown in recent work by Juselius and Takáts (2021) and is discussed broadly in Goodhart and Pradhan (2020). The *direct* relationship of the policy variable is given by β^h. The KBO approach is such that this is the average effect in sample, and is comparable to an equivalently specified estimate without state dependence. The *indirect* relationships are estimated via θ^h and are by construction zero when the policy variable is zero or when controls are at their mean $\bar{x}_{i,t}$. A change in any control from its mean in the population has a γ^h estimated impact on inflation itself, while also affecting the size of the estimated policy transmission through θ^D.

The advantage of this specification is that it provides straightforward ability to quantify heterogeneity of policy, as well as in the underlying assumptions embedded in existing state dependent literature. Much of the literature estimating such effects will opt to estimate coefficients on split samples.[9] While it might be possible to find a somewhat compelling instrument for monetary policy interest rate shocks in some contexts, such as the well known Romer and Romer (2004) narrative series, there is not, to my knowledge, a readily available instrument for exogenous change in monetary aggregates suitable for this larger panel. Further, identification of these shocks would still require causal variation in demographics as mentioned above, which is likely a more challenging task. For this reason, although I will work hard to show that these results are fairly robust, I reiterate my caution against a causal interpretation.

Given the importance of the question at hand I believe that exploring this empirical relationship is worthwhile. Much of the conventional wisdom around the structural shift in money's effect on prices also generally lacks clean empirical causal identification. Because of this the recent empirical advancements in the literature have focused more on identification of interest rate policy shocks. Ignoring money growth may seem reasonable if the underlying structural shift is permanent, as one might expect in a story of policy regime change. I show that demographic channels provide another plausible, and not mutually exclusive, explanation of the decline in importance of monetary aggregates during the great moderation, with different implications going forward as populations continue to evolve.

2.4. Controlling of Population Structure: an Augmented Fair and Dominguez (1991) Methodology

Before moving to results I must first describe how I will control for demographics in Equation 2 using the entire population age structure. One of my two main demographic specifications will be to use a methodology similar to that of Fair and Dominguez (1991), who fit the estimated effects for population age shares across the *entire* age distribution with a low order polynomial. A naive regression might split the population into J bins and add all of these population age shares to the estimating equation. This is problematic for two reasons. The first problem is that all of the population shares sum to one, making it impossible to jointly estimate all of the coefficients along with a regression constant. Additionally, as the number of age bins increases, the population shares become increasingly colinear. This latter point often results in significant parameter instability as one splits the age structure into finer and finer groups. The work of Fair and Dominguez (1991) suggests two assumptions:

1. Letting α_j be the regression coefficient on population share $p_{j,i,t}$ of age group j in

[9] See for example many papers on state dependent fiscal multipliers: Ilzetzki, Mendoza, and Végh (2013), Gechert and Rannenberg (2014), Jordà and Taylor (2016) and others.

country i and time t. Assume that all of these coefficients across the age distribution sum to zero. In other words:

$$\sum_j^J \alpha_j = 0 \tag{3}$$

2. Assume that the estimated age coefficients α_j can be fitted with a K order polynomial. In other words:

$$\alpha_j = \sum_k^K \gamma_k j^k \tag{4}$$

The first point makes the coefficients on J perfectly colinear population shares estimable without dropping the constant from the equation. The second forces that the estimated coefficients across the age distribution to be relatively smooth, and reduces the number of parameters that need to be estimated. How much variation is allowed across the life-cycle depends on the order of the polynomial, as increasing the degree of the polynomial allows for more turning points in the estimated coefficients across age. A notable, but not substantial, difference between my approach is that instead of using population age shares to construct these variables. I use these shares less their population mean, consistent with the KBO approach. By putting these deviations of population age shares from long run means into Equation 1 and applying the two assumptions above I construct demographic variables that estimate the γ_k parameters of the fitted polynomial. These are given by:

$$D_{k,i,t} = \left[\sum_j^J (p_{j,i,t} - \bar{p}_{i,t}) j^k\right] \tag{5}$$

I use a fourth order polynomial in my baseline specification, and so create $D_1 - D_4$ to include in my analysis. I provide evidence that this choice is the best fit in Appendix C. For a detailed derivation of Equation 5, as well as a discussion of the Fair and Dominguez (1991) controls see Appendix A.

2.5. Use of Non-Stationary Demographic Controls

Most of the controls used above are growth rates or first differenced variables, and all non-demographics controls pass an augmented Dickey-Fuller test for stationarity. These variables are generally included with the intention of accounting for the shorter run factors that may be important for determining inflation. However, my demographic controls described above are non-stationary, $I(1)$ series. There may be some concern that results below represent a spurious relationship with cumulative inflation, which though stationary in changes at very short horizons is itself $I(1)$. Montiel Olea and Plagborg-Møller (2021) show that lag-augmented local projections provide valid inference over long horizons with both stationary and non-stationary data. Thus as a first protection against spurious relationships, all estimations will include three lags of inflation as controls. The number of lags used here, or indeed specifications with no lagged inflation, have little impact on my coefficients of interest.

Additionally, I find that the working age population as well as the set four of demographic controls described in Section 2.4 are co-integrated with cumulative inflation, which my outcome of interest. I show these test results in Appendix B. This is consistent with existing work (Bobeica et al., 2017) who study the cointegrated relationship between inflation and age structure. In the presence of a cointegrating relationship the estimations performed in Equation 2 should be *super*-consistent, with coefficients of interest converging faster than they would if the series were stationary. This suggests that these demeaned levels are more appropriate than a stationary $I(0)$ difference of the demographic controls. As a final check, for all estimated models in the paper I run an augmented Dickey-Fuller test on predicted residuals post-estimation with a unit root rejected in all cases with extremely strong (< 0.0001) significance.

3. Money Growth and Inflation: The Long Run Evidence

To set a baseline result, I estimate the direct relationship between excess money growth and inflation through Equation 1. This serves as an average baseline, without heterogeneity across controls, for the empirical relationship between money growth and inflation. In Table 1, each column shows the cumulative impulse response function five years after a change in excess money growth. This window gives an idea of the *medium-to-long term* correlation with inflation.

The specifications in Table 1 show how this relationship changes from a naive univariate panel fixed effects specification after the inclusion of macroeconomic,[10] financial,[11] and my two sets of demographic controls. Standard errors are clustered at the country level. While inclusion of both sets of controls lower point estimates, the coefficient on money growth remains large in the full specification. Roffia and Zaghini (2007) show that accommodation of excess money growth by financial variables is important in the short run. Given the importance of longer run credit trends (as described in Schularick and Taylor (2012)

[10]Current and two lags of the demeaned values of: GDP growth, per capita consumption growth, change in investment as a share of GDP, population growth rate, change in current account to GDP ratio, log change in debt to GDP ratio, dummies for WW1 and WW2, as well as interactions between WW1/WW2 and demographic vars.

[11]Current and two lags of demeaned values of: change in log equity prices, dummy for financial crises, growth in credit per GDP, and the change in the short term bill rate.

Table 1: *Inflation Response to Excess Money Growth: Five Year Cumulative Effect*

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
m_{t-1}	1.339***	0.782**	0.739**	0.634**	0.783**	0.624***	0.632**
	(0.386)	(0.328)	(0.305)	(0.223)	(0.321)	(0.203)	(0.222)
dmWAP			-1.055*			-0.221	
			(0.542)			(0.357)	
D4				-0.005***			-0.005***
				(0.001)			(0.001)
D3				0.154***			0.145***
				(0.033)			(0.031)
D2				-1.451***			-1.390***
				(0.347)			(0.318)
D1				4.782***			4.670***
				(1.352)			(1.210)
Macro $X_{i,t}$	No	Yes	Yes	Yes	Yes	Yes	Yes
Financial $X_{i,t}$	No	No	No	No	Yes	Yes	Yes
FStat: D1-D4				9.27			10.36
AIC	12357	11956	11930	11348	11895	11526	11285
BIC	12362	12034	12008	11431	11973	11604	11368
R2	0.14	0.38	0.39	0.60	0.40	0.55	0.62
N	1342	1342	1342	1342	1342	1342	1342

Table reports estimations of IRFs for money growth under working age population, and fourth order polynomial specifications. Standard errors clustered at the country level are in parenthesis with $*p < 0.10, **p < 0.05, ***p < 0.01$. $X_{i,t}$ significance. Regressions include three lags of the dependent variable as well two lags and contemporaneous values of all additional included controls.

and others) these are likely also important for the long run transmission. My preferred specification throughout the rest of this paper will include all of these controls. To see how the coefficient on money growth evolves over time, I plot the impulse response function for the specification in Column 7 over a fifteen year window in Figure 3.

There are two things worth noting about Figure 3. The first is that including the full set of macro, financial, and demographic controls, there appears to be a convergence to the unity relationship implied by the quantity theory in the historical data over about twelve years. Documenting this fact in a long sample is one contribution of the paper. The second takeaway is that the shot-to-medium estimates for transmission money growth to inflation are substantial. The current year estimated impact of a lagged percentage point change in money growth is small, but economically significant implying that each percentage point of money growth will contribute to about 0.11 percentage points of inflation in the next year. This rises persistently and is about 0.63 within five years. While it takes quite a long time to reach the predictions of long run theory, there are non-trivial correlations between excess money growth rates and inflation for all horizons, and across all specifications listed in Table 1.

The first demographic control, dmWAP, is the deviation of the working age population

Figure 3: *Inflation Response Unit Change in Lagged Excess Money Growth: Full Set of Controls*

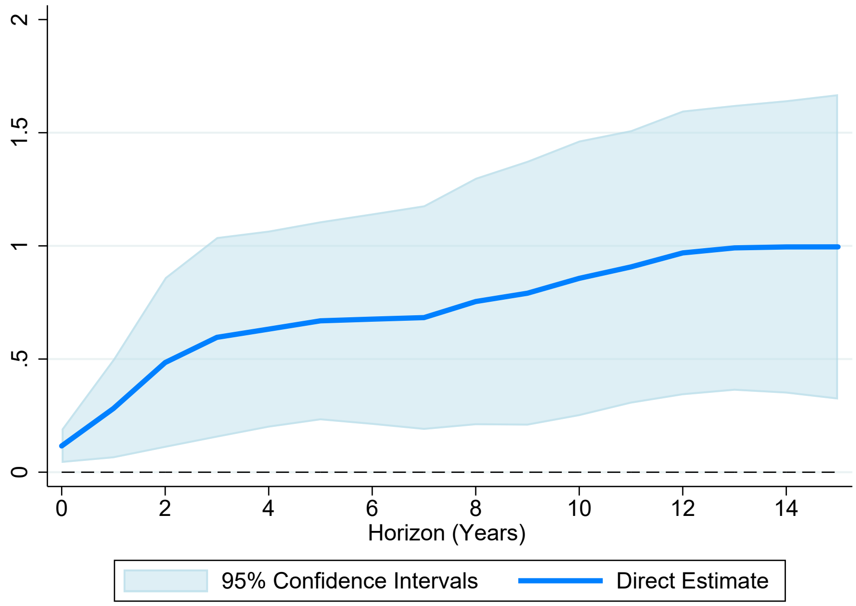

from its long run mean. The working age population has the expected sign. More workers loosen labor markets and reduce wage/price pressure, while consumption associated with dependents may put upward pressure on prices. However, in my specification with financial controls the statistical significance of this estimate disappears. The more flexible estimation of the age structure is highly significant, lending some suggestive evidence that there is more variation in age structure that is missed by a single statistic. Importantly, variation across working age might make age specific relationships hard to capture with a single variable like the WAP alone.

While my Fair and Dominguez (1991) polynomial controls are each strongly significant, with very strong joint significance, one downside of using these is that their regression coefficients have little intuitive meaning themselves. For example, the estimated coefficient for D1 represents the sensitivity of the linear (k=1) term in Equation 4. For a better understanding of age specific relationships, I calculate the effect of each of the J the individual demeaned population shares, which can be backed out by using the regression estimates in Equation 4.

$$\hat{\alpha}_j = \hat{\gamma}_0 + \hat{\gamma}_1 j + \hat{\gamma}_2 j^2 + \hat{\gamma}_3 j^3 + \hat{\gamma}_4 j^4$$

Where the $\hat{\gamma}$ terms above are the coefficients estimated in Table 1 for each demographic control, with $\hat{\gamma}_0$ defined as a function of these due to Equation 3. In Figure 4, I plot these point estates as well as confidence bands calculated via the delta method. These results are

Figure 4: *Inflation Response to Population Share: $h = 4$*

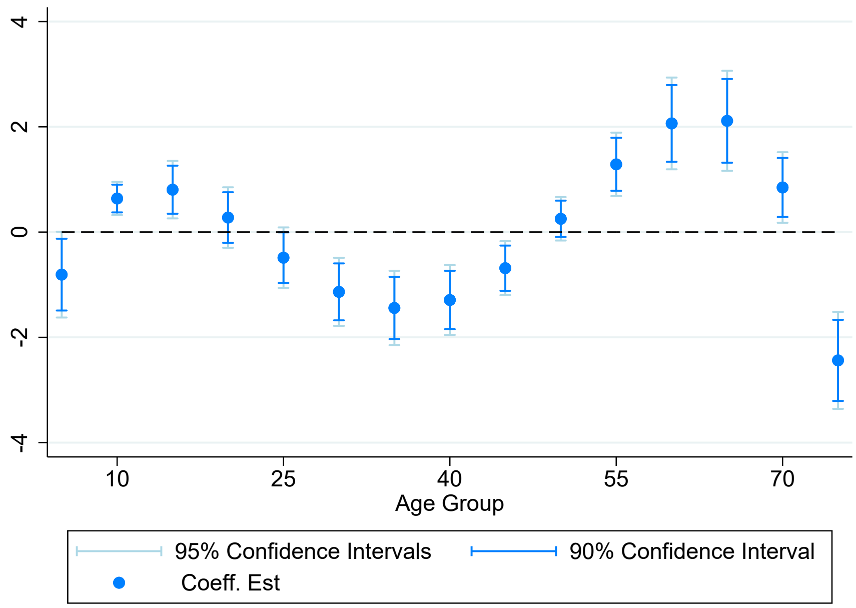

fairly closely aligned with what might be expected by theory. Young and old age groups have generally positive correlation with inflation, while early-mid career workers have significantly negative coefficients. The negative values for the very youngest and the very oldest groups are slightly puzzling, but for the oldest age group is consistent with prior findings of Juselius and Takáts (2021), who's work use a similar Fair and Dominguez (1991) control in a static framework,[12] and find quite similar picture across the age spectrum.

It is important to note that these results tell us nothing about the effect of aging on transmission of money growth. While an age composition relationship is undoubtedly important for understanding inflationary pressures, it is not the primary focus of this paper. Prior work has suggested that low inflation regimes may have different transmission of policy, by providing disinflationary pressure. This direct relationship may be important for supporting low inflation monetary regimes during the great moderation, but does not provide evidence of whether transmission of that policy has shifted. It is also useful here to see that aged 55-65 individuals appear to have opposite sign effects than 20-45 year olds, suggesting that population age structure effects might not be well approximated by the WAP, which is commonly used as a proxy for age structure. Having established this baseline, I move to estimating the *indirect* relationships from Equation 2 with interactive terms for demographics and all other controls.

[12]Figure 4 is essentially a replication of their headline results to a dynamic five year projection.

4. Population Age Structures and Monetary Transmission

I now present evidence that population aging is strongly related to the transmission of money growth to inflation. I estimate the impulse response functions defined in Equation 2, as suggested by Cloyne et al. (2020). My baseline specifications will be the full set of macroeconomic and financial controls, but now also control for interaction of the lag of excess money growth and each of these covariates. My primary outcome of interest is not the direct effect of aging on inflation (though these can still be seen), but rather this interactive term.

4.1. Working Age Population

I first present results for working age population (WAP). While this is a crude estimate of the age structure, and evidence from Figure 4 might suggest that it misses important turning points, it allows for quite simple interpretation and therefore provides useful introduction to population estimates and their relationship with monetary transmission. In Table 2, I show the impulse response function at up to five year horizons for these estimates. The direct coefficient of lagged excess money growth accumulated over this whole horizon (0.51) is of similar magnitude, though smaller than the equivalent estimate in Table 1 (0.62), after adding interactions of money growth with WAP, as well as with all other covariates.

Table 2: *Response of Inflation to Excess Money Growth: Working Age Population*

	0	1	2	3	4
m_{t-1}	0.09***	0.22**	0.40**	0.49**	0.51**
	(0.03)	(0.08)	(0.14)	(0.17)	(0.18)
dmWAP	0.08**	0.31**	0.48***	0.57**	0.55*
	(0.03)	(0.11)	(0.16)	(0.22)	(0.29)
$m_{t-1}\times$ dmWAP	-0.02***	-0.08***	-0.11***	-0.14***	-0.15***
	(0.01)	(0.02)	(0.03)	(0.04)	(0.05)
Macro $X_{i,t}$	✓	✓	✓	✓	✓
Financial $X_{i,t}$	✓	✓	✓	✓	✓
Interact $X_{i,t}$	✓	✓	✓	✓	✓
AIC	6893	9134	10028	10697	11229
R2	0.70	0.65	0.65	0.64	0.64
N	1342	1342	1342	1342	1342

Table reports estimations of IRFs for money growth under working age population, and fourth order polynomial specifications. Heteroscedastic standard errors are in parenthesis with $*p < 0.10, **p < 0.05, ***p < 0.01$. $X_{i,t}$ significance. Regressions include three lags of the dependent variable as well two lags and contemporaneous values of all additional controls and their interactions with lagged money growth.

The direct coefficients on the WAP now tell a different story. The impact of age *composition* is now positive and quite large. This seems contradictory, as theories connecting WAP to inflation through labor market tightness, such as those put forward by Goodhart and Pradhan (2020), suggest a negative effect. A potential explanation could be that hump shaped life-cycle consumption, peaking late in working life might explain some inflationary relationship of the WAP. It also seems plausible that some of the estimated effects of age that many have found could work through indirect channels such as those I find here. As above, any implications for inflation that vary throughout working life, as in Figure 4, could make this relationship difficult to cleanly estimate as it cannot distinguish between a large WAP with many 20-40 year-olds and one with high concentration in late career. In Appendix C Table 7, I also show that while the signs of these coefficients are fairly consistent, their significance depends on inclusion of controls (especially for the direct relationship), suggesting that one should be wary of explanations of aging and inflation that use WAP alone. In the following section, where I use the entire age distribution, my results for the composition effect of age structure on inflation directly will be more or less in line with the inflationary aging hypothesis in the literature while also being strongly

Figure 5: *Heterogeneous Effect of Excess Money Growth and Inflation: WAP*

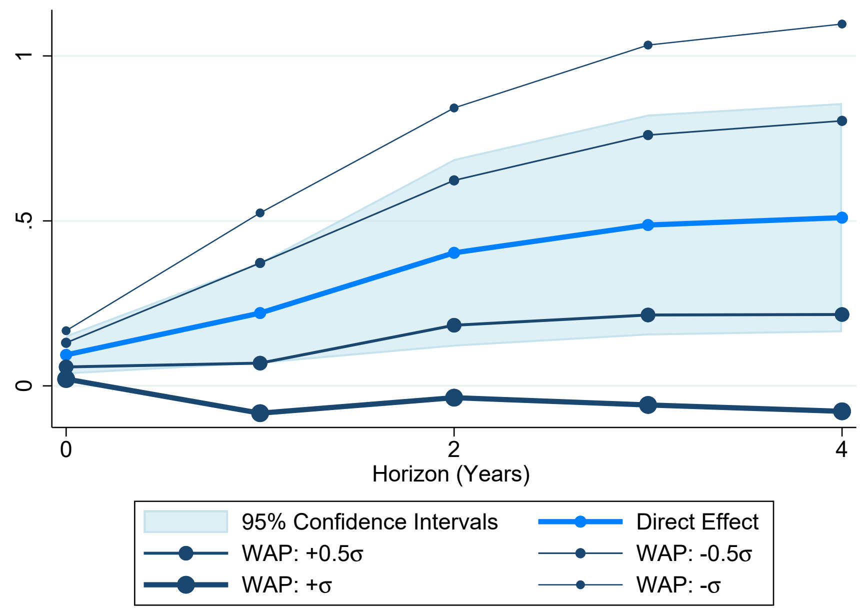

robust to choice of control set.

How large are these transmission mechanisms? Figure 5 shows the effect of increasing/decreasing the demeaned WAP by one-half and one standard deviations. The thickest line/markers represents an increase of working age population to 3.93 percentage points above its long run mean, while the smallest line/markers shows a symmetric decrease. The point estimates suggest that a high working age population (thicker curves/markers) brings transmission of money growth down substantially. Taking the United States as an example, the WAP has been persistently above its long run mean starting from 1977, and remains so today (though notably not persistently increasing over that period). In the latest year of the sample, 2017, the working age population in the united states is about 3.06 percentage points larger than the historical average of 56.7% in the sample, but has begun falling somewhat rapidly off a peak that occurred in 2011.

These are better understood in context by plotting similar variation to Figure 5, but instead feeding empirical values for a given country rather than standard deviations. In Figure 6, I present this age decomposition for three countries: the United States, France, and the United Kingdom. Each left-hand panel contains impulse responses of inflation to money growth over a five year window for one of these countries. For reference the *direct* money growth relationship, which is an average estimate across the sample, is plotted, along with conditional estimates for the age distribution of that country in three historical periods: 1970, 1990, and 2017; and two future projections: 2030 and 2050. Data on future demographics comes from UN (2019) projections. I also plot the evolution of the demeaned working age population in each country in the right-hand panels.

In Figure 6b, I show the demeaned value of the WAP for the United States to provide a clear idea of what drives the conditional results in Figure 6a. In particular, the effect of the baby boomer cohort, which represented an unprecedented (in this sample) increase of young dependents in the 1950s and 60s, pushed WAP down in the post-war period with a dramatic reversal over the course of the great moderation. In the decades to come, a return toward the mean WAP will imply gradual return to average estimated money pass through to inflation. One thing worth noting is that a persistent decrease in young dependents in the years following the boomer cohort both exaggerated the rise of WAP during the great moderation, and slows the speed of decline as boomers enter retirement. Similar historical evidence can be seen in France and the UK, though in France working age population has already reached the mean value in the sample, and both countries will see declines much more rapidly than the United States. Increasing longevity will provide negative pressure on the working age population for all of these countries going forward, further increasing estimated money pass through to inflation implied by these estimates.

These results give suggestive evidence that demographic channels increased transmission of money growth to prices in the 1970s, weakened it over subsequent years, and may increase it again as boomers fully leave the WAP. However, the WAP is an imprecise measurement of aging. Using just one statistic for population age structure assumes a number of things. First, it forces any transmission effects to be constant across the life-cycle. We know from work such as Wong (2019) that this is unlikely to be the case with interest rate policy, and Figure 1 suggests it is unlikely the case for life-cycle variation in money demand. Additionally, using WAP forces old and young dependents to have symmetric effects. This seems unlikely, especially in the face of lengthening of time-span in retirement, which itself might alter the consumption behavior of older retirees.

4.2. Estimates Over the Full Age Distribution

I now present estimates that control for the entire age distribution. The goal is to allow for a more flexible approach than is allowed by using a single variable to capture population age structure dynamics. I use a fourth order ($k = 4$) polynomial to fit 16 population age shares in my baseline model. In Appendix C, I show that this selection likely provides the best trade-off between improved in sample fit (based on AIC) as well as issues associated with overfitting age coefficients which is common among Fair and Dominguez (1991) controls. The estimates of impulse response functions controlling for this specification are in Table 3. As before, I present the estimation including all covariates, country fixed effects, and

Figure 6: *Response of Inflation to Excess Money Growth, KBO Decomposition: Working Age Population*

(a) *Conditional IRF: USA*

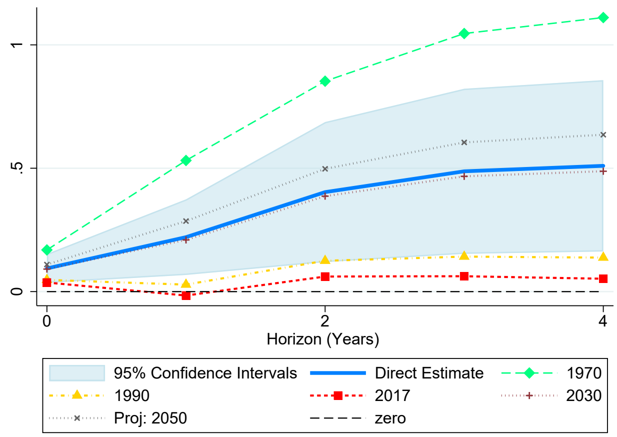

(b) *Demeaned WAP: USA*

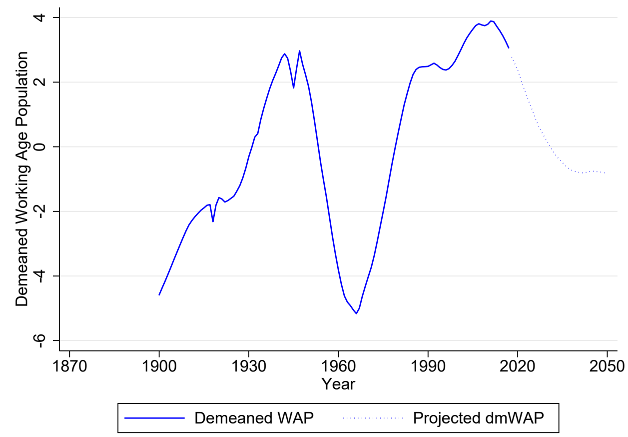

(c) *Conditional IRF: France*

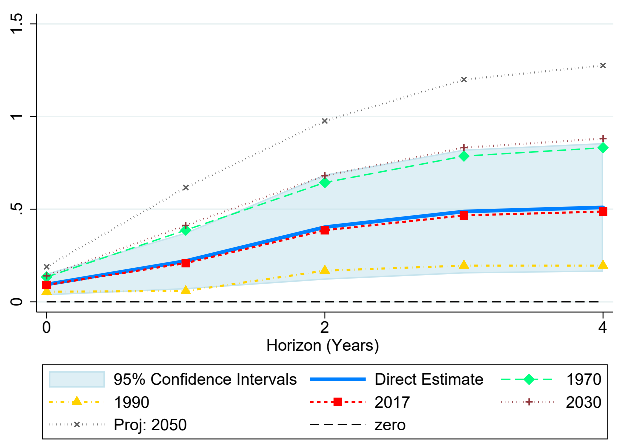

(d) *Demeaned WAP: France*

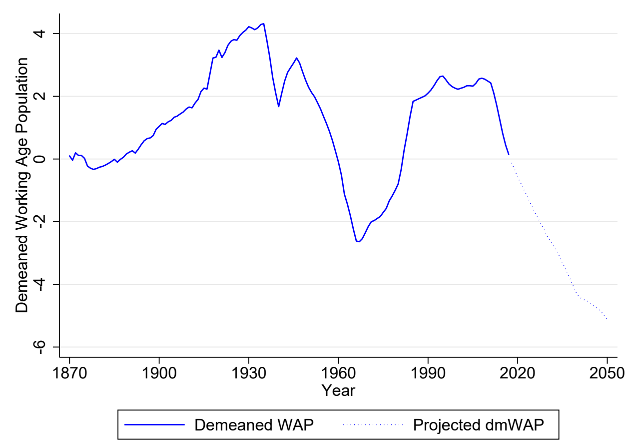

(e) *Conditional IRF: United Kingdom*

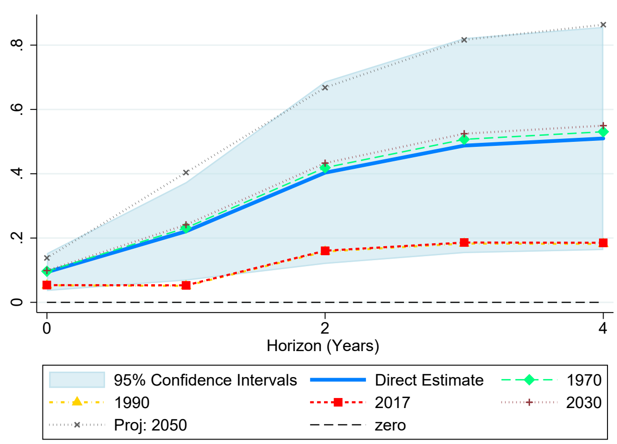

(f) *Demeaned WAP: United Kingdom*

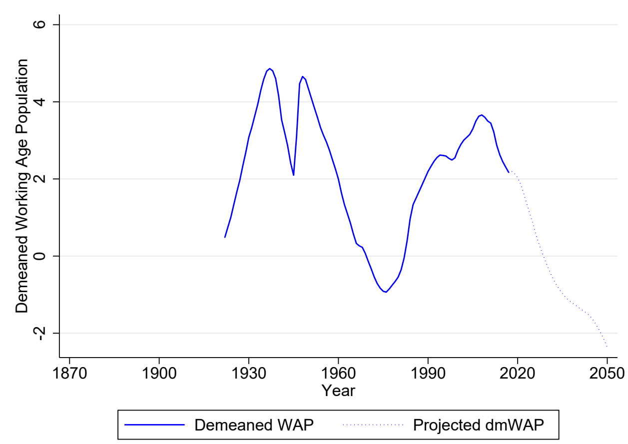

Notes: Direct estimates are the average estimate of excess money growth on inflation when population is at long run sample mean. Conditional IRFs reflect this relationship conditional on age distribution of each country in a given year.

interactive effects for the four demographic variables and all other controls.

Table 3: *Response of Inflation to Excess Money Growth: Full Age Distribution*

	0	1	2	3	4
m_{t-1}	0.09**	0.25**	0.45**	0.54**	0.59**
	(0.03)	(0.09)	(0.17)	(0.21)	(0.21)
D1	0.09	0.63*	1.45**	2.25**	3.20***
	(0.14)	(0.34)	(0.59)	(0.77)	(0.96)
D2	-0.14	-1.36	-3.41*	-5.50**	-8.15***
	(0.41)	(1.05)	(1.68)	(2.19)	(2.74)
D3	0.08	1.15	3.05*	5.04**	7.65**
	(0.43)	(1.11)	(1.72)	(2.25)	(2.83)
D4	-0.01	-0.33	-0.91	-1.53*	-2.37**
	(0.14)	(0.37)	(0.57)	(0.75)	(0.95)
$m_{t-1} \times D1$	0.07***	0.14***	0.20**	0.25*	0.26*
	(0.02)	(0.04)	(0.08)	(0.12)	(0.13)
$m_{t-1} \times D2$	-0.23***	-0.50***	-0.72***	-0.94***	-1.04***
	(0.05)	(0.12)	(0.21)	(0.31)	(0.34)
$m_{t-1} \times D3$	0.24***	0.56***	0.84***	1.11***	1.25***
	(0.06)	(0.13)	(0.22)	(0.31)	(0.36)
$m_{t-1} \times D4$	-0.08***	-0.19***	-0.30***	-0.40***	-0.46***
	(0.02)	(0.05)	(0.07)	(0.11)	(0.12)
Macro $X_{i,t}$	✓	✓	✓	✓	✓
Finance $X_{i,t}$	✓	✓	✓	✓	✓
Interact $X_{i,t}$	✓	✓	✓	✓	✓
F-test: D1-D4	9.50	12.53	10.91	7.70	8.08
F-test: D1-D4$\times m_{t-1}$	7.04	8.38	9.95	13.50	14.91
R2	0.71	0.67	0.67	0.67	0.68
AIC	6843	9041	9922	10566	11041
N	1342	1342	1342	1342	1342

Table reports estimations of IRFs for money growth under working age population, and fourth order polynomial specifications. Clustered standard errors are in parenthesis with $\dagger p < 0.15, *p < 0.10, **p < 0.05, ***p < 0.01$. $X_{i,t}$ significance. Regressions include three lags of the dependent variable as well two lags and contemporaneous values of all additional controls. F-tests are for joint significance of demographic controls (one for direct controls and another for interactions with lagged money growth).

The cumulative excess money growth coefficient at $h = 4$, a five year horizon, is nearly the same as the equivalent estimate Table 1, with average estimates of 59% correlation between excess money growth and inflation. While the regression coefficients in Table 3 on both the demographic variables directly and their interactions are not easily interpretable economically (as before), it's clear from reported F-tests, that both the direct estimates and indirect interactions are jointly significant at all horizons. The significance of all four interaction terms with demographics suggests that the model found useful variation in each term of the polynomial in improving fit of the prediction across the age distribution. I show in Appendix C that four terms is ideal for this setup, but results are similar with fewer/more.

The compositional relationships of population age structure on inflation directly are weaker at short horizons. This is in contrast to strong significance of all demographic terms in the specification of Table 1 Column 7. Short horizons are omitted from that table for exposition, but all four demographic variables were significant at the 1% level at all fifteen horizons pictured in Figure 3. While it is perhaps reasonable that aging takes time to influence inflation directly, much of the variation picked up by population age structure coefficients in Table 1 comes through these indirect channels of the population relationship with excess money growth. The sensitivity of inflation to money, conditional on age structure, is strongly significant at all horizons for each of these variables. Moreover, unlike estimates in Table 2, these parameters are fairly insensitive to model specification as shown in Appendix C, Table 7.

As in the baseline average estimates from Table 1, I can plot the *compositional* relation between age and inflation across the life cycle by backing out the implied coefficients as in Figure 4. More important for my work the are *indirect* impact of age structure coming through changes in excess money growth. These are estimated in the same way, but using the interaction term and as such are zero when the policy variable is zero. In Figure 7a, the solid blue line represents the *composition* terms: the average relationship between an age group and inflation unrelated to excess money growth. Though some significance is lost, the overall coefficient estimates by age group are roughly the same as in Figure 4. These once again match well with results from Juselius and Takáts (2021) whose work studies this relationship directly. In addition to these, I show the net effect of adding *indirect* impact of age: those which come through changes in money growth. These are constructed from the composition coefficients by age group and adding/subtracting a half standard deviation change in excess money growth times the age specific term implied by the interaction coefficients. Figure 7b simply shows these *indirect* coefficients alone for a positive money growth shock, which are the difference between the dashed green and solid blue lines in Figure 7a.

The total impact of aging, including indirect channels, shows that the indirect effects more or less reinforce the age specific coefficients. This is somewhat easy to see from the fact that each interaction in Table 3 has the same sign as its non-interactive counterpart. These indirect coefficients are positive for young dependents and very early career workers,

Figure 7: *Age Specific Estimates at H = 4: Composition + Indirect Age Coefficients*

(a) *Composition + Indirect Coefficients* **(b)** *Indirect Coefficients Alone ($+0.5\sigma m_{t-1}$ Shock)*

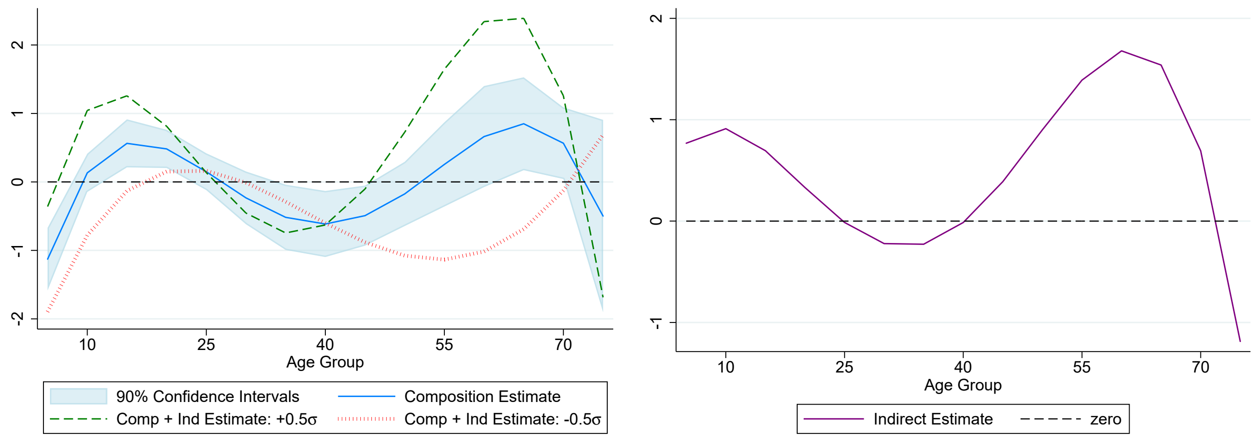

negative for early-to-mid career workers and then strongly positive again for late career (> 55) workers and retirees. Unlike the direct (composition) age coefficients, these do not sum to zero, and what matters is not whether a given coefficient is positive or negative, but their relative size. Moving from a population structure with large weight in young or old populations, to one with weight in the early working ages population will reduce transmission, even leading to a negative net effect, as will be the case below. This could be true even if all of the coefficients in Figure 7b were positive, as the deviations from population sample mean always sum to zero. As such an increase in one group necessarily reflects a decrease in another, and what matters for the net effect of a change in population age structure are how relative weights from population deviations interact with coefficients from Figure 7b.

In Figure 7b, the indirect effect drops substantially for young-to-middle aged workers and then begins to sharply rise for workers in mid-to-late career, with a peak in the five year groups on either side of retirement. This quite similar to the stylized picture of money holdings by age from the SCF in Figure 1. Here though I also see the effect of young dependents, who might exert pressure on money demand through spending needs of their parents. A theory of money, needed to finance consumption, may explain these dynamics. Fernández-Villaverde and Krueger (2007) show that consumption has a dramatic hump-shape, peaking somewhere in the mid-late fifties and remaining high throughout early retirement before falling during later retirement years. This is more or less the same picture I see for the indirect effect in the second half of life here. My empirical findings suggest that more theoretical work should be done to investigate how late career workers and early retirees might be an important conduit for monetary transmission to inflation, by generating these life-cycle profiles. It seems likely based on these pictures that some wealth and consumption dynamics, which match these shapes could play an important role.

There may be other mechanisms at work. If more experienced workers are more important for the wage pressures through bargaining this might act as a push factor for inflation. This is a main hypothesis regarding inflationary aging, but may be important for priming the monetary transmission if these wages affect money demand. Other explanations that have been put forth, focused on interest rate policy, suggest that households with large accumulation of wealth become *more* sensitive to changes in interest rates and therefore more likely to strengthen the effect of policy. However, this is the opposite of the effect expected in work such as Wong (2019),[13] where young households are more sensitive via a different, borrowing constraint, channel. One could imagine such mechanisms for safe interest rate transmission to potentially operate similarly to money if it was modeled as a competing asset as in Aoki et al. (2019).

The visualizations in Figure 7 are useful, but do not yet provide a clear idea of how recent demographic trends have contributed to monetary transmission. I conduct the same KBO decomposition as with my WAP estimates above, showing the impulse response functions for a one percentage point increase in excess money growth conditional on the age distribution of: the United States, France and the United Kingdom. These results are plotted in Figure 8.

In Figure 8a each line represents estimates conditional on a particular realization of the population age distribution for the United States. These then correspond to four values for the constructed demographic variables, which capture movements in deviations of population age shares from their long run mean. As such, each IRF represents the *direct* estimate of the inflation response to excess money growth (plotted by the blue line) plus the contribution from each of these four *indirect* demographic channels. The first twenty years of this exercise offer the same, though more dramatic, picture of the change in this relationship than those seen in the previous section. The great inflation of the 1970s are a time when the estimated impact of aging on the money-inflation relationship is strongly positive, with changes in excess money growth rate accommodated by aging such that a one-to-one relationship materialized within three years. During the great moderation, represented here by the estimated response in 1990 (and similar though not plotted 2000) this relationship completely disappears, as demographic headwinds create an almost puzzlingly strong negative money-inflation relationship.

What drives the dramatic shift from 1970 to 1990? To understand more clearly I show

[13]Notable of course is that in her work it is not money, but interest rate policy that matters, which perhaps could have different transmission.

Figure 8: *Response of Inflation to Excess Money Growth, KBO Decomposition: Full Age Controls*

(a) *Conditional IRF: United States*

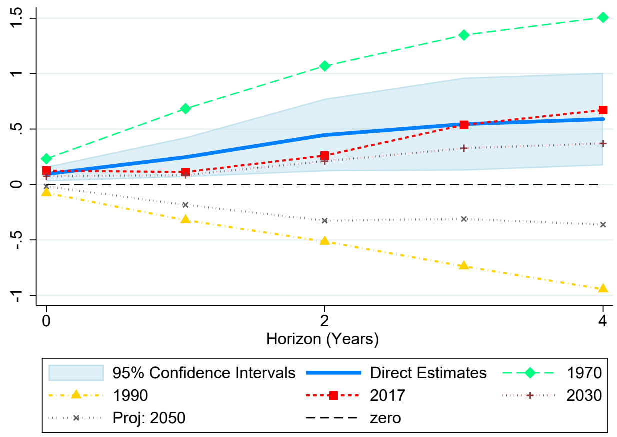

(b) *Population Shares: United States*

(c) *Conditional IRF: France*

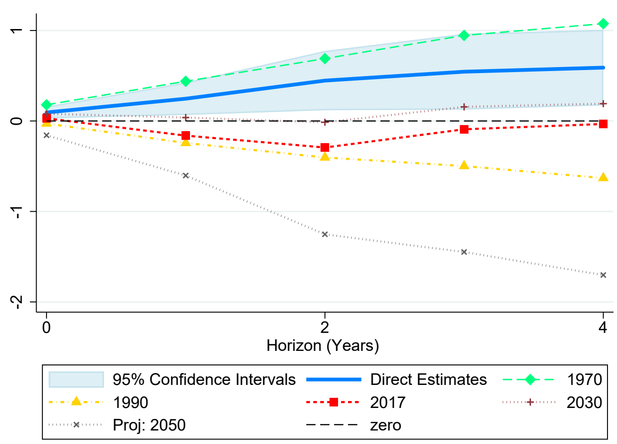

(d) *Population Shares: France*

(e) *Conditional IRF: United Kingdom*

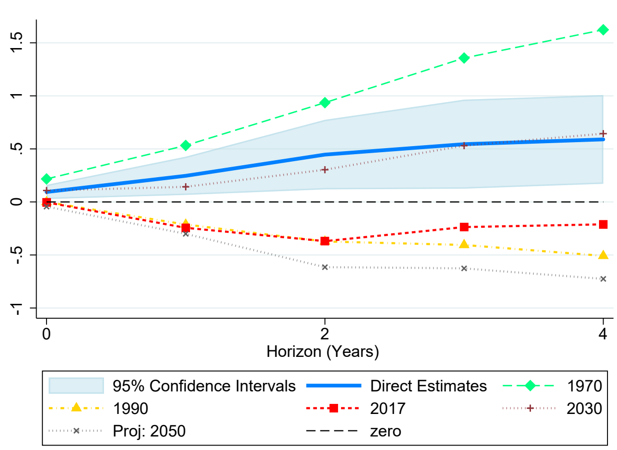

(f) *Population Shares: United Kingdom*

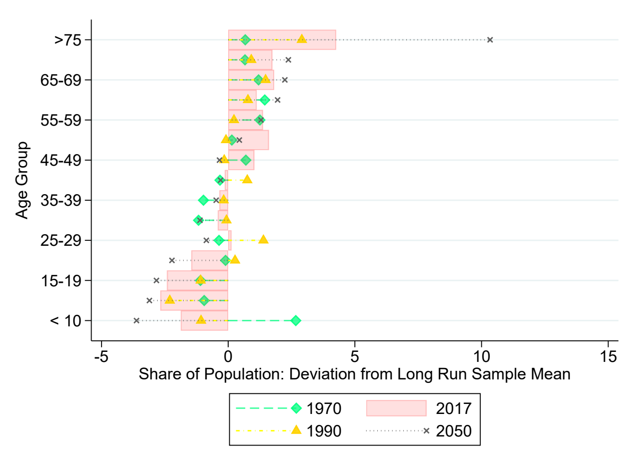

Notes: Direct estimates are the average estimate of money growth on inflation when population is at long run sample mean. Conditional IRFs are effects of money growth rates conditional on age distribution of each country in a given year.

the demeaned population shares for each of these years in the right-hand panels of Figure 8. Looking at the United States in Figure 8b there are a few large swings from the 1970s (green diamonds) to the 1990s (gold triangles). The first is that the population moves from one with a large share of young dependents, by historical standards, to one with relatively few. These swing from higher than historical sample averages (here the zero line) when the boomers make up this group in the 1970s (aged between 6 and 24 in 1970) to lower for all but the youngest group.[14] Looking back on the indirect coefficients across age groups plotted in Figure 7b, these are ages with quite a strong positive impact on the money-inflation relationship so this reduction shrinks the effect. Additionally, there is an increase on the younger working aged group as the boomer generation has moved into this category. By 1990 stage boomers are between 26 and 44 years old, occupying the early-mid working life trough in Figure 7b. Meanwhile population aging has not yet begun to take off in the United States, with the relatively small silent generation making up late career workers who have the strongest positive impact on the money-inflation relationship. The great moderation, it would seem, was the perfect demographic storm for a movement from strong to weak money pass through to prices, largely due to the size of the boomers.

By 2017, the boomer cohort is at the peak of the money growth relationship aged 53-71. In the United States this would suggest a resurgence of money transmission. This effect is limited by the massive increase in the oldest age cohorts, who contribute negatively, due to continued increases in life expectancy over this period, as well as the still small groups of young dependents. According to the OECD, life expectancy at age 65 for men in the United States increased by 2.9 years from 1990 and 2017 from 15.1 to 18 years. This increase is even more dramatic in the United Kingdom (14 to 18.8 years) and France (15.7 to 19.6 years). The combination of having relatively fewer young dependents and relatively more old age retirees leads the effects in these two countries to remain quite low in the most recent data. Projecting forward it is indeed the increasing importance of the eldest age cohort and continued decreasing share of young dependents that drive further downward pressure on money-inflation estimates.

It is clear from the right-hand panels of Figure 8 that the 75 age group has an outsized effect on future projections. As a result, I wish to caution putting too much weight on point estimates here. A known issue of using demographic controls is that over-fitting of the polynomial over the age distribution can tend to lead to extreme values for the youngest and oldest age groups. Further these out-of-sample projections feed in unprecedented increases in this age group, with no comparable historical data. While decreasing money demand

[14]In 1990 young dependents are populated by the relatively small Generation X and the eldest of the larger millennial cohort.

in old age is consistent with the stylized picture in Figure 1, the reversal for the oldest group is dramatic. It is true that the extreme shift of the estimates of future projections in Figure 7b may be subject to such over-fitting. I explain my choice of polynomial order in Appendix C, showing that the fourth degree balances an improved model fit coming from the increase in parameters (measured by AIC), with a worsened out-of-sample fit measured through a method developed in Bilger and Manning (2015). Moreover, I show that the implied age coefficients of the indirect estimates by age groups for higher orders look broadly similar, with more extreme values for old age groups, suggesting that $k = 4$ captures the relevant age variation without allowing for these swings to become overly extreme. In Appendix C.3, I present estimates and projections using lower order versions of this polynomial. These give a broadly consistent story of the changing transmission from 1970 to present, with less extreme negative impacts from old age groups in future projections.

4.3. Where is the Missing Inflation?

Where is the missing inflation? My two separate demographic controls paint similar pictures over a long run view, but give conflicting answers in the years since the global financial crisis. If one thinks that the estimates in Section 4.2 are potentially over-fitting the age data, then they might be content with the idea that working age population suggests that demographic headwinds have still been in place over the last decade, perhaps reconciling seemingly weak inflation in response to relatively loose monetary policy. However it is *possible* that the age distribution has already become accommodating and that other forces have been working in the other direction.

Coibion, Gorodnichenko, and Ulate (2019) suggest that an expectations-augmented Phillips curve can account for inflation weakness post global financial crisis in a wide range of countries. With the deflationary pressures of the Great Recession lasting until about 2018. Part of the inability of standard Phillips curves to show this is that economic slack is not well accounted for by traditional measures such as the unemployment rate. Falling participation, for example, might signify more slack in the economy that would be missed otherwise. My present analysis isn't able to account for such expectations, nor even for dynamics in employment. As a result if I were to try to forecast inflation using my model I would, like the traditional Phillips curve, significantly overestimate inflation.

To get a sense of this I plot the predicted current year inflation rate using my specification from Table 3 for a subsample of countries in Figure 9. To be clear I expect this to be a poor estimate of inflation. This paper is not making any attempts at forecasting, and has explicitly focused on attempting to estimate medium-to-long term policy interactions. However, many

Figure 9: *Prediction of Current Inflation*

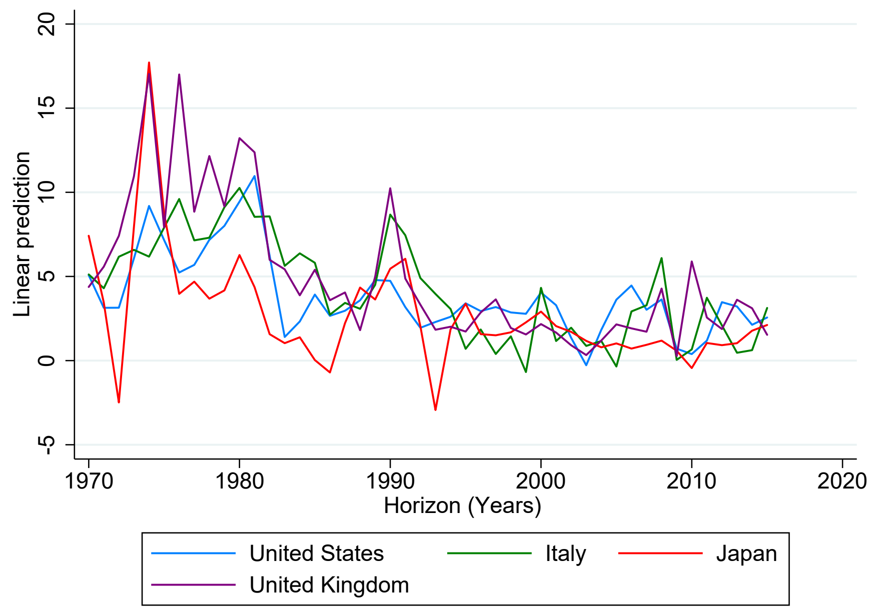

of my macroeconomic and financial controls are quite important in determining inflation at shorter horizons and can give some sense as to whether my estimated age-structure transmission mechanism, which is quite strong by 2017 in the analysis of subsection 4.2, would have counter-factually predicted a large inflation takeoff. It's clear from Figure 9 that in spite of my sharp predicted increase in the transmission of money growth to inflation over this period, the model has not predicted any such takeoff from inflation levels from what might have been expected during the great moderation.

These results are reassuring in that they provide some context such that the dramatic transmission mechanisms plotted for the United States in Figure 8a do not necessarily imply a takeoff of inflation in their own right. Rather they suggest that age structure has shifted toward accommodation of money transmission ceteris paribus. Of course many factors might be working against this in the short run, and those discussed in Coibion et al. (2019) suggest that an overhang from the global financial crisis is plausible. However, with inflation response to the Covid crisis appearing more robust than many expected my results suggest that policy makers should be careful to disregard monetary aggregates as a potential driver, particularly in the United States where demographics suggest a robust relationship at present and where Fed interventions have led to incredible spikes in M2.

5. Estimates Using Divisa Monetary Aggregates

The above estimates use simple sum monetary aggregates of broad money (M2). These are used primarily out of necessity as they are available for a large panel of advanced economies going back to 1870. The length of this panel allows for an estimation that spans generations, capturing more than the period after the birth of the large boomer cohort, while also leveraging the large degree of variability in timing and magnitude of population changes across countries. However, there are known issues with these standard monetary aggregates, which fail to capture the fact that monetary assets are not perfect substitutes and therefore should not be aggregated without first constructing weights that are dependent not only on quantities of money, but also user costs of various monetary assets. The Divisia method of such an aggregate was proposed in Barnett (1980). This built on earlier work from Barnett (1978), who calculates user costs of money. Barnett (1980) proposes the use of either Divisia or Fisher ideal index methods for aggregation of monetary quantities, which take the relative costs of various types of money into account in aggregating over simple sum measures. This has since been shown to outperform official sum monetary indices in a number of contexts.[15] Barnett et al. (1984) show that unlike simple-sum measures of money, measures of money demand using Divisia aggregates show relatively strong parametric stability. This suggests that at least some of the conventional understanding of a shift in money demand appears to come from improper measurement.

In this section I show that my results on demographic impact on money growth transmission hold when when using this Divisia measure of money for United States from 1967-2017. Monthly data are available from the Center for Financial Stability. While results are broadly similar across different Divisia aggregates I present results for M2 as they are the counterpart to the measures of "broad" money available in the Jordà et al. (2017) data used in my main specifications. The data in my above analysis above is annual, and while many controls are available at higher frequency elsewhere, the population-by-age data from the HMD are not. Moreover slow moving demographics have little meaningful variation at higher than annual frequencies. As such I use annual growth in Divisia money to estimate a non-panel version of Equation 2. Because of the smaller sample I also am limited in controls, including two lags of demeaned changes in: bill rates, log real consumption, investment-to-gdp, and current account-to-gdp, and population growth rate. The results of this estimation for both the WAP and polynomial controls are in Table 4.

[15]See for example: Barnett et al. (1984) and Belongia and Ireland (2016)

Table 4: *Excess Divisia Money Growth and Inflation*

	WAP					Full Age (k=2)				
	0	1	2	3	4	0	1	2	3	4
m_{t-1}	0.31***	0.77***	0.99***	0.95**	0.95**	0.35**	0.93***	1.19***	1.14***	1.07**
	(0.10)	(0.22)	(0.31)	(0.32)	(0.37)	(0.12)	(0.24)	(0.32)	(0.32)	(0.39)
WAP	0.11	-0.17	-0.76†	-1.75***	-2.74***					
	(0.12)	(0.38)	(0.46)	(0.41)	(0.43)					
D1						0.14	0.15	-0.23	-0.90**	-1.60***
						(0.12)	(0.26)	(0.36)	(0.41)	(0.53)
D2						-0.10	-0.12	0.14	0.62*	1.10**
						(0.10)	(0.19)	(0.26)	(0.31)	(0.40)
$m_{t-1}\times$ WAP	-0.07**	-0.18**	-0.29***	-0.33***	-0.31***					
	(0.03)	(0.08)	(0.10)	(0.10)	(0.10)					
$m_{t-1}\times$ D1						-0.04	-0.12*	-0.18*	-0.19*	-0.14
						(0.03)	(0.06)	(0.08)	(0.09)	(0.12)
$m_{t-1}\times$ D2						0.02	0.08†	0.12*	0.12*	0.09
						(0.02)	(0.04)	(0.06)	(0.06)	(0.08)
F-test: D1-D2						0.91	0.34	0.39	3.39	6.94
F-test: D1-D2 $\times m_{t-1}$						3.20	3.18	4.97	6.78	5.00
AIC	121	185	211	218	233	120	184	212	224	243
R2	0.96	0.95	0.96	0.97	0.97	0.96	0.95	0.96	0.97	0.97
N	43	43	43	43	43	43	43	43	43	43

Table reports estimations of IRFs for response of inflation to Divisia money growth under working age population, and second order polynomial specifications. Heteroskedastic standard errors are in parenthesis with $\dagger p < 0.15, *p < 0.10, **p < 0.05, ***p < 0.01$. $X_{i,t}$ significance. Regressions include three lags of the dependent variable as well two lags of demeaned changes in bill rates, log real consumption, investment per gdp, current account per-gdp and population growth.

The immediate takeaway is that the average estimates for excess money growth, measured here as the difference between growth rate in Divisia M2 and output growth, increases to the expected unitary value within three years. This better fit with simple quantity theory intuition is expected given the existing literature showing that these measures of money are more closely linked to inflation, and better estimate money demand. In spite of this the pass through of excess Divisia money growth still appears to have strong state dependence with demographics, with working age population in particular correlating strongly. The much smaller sample significantly weakens the polynomial effect, though demographic interactions are still jointly significant at all horizons at the 10% level and well below 5% level for $h > 1$. I use a second order polynomial, though a third yields quite similar results, from both information criteria and overfitting concerns studied in Appendix C. In particular the age specific coefficients become quite volatile in this small sample for polynomials of order greater than 3. This is likely due to fitting a much smaller amount of demographic variation. However the working age population estimates are larger and more significant in this small sample than in the baseline estimates.

Figure 10 shows that the implied change in inflation response to changes in excess money looks quite similar to the full sample analog in Figure 8a, scaled by the fact that

Figure 10: *Response of Inflation to Excess Money Growth, KBO Decomposition: Divisia M2 Aggregates*

(a) *Conditional IRF (WAP): United States*

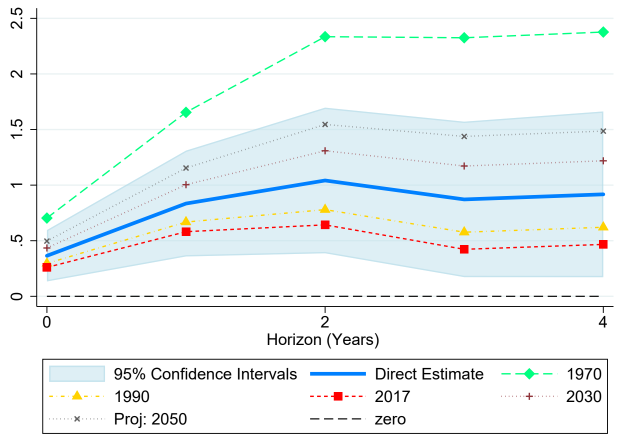

(b) *Demeaned WAP: United States*

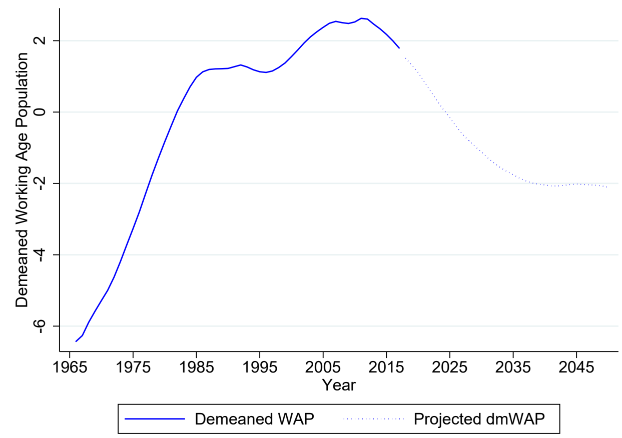

(c) *Conditional IRF (3rd Order Poly): United States*

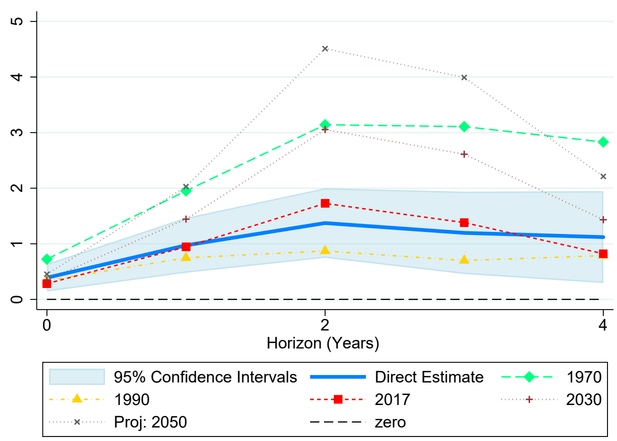

(d) *Demeaned Population Shares: United States*

Notes: Direct estimates are the average estimate of excess money growth on inflation when population is at long run sample mean. Conditional IRFs reflect this relationship conditional on age distribution of each country in a given year.

the average coefficient on excess money growth now rises much faster. While I caution against putting too much weight on these point estimates, in particular those projecting into the future, they are suggestive that demographic tail-winds may have provided strong inflationary support for changes in money growth during the 1970s, and headwinds as the boomers entered the workforce in full, and may again support money and inflation in the future. More work studying the age-state dependence of money using these Divisia aggregates would be valuable, particularly if they could be extended to historical samples where the single cohort is playing less of an outsized role on parameter estimates as with my estimates in the full specification.

6. Conclusion

Age structure does an excellent job accounting for a boom-bust cycle of the medium-to-long transmission of excess money growth rates to inflation over the past fifty years. This appears to be a feature of three forces: falling fertility (leading to smaller shares of young dependents), increasing old age dependents through increased life expectancy, and the movement of the baby boomer cohort through the age distribution. If the results outlined above are to be taken seriously then policy makers should be cautious of the idea that the policy environment of the last thirty years is the norm. While decreased fertility and increasing longevity are likely going to continue in advanced economies, the aging of the boomers may be allowing for a return of the money-inflation relationship at present.

There are a number of caveats that should be made clear. While my results appear to be quite robust, this analysis falls short of causal identification. Finding an identification strategy that combines exogenous money growth shocks with plausibly exogenous demographic shocks would add credibility to any such estimates, but are not readily available. Because of this, I see these results as also calling for further theoretical work to explore potential mechanism more concretely, while allowing for a robust quantitative analysis. These of course can work together, as clearer ideas as to the mechanisms that might explain the aggregate data should help point to areas where micro data may shed light on how such mechanisms may operate.

While I provide some evidence in support of the theses presented in Goodhart and Pradhan (2020), I do not focus specifically on how aging affects inflation directly, rather exploring how population aging may work through indirect channels of money transmission. These indirect channels appear important, and should not be ignored in this broader discussion of the effects of aging. While all estimates explain the transition from great inflation to great moderation, the implications for the future are nuanced. If we consider working age population, there is some time before the United States would face money driven inflation, but aging will continue to push toward stronger money transmission. Using more information from the age distribution, it is possible that the United States is already experiencing positive demographic pressure on the money-inflation relationship, but that significant aging beyond this may actually reduce this effect as the group of individuals over 75 grows substantially. Policy makers should ignore money at their own risk. While the widely accepted notion that a low inflation policy regime has removed the need to carefully track monetary aggregates, this paper suggests that aging may just as easily explain some of this shift. If this is indeed true, then we cannot expect the relationship to remain constant over the coming decades.

References

Aksoy, Y., H. S. Basso, R. P. Smith, and T. Grasl (2019). Demographic structure and macroeconomic trends. *American Economic Journal: Macroeconomics 11*(1), 193–222.

Aoki, K., A. Michaelides, and K. Nikolov (2019). Inflation, money demand and portfolio choice. *Available at SSRN 2399362*.

Assenmacher-Wesche, K. and S. Gerlach (2007). Money at low frequencies. *Journal of the European Economic Association 5*(2-3), 534–542.

Barnett, W. A. (1978). The user cost of money. *Economics letters 1*(2), 145–149.

Barnett, W. A. (1980). Economic monetary aggregates an application of index number and aggregation theory. *Journal of econometrics 14*(1), 11–48.

Barnett, W. A., E. K. Offenbacher, and P. A. Spindt (1984). The new divisia monetary aggregates. *Journal of Political Economy 92*(6), 1049–1085.

Belongia, M. T. and P. N. Ireland (2016). Money and output: Friedman and schwartz revisited. *Journal of Money, Credit and Banking 48*(6), 1223–1266.

Benati, L. (2005). Long-run evidence on money growth and inflation. *Bank of England Quarterly Bulletin, Autumn*.

Bilger, M. and W. G. Manning (2015). Measuring overfitting in nonlinear models: a new method and an application to health expenditures. *Health economics 24*(1), 75–85.

Blinder, A. S. (1973). Wage discrimination: reduced form and structural estimates. *Journal of Human resources*, 436–455.

Bobeica, E., C. Nickel, E. Lis, and Y. Sun (2017). Demographics and inflation. Technical report, ECB Working Paper.

Cloyne, J. S., Ò. Jordà, and A. M. Taylor (2020). Decomposing the fiscal multiplier. Technical report, National Bureau of Economic Research.

Coibion, O., Y. Gorodnichenko, and M. Ulate (2019). Is inflation just around the corner? the phillips curve and global inflationary pressures. In *AEA Papers and Proceedings*, Volume 109, pp. 465–69.

de Albuquerque, P. C., J. Caiado, and A. Pereira (2020). Population aging and inflation: evidence from panel cointegration. *Journal of Applied Economics 23*(1), 469–484.

Edo, A. and J. Melitz (2019). The primary cause of european inflation in 1500-1700: Precious metals or population? the english evidence.

Eggertsson, G. B., N. R. Mehrotra, and J. A. Robbins (2019). A model of secular stagnation: Theory and quantitative evaluation. *American Economic Journal: Macroeconomics 11*(1), 1–48.

Fair, R. and K. Dominguez (1991). Effects of a changing us age distribution on macroeconomic equations. *The American economic review*, 1276–1294.

Fernández-Villaverde, J. and D. Krueger (2007). Consumption over the life cycle: Facts from consumer expenditure survey data. *The Review of Economics and Statistics 89*(3), 552–565.

Fratianni, M., M. Gallegati, and F. Giri (2021). International historical evidence on money growth and inflation: The role of high inflation episodes. *The BE Journal of Macroeconomics*.

Friedman, M. and A. J. Schwartz (1963). Money and business cycles. *The Review of Economics and Statistics*, 32–64.

Gagnon, E., B. K. Johannsen, and D. Lopez-Salido (2021). Understanding the new normal: The role of demographics. *IMF Economic Review*, 1–34.

Gechert, S. and A. Rannenberg (2014). Are fiscal multipliers regime-dependent? a meta regression analysis. Technical report, IMK working paper.

Gertler, P. and B. Hofmann (2018). Monetary facts revisited. *Journal of International Money and Finance 86*, 154–170.

Goodhart, C. (1989). The conduct of monetary policy. *The Economic Journal 99*(396), 293–346.

Goodhart, C. and M. Pradhan (2020). *The great demographic reversal: Ageing societies, waning inequality, and an inflation revival*. Springer Nature.

Human Mortality Database (2019). University of California Berkeley, and Max Planck Institute for Demographic Research. https://www.mortality.org.

Ilzetzki, E., E. G. Mendoza, and C. A. Végh (2013). How big (small?) are fiscal multipliers? *Journal of monetary economics 60*(2), 239–254.

Jordà, Ò. (2005). Estimation and inference of impulse responses by local projections. *American economic review 95*(1), 161–182.

Jordà, Ò., K. Knoll, D. Kuvshinov, M. Schularick, and A. M. Taylor (2019). The rate of return on everything, 1870–2015. *The Quarterly Journal of Economics 134*(3), 1225–1298.

Jordà, Ò., M. Schularick, and A. M. Taylor (2017). Macrofinancial history and the new business cycle facts. *NBER macroeconomics annual 31*(1), 213–263.

Jordà, Ò., M. Schularick, and A. M. Taylor (2020). The effects of quasi-random monetary experiments. *Journal of Monetary Economics 112*, 22–40.

Jordà, Ò. and A. M. Taylor (2016). The time for austerity: estimating the average treatment effect of fiscal policy. *The Economic Journal 126*(590), 219–255.

Juselius, M. and E. Takáts (2021). Inflation and demography through time. *Journal of Economic Dynamics and Control*, 104136.

Kitagawa, E. M. (1955). Components of a difference between two rates. *Journal of the american statistical association 50*(272), 1168–1194.

Kopecky, J. and A. M. Taylor (2022). The savings glut of the old: Population aging, the risk premium, and the murder-suicide of the rentier. Technical report, National Bureau of Economic Research.

Montiel Olea, J. L. and M. Plagborg-Møller (2021). Local projection inference is simpler and more robust than you think. *Econometrica 89*(4), 1789–1823.

Oaxaca, R. (1973). Male-female wage differentials in urban labor markets. *International economic review*, 693–709.

Plagborg-Møller, M. and C. K. Wolf (2021). Local projections and vars estimate the same impulse responses. *Econometrica 89*(2), 955–980.

Reis, R. (2021). Losing the inflation anchor.

Roffia, B. and A. Zaghini (2007). Excess money growth and inflation dynamics. *International Finance 10*(3), 241–280.

Romer, C. D. and D. H. Romer (2004). A new measure of monetary shocks: Derivation and implications. *American Economic Review 94*(4), 1055–1084.

Rudd, J. B. (2022). Why do we think that inflation expectations matter for inflation?(and should we?). *Review of Keynesian Economics 10*(1), 25–45.

Sargent, T. J. and P. Surico (2011). Two illustrations of the quantity theory of money: Breakdowns and revivals. *American Economic Review 101*(1), 109–28.

Schularick, M. and A. M. Taylor (2012). Credit booms gone bust: Monetary policy, leverage cycles, and financial crises, 1870-2008. *American Economic Review 102*(2), 1029–61.

Teles, P., H. Uhlig, and J. Valle e Azevedo (2016). Is quantity theory still alive? *The Economic Journal 126*(591), 442–464.

UN (2019). World population prospects. *Department of Economic and Social Affairs, Population Division (2019) Online Edition. Rev. 1.*

Wang, L. and T. Zhu (2021). Population aging and money demand. *Economics Letters 206*, 109984.

Wong, A. (2019). Refinancing and the transmission of monetary policy to consumption. *Unpublished manuscript*.

A. Construction of Demographic Controls

My estimates for effects across the entire age distribution build on the methodology of Fair and Dominguez (1991). Suppose one wanted to estimate the following relationship:

$$y_{i,t} = X_{i,t}\beta + \sum_{j}^{J} \alpha_j p_{j,i,t} + \mu_i + \nu + \epsilon_{i,t} \tag{6}$$

Where $y_{i,t}$ is any outcome of interest, $X_{i,t}$ is an arbitrary vector of controls, $\epsilon_{i,t}$ is error, μ_i country fixed effects, and ν a constant. The variables $p_{j,i,t}$ are shares of the population, divided into J bins. This is inestimable due to the perfect colinearity of the population shares $p_{j,i,t}$. Additionally, while one would prefer to take a granular approach to modeling population shares (allowing for a large number of, J, groups), these shares are highly colinear with one another, more so as their number increases. Finally in smaller samples it may be undesirable to fit such a large number of coefficients, particularly as in my case when interaction terms would also be needed. Fair and Dominguez (1991) proceeds by making the following two assumptions:

1. Letting α_j be the coefficient on population share $p_{j,i,t}$ of age group j in country i and time t. Assume that all of the effects of these coefficients across the age distribution sum to zero. In other words:
$$\sum_{j}^{J} \alpha_j = 0$$
2. Assume that the age coefficients α_j can be fitted with a K order polynomial. In other words:
$$\alpha_j = \sum_{k}^{K} \gamma_k j^k \tag{7}$$

All three problems are addressed. First it transforms the problem of estimating J coefficients into one of estimating K, as I will show in a moment. Second, the assumption that all age effects, α_j, sum to zero makes them jointly estimable, without having to drop the regression constant. Finally, by forcing the age effects to lie on a polynomial, the model requires that there be relatively smooth transitions from the effect of one age group to another. If one were to take only the first assumption (or omit the constant of the regression) it would be possible to estimate the effects in Equation 6. The results of such estimations when using five or ten year population age groups often lead to highly unstable coefficients that alternate signs quickly from one age group to the next. This is due to the high degree of colinearity between one age group and those immediately around it, which increases in J.

To see how this methodology works. Substitute α_j into Equation 6. I will assume in what follows a third order polynomial for exposition. This yields:

$$y_{i,t} = X_{i,t}\beta + \sum_{j}^{J} \left[(\gamma_0 + \gamma_1 j + \gamma_2 j^2 + \gamma_3 j^3) p_{j,i,t} \right] + \mu_i + \nu + \epsilon_{i,t} \tag{8}$$

The summing over both sides of Equation 7, and using the assumption that the sum of α_j must be zero, it can easily be shown that γ_0 is equal to:

$$\gamma_0 = -\frac{1}{J}\left[\gamma_1 \sum_j j + \gamma_2 \sum_j j^2 + \gamma_3 \sum_j j^3\right]$$

Thus far I have followed exactly the methodology of Fair and Dominguez (1991). From here one simply plugs the above expression for γ_0 into Equation 8, and given the fact that the first term, $\sum_j \gamma_0 p_{j,i,t} = \gamma_0$ by the fact that the population parameters sum to one, then this can be rearranged as:

$$y_{i,t} = X_{i,t}\beta + \gamma_1 \sum_j^J \left(p_{j,i,t} j - \frac{\sum j}{J}\right) + \gamma_2 \sum_j^J \left(p_{j,i,t} j^2 - \frac{\sum j^2}{J}\right) + \gamma_3 \sum_j^J \left(p_{j,i,t} j^3 - \frac{\sum j^3}{J}\right) + \mu_i + \nu + \epsilon_{i,t} \tag{9}$$

With the terms in parenthesis being the demographic variables. Because I wanted my estimates to follow the demeaned structure needed in the Kitagawa-Blinder-Oxaca (KBO) decomposition. I opted to instead estimate the following.

$$y_{i,t} = X_{i,t}\beta + \sum_j^J \alpha_j (p_{j,i,t} - \bar{p}_{j,t}) + \mu_i + \nu + \epsilon_{i,t} \tag{10}$$

The first few steps are exactly the same so I don't replicate them here. The problem however is just to rearrange terms slightly differently in order to construct the correct demographic variables for estimation. This simplifies to:

$$Dk_{i,t} = \left[\sum_j^J (p_{j,i,t} - \bar{p}_{i,t}) j^k - \sum_j^J j^k \sum_j^J (p_{j,i,t} - \bar{p}_{i,t})\right] \tag{11}$$

But only this first term is needed given that $\sum_j \gamma_0 (p_{j,i,t} - \bar{p}_{j,t}) = 0$, and not γ_0 as before. So rather the estimation becomes

$$y_{i,t} = X_{i,t}\beta + \gamma_1 \sum_j^J (p_{j,i,t} - \bar{p}_{j,t}) j + \gamma_2 \sum_j^J (p_{j,i,t} - \bar{p}_{j,t}) j^2 + \gamma_3 \sum_j^J (p_{j,i,t} - \bar{p}_{j,t}) j^3 + \mu_i + \nu + \epsilon_{i,t} \tag{12}$$

Where these $\sum_j (p_{j,i,t} - \bar{p}_{j,t}) j^k$ terms are the demographic controls used in the estimations in the paper.

B. Cointegration of Population Age Structures and Inflation

In this appendix, I present results of a cointegrating relationship between inflation and age structure in my data. All controls other than my two demographic specifications are stationary. Notably neither the demeaned working age population series, nor the Fair and Dominguez (1991) style demographic controls are stationary and are rather $I(1)$ processes.

I show this in Table 5, which reports the inverse χ^2 statistic for both Phillips-Perron and augmented Dickey-Fuller, Fisher-type panel unit root tests. P-values for these are shown in parentheses. The first half of this table shows that indeed I cannot reject a unit root for the cumulative inflation (my dependent variable used to construct local projection horizons), nor for any of the age structure variables. The lone exception is the Phillips-Perron statistic on the demeaned working age population series. All other controls are not reported as they are strongly stationary for all tests.

These variables are however appear to be stationary in differences as is shown by the reported. This is strongly rejected in the Phillips-Perron in all tests, failing to reject for D3 and D4 for the augmented Dicky-Fuller.

Table 5: *Unit Root Test: Cumulative Inflation and Age Structure Controls*

	Phillips-Perron	Augmented Dickey-Fuller
Cumulative Inflation	21.11	10.29
	(0.98)	(1.00)
dmWAP	69.30	13.51
	(0.00)	(0.99)
D1	4.30	3.46
	(1.00)	(1.00)
D2	2.52	2.84
	(1.00)	(1.00)
D3	2.11	2.93
	(1.00)	(1.00)
D4	1.68	3.33
	(1.00)	(1.00)
First Difference		
d.(Cumulative Inflation)	456.22	247.93
	(0.00)	(0.00)
d.(dmWAP)	123.61	97.95
	(0.00)	(0.00)
d.(D1)	95.17	62.89
	(0.00)	(0.00)
d.(D2)	94.62	50.97
	(0.00)	(0.00)
d.(D3)	92.07	44.81
	(0.00)	(0.15)
d.(D4)	88.39	40.63
	(0.00)	(0.27)

The purpose of the above was to show that my dependent variable and demographic controls are I(1) processes. As mentioned in the text, there have been a number of papers showing that there is a long run, cointegrating, relationship between population age structure and inflation. I now show that indeed this appears to be the case. In the effort of transparency, I report nine different tests of panel cointegration and their corresponding p-values in Table 6. These are shown for both specifications of demographic controls, WAP and D1-D4 jointly. In seven out of nine of these tests I reject the null hypothesis of no cointegration.

Table 6: *Test of Cointegrating relationship between cumulative inflation and and Demographic variables*

	Pedroni			Kao					Westerlund
	mPP	PP	aDF	mDF	DF	aDF	umDF	uDF	VR
dmWAP	4.70	6.88	11.29	3.15	4.27	0.40	4.60	8.56	1.25
	(0.00)	(0.00)	(0.00)	(0.001)	(0.00)	(0.34)	(0.00)	(0.00)	(0.11)
D1-D4	0.58	4.09	10.64	1.88	1.94	-0.63	3.84	5.30	-2.61
	(0.28)	(0.00)	(0.00)	(0.03)	(0.03)	(0.27)	(0.00)	(0.00)	(0.01)

The existence of such a relationship should ease concerns of spurious results in my estimates as well as imply that coefficients are *super*-consistent. Differencing these non-stationary series would not be appropriate in my estimations.

C. Model Selection

In this appendix I discuss choice of model selection. I show that my objects of interest, demographic interactions with money growth rates are relatively insensitive to choice of control sets for my preferred Fair and Dominguez (1991) style polynomial controls. Direct impacts of money growth on cumulative inflation are similar to those in Table 1 using standard local projections (ie without the KBO interactions between controls). In Table 7 I run my two main specifications at the full five year cumulative horizon under various model specifications.

There are a few conclusions from Table 7. Inclusion of controls: lagged dependent variables, country fixed effects, and lags and contemporaneous values of macro and financial variables generally improve the explanatory power of the model, both in terms of the adjusted R^2 terms and AIC. While the effects of WAP are insignificant in the simple specification of the model without additional controls, the polynomial demographic terms and interactions are significant throughout, and after inclusion of lagged dependent variables, see relatively small changes in their magnitude. The fully saturated specifications are used in the baseline figures of the paper and correspond to columns 5 and 10, but the quantitative implications for the polynomial controls not greatly changed across most of these specifications.

Table 7: *Five Year $h = 4$ Cumulative Response of Inflation, Various Models*

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
m_{t-1}	1.11***	0.44**	0.38**	0.53***	0.51**	1.15***	0.63***	0.56**	0.62**	0.59**
	(0.27)	(0.18)	(0.17)	(0.16)	(0.18)	(0.24)	(0.19)	(0.19)	(0.22)	(0.21)
dmWAP	0.21	-0.11	-0.16	0.34	0.55*					
	(0.27)	(0.22)	(0.30)	(0.33)	(0.29)					
D1						4.29***	3.46***	3.78***	3.44***	3.20***
						(1.20)	(0.80)	(0.88)	(1.13)	(0.96)
D2						-11.14***	-9.31***	-10.49***	-9.14**	-8.15***
						(3.58)	(2.44)	(2.63)	(3.11)	(2.74)
D3						10.74**	9.19***	10.55***	8.88**	7.65**
						(3.79)	(2.63)	(2.82)	(3.16)	(2.83)
D4						-3.42**	-2.98***	-3.45***	-2.82**	-2.37**
						(1.28)	(0.91)	(0.97)	(1.05)	(0.95)
$m_{t-1}\times$ dmWAP	-0.07	-0.01	-0.01	-0.13***	-0.15***					
	(0.10)	(0.08)	(0.07)	(0.04)	(0.05)					
$m_{t-1} \times D1$						0.35**	0.20	0.18	0.25**	0.26*
						(0.16)	(0.14)	(0.13)	(0.11)	(0.13)
$m_{t-1} \times D2$						-1.34***	-0.83**	-0.73*	-0.99***	-1.04***
						(0.43)	(0.38)	(0.36)	(0.30)	(0.34)
$m_{t-1} \times D3$						1.59***	1.02**	0.88**	1.20***	1.25***
						(0.46)	(0.40)	(0.38)	(0.31)	(0.36)
$m_{t-1} \times D4$						-0.57***	-0.37**	-0.32**	-0.44***	-0.46***
						(0.16)	(0.14)	(0.14)	(0.11)	(0.12)
Lags $\Delta\pi_0$		✓	✓	✓	✓		✓	✓	✓	✓
CountryFE			✓	✓	✓			✓	✓	✓
Macro $X_{i,t}$				✓	✓				✓	✓
Financial $X_{i,t}$					✓					✓
R2	0.42	0.50	0.48	0.61	0.64	0.54	0.60	0.58	0.66	0.68
AIC	11969	11763	11703	11318	11229	11676	11487	11444	11132	11041
F-test (dem int.)						8.31	7.47	6.94	16.00	14.91
p (F-test)						0.00	0.00	0.00	0.00	0.00
N	1342	1342	1342	1342	1342	1342	1342	1342	1342	1342

Table reports estimations cumulative five year response for money growth under various specifications. Clustered standard errors are in parenthesis with $*p < 0.10, **p < 0.05, ***p < 0.01$. Lags $\delta\pi_0$ refer to inclusion of three lags of the dependent variable (change in inflation at time zero). Macro controls refer to two lags and contemporaneous demeaned values of population growth, deficits to GDP, and the change in: log output, log real consumption, investment to gdp, global GDP, log debt to gdp ratio. Financial controls include two lags and the contemporaneous demeaned values of log real equity, change in credit to GDP and the change in the short term bill rate. F test and corresponding p-value is the test of joint significance of the four polynomial demographic variables interacted with lagged money growth. D2/D3/D4 scaled by 10/100/1000 for readability. All models contain dummies for WW1 and WW2 and their interaction with demographic variables.

C.1. Choice of Demographic Polynomial Length

The choice of the order of the polynomial for fitting the demographic effects balances the desire for flexibility over age specific parameters with both tractability and danger of over-fitting the effects of particular age groups. To investigate my choice of a fourth order effect I study two measures of model fit on the full specification used in the paper. The first is an Akaike information criterion which balances goodness-of-fit via a likelihood function with a penalty function for additional parameters. This function, while useful for assessing information content from additional controls, is strictly decreasing for polynomial orders from 1 to 6. In an attempt to also understand issues related to overfitting I also include a measure of out-of-sample fit of the model using the method of Bilger and Manning (2015). This method produces estimates of overfitting by randomly selecting a sample from the data and testing the amount of bias (measured as a percent) from estimating predictions on the remaining sample. Their statistic for out of sample fit $(1-\delta)$ is the percentage of bias of the model in these out-of-sample predictions.

In Table 8 I show these two statistics for various polynomial choices when the population is subdivided into sixteen groups (as in the paper) and in a smaller division that uses eight[16] groups. I also consider a version of the model that expands the J groups to 75, but note that this is estimated on a sub-sample of my data. It's clear that AIC is falling for any increase in the order of the polynomial. Using AIC fewer age groups is a better fit for short term response ($h = 0$), but worse in general for the full impulse response horizon ($h = 4$). It is not possible to estimate the $J = 8$ coefficients for higher order polynomials as the demographic variables become nearly perfectly colinear due to fitting too few age groups with too many parameters. Using more age groups would require data at a finer level than the 5-year cohorts. This is possible for data from the HMD but supplementary data used for the United States from historical censuses, as well as UN projections used in the paper come only at the level of five year age groups. The parameter of overfitting is minimized either for $k = 3$ or $k = 4$, again depending on the impulse response horizon. Because I am more interested in the longer term effects I prefer to choose polynomial degree based on the five year horizon, and thus prefer $k = 4$, but show that the estimates coming from a third order fit are fairly similar in subsection C.3 if instead the third (or indeed second) order effects are estimated.

Table 8: *AIC and Out of Sample Fit Tests: Full Sample*

IRF Horizon, $h = 0$

	Polynomial Order (k)											
	1		2		3		4		5		6	
	AIC	$1-\delta$	AIC	$1-\delta$	AIC	$1-\delta$	AIC	$1-\delta$	AIC	$1-\delta$	AIC	$1-\delta$
J=75	6896	23.06	6887	24.78	6864	26.72	6799	26.32	6712	22.62	6699	27.31
J=16	6891	22.15	6883	24.76	6874	26.64	6833	30.27	6815	31.12	6703	33.37
J=8	6898	21.89	6878	23.61	6872	24.56	6826	29.33	-	-	-	-

IRF Horizon, $h = 4$

	Polynomial Order (k)											
	1		2		3		4		5		6	
	AIC	$1-\delta$	AIC	$1-\delta$	AIC	$1-\delta$	AIC	$1-\delta$	AIC	$1-\delta$	AIC	$1-\delta$
J=75	11511	60.09	11271	42.26	11171	35.70	11008	32.74	10976	39.33	10910	45.92
J=16	11529	61.20	11261	44.00	11158	34.87	11072	32.36	10986	38.45	10914	53.54
J=8	11655	64.78	11251	45.02	11214	41.73	11103	30.68	-	-	-	-

I note that the overfitting estimates used above are not particularly suited to answer whether or not the model is overfitting the age distribution, which is a problem not of predictive power, but rather whether a few age groups are given unrealistically large weights. The procedure gives some evidence as such overfitting should lead to poor out of sample predictions, at the expense of better in-sample predictions, but such bias may be hard to detect using in-sample data particularly when population aging creates significant increases/decreases in the tails of the age distribution where polynomial weights could become quite large. A reasonable further check is to plot the age-specific coefficients for

[16] Population less than 15, greater than 75 and ten year groups between them.

Figure 11: *Full Sample Age Specific Coefficients with Varying k ($h = 4$, $J = 16$)*

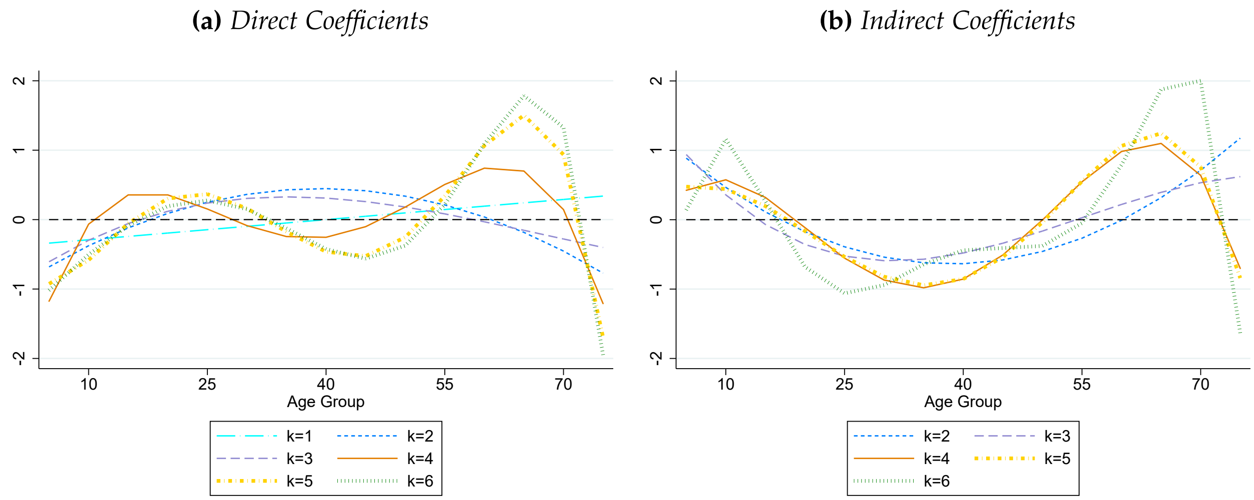

each of these polynomials, to see if individual age groups are given an outsized coefficient impact.

In Figure 11, I plot the age specific coefficients for each of these polynomial lengths using the $J = 16$ age groups from the paper. Unlike the direct coefficients in Figure 11a the indirect coefficients are not required to sum to zero. They are of course zero when excess money growth is zero, and so as in the paper are presented here with a 0.5 standard deviation change in excess money growth. To aid in comparison I demean these, which simply centers them around zero. Because the estimated impacts of these interactions are a weighted sum of the demeaned demographic shares (which sum to zero) this constant difference between these figures has no impact on their estimated impact as the more relevant information is the relative magnitude of different age groups, which can be more readily seen in this form.

Two things are worth noting about Figure 11. The first is that the general shape for both groups seems to keep the three turning points over the age distribution (first possible with $k = 4$) when increasing k to 5 and 6, with some slight added curvature in the indirect effect. While the direct effects are not the object of interest there appears to be significant added value of moving from $k = 3$ to $k = 4$. The second thing to note is that increasing k above 4 does appear to create some extreme values for the oldest age groups, potentially over-fitting the data here. This is particularly pronounced for the indirect coefficients at $k = 6$. Since increasing k to 5 has virtually no effect on the indirect coefficients, but provides worse predictive power by the Bilger and Manning (2015) measure, I prefer $k = 4$ for the paper.

To show that my coefficients of interest, the interaction term (or *indirect effect*) between demographics and money growth are not greatly impacted by radically increasing J, I also perform this exercise for the $J = 75$ case, fitting the one-year age specific demographic coefficients for each age group less than 75 years old. While the direct coefficients now looks quite different, driven mostly by some undesirable extreme values at the two tails, the interaction coefficients are quite similar to those from the $J = 16$ specification. For my baseline specification, and indeed orders of polynomial $k \geq 3$ the age coefficients in Figure 12b are quite close to those from Figure 11b.

Figure 12: *Full Sample Age Specific Coefficients with Varying k ($h = 4$, $J = 75$)*

(a) *Direct Effects* **(b)** *Indirect Effects*

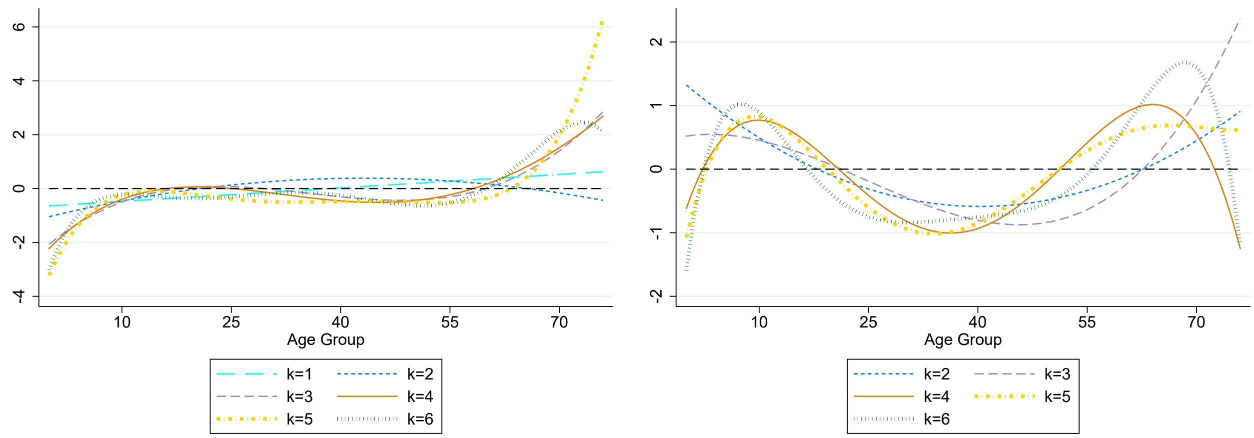

C.2. Model Choice: Divisia Data

The model run on Divisia data for the post-war United States has a much shorter sample than the rest of the estimates, and of course lacks the panel variation. As such population shares are relative to a different mean value than in the baseline estimation and coefficients are not directly comparable. Because of these differences I rerun the above to find the demographic variables most appropriate for use. Another consideration is that model power is significantly lower in this short annual sample, so all things equal a lower order choice will allow for addition of more controls to keep the exercise closer to the baseline panel estimates.

Table 9 shows the AIC and out-of-sample fit parameters for the Divisia money models. Here because of the smaller sample there are much lower AIC terms, but a similar pattern to those in the full model persists. However, out of sample fit rises for the full impulse response horizon for every polynomial order greater than two. Further, in Figure 13 I show that parameters become highly unstable for polynomial orders greater than three. Here because there is so little population variation it is likely that the model is even more sensitive to overfitting coefficients. As a result I choose $k = 2$ for my Divisia money growth regressions in the paper, though the quantitative implications of $k = 3$ are nearly identical, as should be somewhat clear from Figure 13.

C.3. Estimates With Lower Order Polynomials (K=2 & K=3)

The overfitting concern discussed above is a common issue with this methodology, and while adding additional degrees to the polynomial should help improve the fit of the age process, it may not be useful for out of sample projections. It's clear from the right hand panels of Figure 8 that the fall in < 10 and rise in > 75 age groups is having a large effect in projections to 2050.

The first thing to note is that removing the fourth order demographic term does reduce

Table 9: *AIC and Out of Sample Fit Tests: Divisia Money USA*

IRF Horizon, $h = 0$

	Polynomial Order (k)											
	1		2		3		4		5		6	
	AIC	$1-\delta$	AIC	$1-\delta$	AIC	$1-\delta$	AIC	$1-\delta$	AIC	$1-\delta$	AIC	$1-\delta$
J=16	70	47.39	62	55.32	60	63.75	43	55.51	37	69.45	10	82.12
J=8	53	49.10	45	55.55	36	63.03	27	69.27	-	-	-	-

IRF Horizon, $h = 4$

	Polynomial Order (k)											
	1		2		3		4		5		6	
	AIC	$1-\delta$	AIC	$1-\delta$	AIC	$1-\delta$	AIC	$1-\delta$	AIC	$1-\delta$	AIC	$1-\delta$
J=16	215	76.49	175	60.75	155	61.13	136	62.22	131	74.44	77	91.87
J=8	206	71.4	191	63.76	168	56.67	147	58.34	-	-	-	-

Figure 13

(a) *Demographics Direct Effects: USA Divisia* **(b)** *Demographics Indirect Effects: USA Divisia*

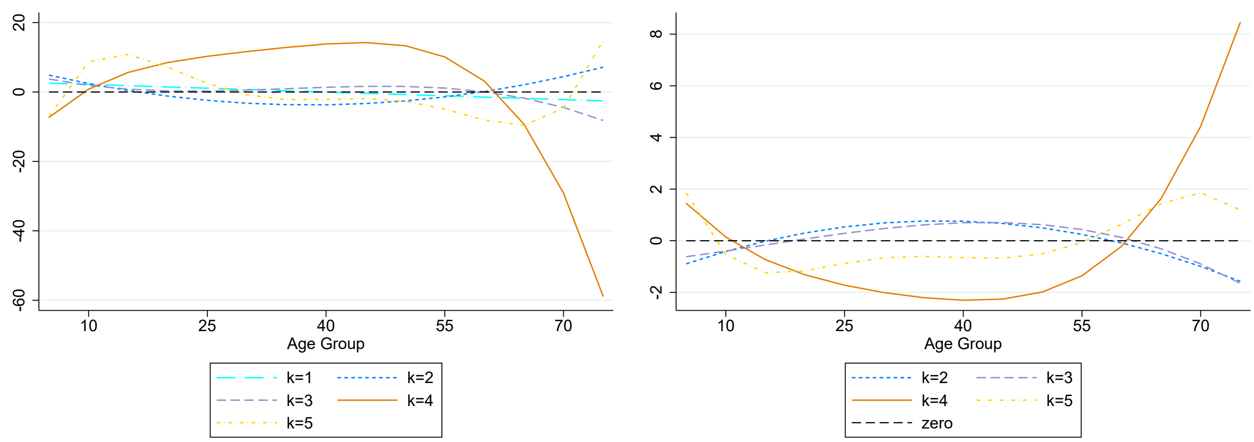

significance substantially. While many of my interaction terms remain significant,[17] it is clear that the fourth order estimation fits substantially better. This is to be expected, and even by design, as the goal was to mitigate some of the over-fitting that allowed for the oldest cohort to have such a singular effect on estimates in the post-baby boom retirement period. In addition to this, some of the weakness seems to stem from modeling an even number of turning points across age, when larger k seems to imply that an odd number across age is preferred. This can in part be seen when instead I use a second order polynomial which I include in Table 10, where again these interactions are strongly significant at every horizon. My intuition is that the best fit for the data requires oldest retirees and young dependents to have same signed direct effects, but that the third order effect struggles to accommodate this while also fitting changes over the working age. The direct effect of excess money is larger at every horizon, suggesting perhaps that the better fit of demographic data was explaining some of this direct money transmission.

To see that these changes in coefficients, though substantial, imply quite similar demo-

[17] Indeed the age structure is jointly significant at all horizons including $h = 0$ at the 5% level.

Table 10: *Inflation and Money Growth Full Age Specification: Impulse Response Function*

	Second Order Estimation					Third Order Estimation				
	0	1	2	3	4	0	1	2	3	4
m_{t-1}	0.10***	0.24***	0.42***	0.53***	0.57**	0.12***	0.30***	0.53***	0.68***	0.76***
	(0.03)	(0.08)	(0.14)	(0.18)	(0.20)	(0.03)	(0.08)	(0.14)	(0.19)	(0.20)
D1	0.01	0.09*	0.14*	0.18*	0.20	0.05	0.23**	0.45***	0.59***	0.65**
	(0.02)	(0.04)	(0.07)	(0.09)	(0.12)	(0.04)	(0.10)	(0.14)	(0.19)	(0.23)
D2	-0.01	-0.05	-0.08	-0.11	-0.12	-0.07	-0.29	-0.64**	-0.86**	-0.95**
	(0.01)	(0.03)	(0.05)	(0.07)	(0.08)	(0.06)	(0.18)	(0.28)	(0.38)	(0.44)
D3						0.03	0.10	0.26*	0.35*	0.39*
						(0.03)	(0.09)	(0.14)	(0.19)	(0.22)
$m_{t-1}\times$ D1	-0.01**	-0.03***	-0.04***	-0.05***	-0.06***	-0.02	-0.07**	-0.13**	-0.19**	-0.24***
	(0.00)	(0.01)	(0.01)	(0.02)	(0.02)	(0.01)	(0.03)	(0.05)	(0.07)	(0.07)
$m_{t-1}\times$ D2	0.01**	0.02***	0.02**	0.03**	0.04**	0.02	0.08	0.19**	0.27**	0.35***
	(0.00)	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.05)	(0.08)	(0.11)	(0.12)
$m_{t-1}\times$ D3						-0.01	-0.03	-0.07*	-0.11*	-0.14**
						(0.01)	(0.02)	(0.04)	(0.05)	(0.05)
Macro $X_{i,t}$	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Finance $X_{i,t}$	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Interact $X_{i,t}$	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
R^2	0.688	0.634	0.627	0.620	0.624	0.690	0.643	0.642	0.639	0.649
Observations	1354	1354	1354	1354	1354	1354	1354	1354	1354	1354

graphic forces, I repeat the KBO exercise from the baseline estimates in the paper using the same empirical population values. As before each economy experiences quite strong money transmission to prices in the 1970s, with substantial weakening by 1990. There are a few differences from this specification. The predictions for the United States and other countries don't look quite as different in the present year. This is due to the positive effects from late career and early retiree boomers not being as strongly offset in France and the United Kingdom, by already large groups of older retirees. Finally, while projections suggest a strengthening through of the transmission in the future the size of future projections is now muted.

The point of this exercise is to show that while aging continues to give the same qualitative picture about monetary transmission in past data, that one should be cautious about the predictions going forward with any one specification for the age distribution. However it is reassuring that changes to k generally provide similar insight into how past movements in population have affected money growth transmission.

Figure 14: *KBO Decomposition: Lower Order Polynomials*

(a) *Conditional IRF 2nd Degree: United States* **(b)** *Conditional IRF 3rd Degree: United States*

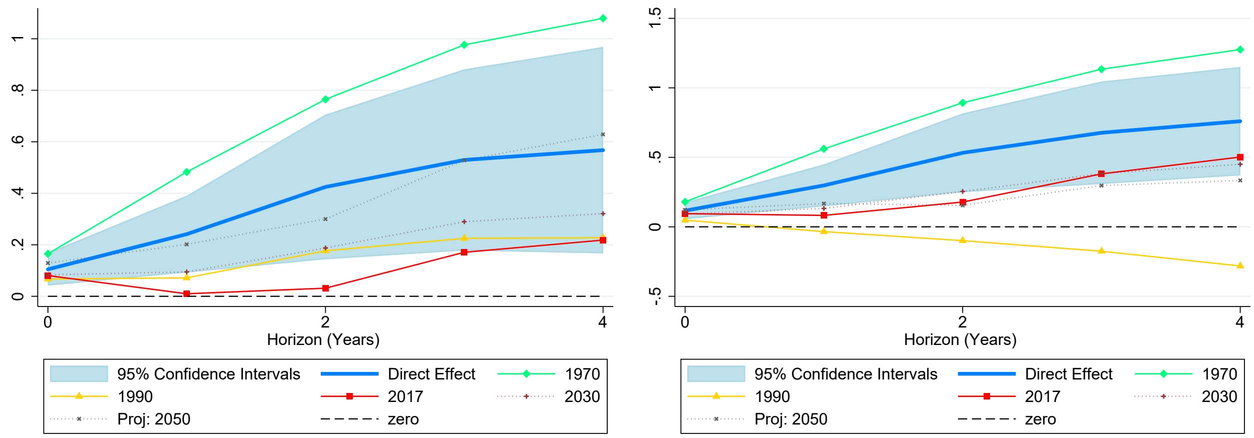

Notes: Direct effects are the average effect of money growth on inflation when population is at long run samsple mean. Conditional IRFs are effects of money growth rates conditional on age distribution of each country in a given year.

D. Estimations With Time Fixed Effects

As discussed in Section 2 I attempt to control for time fixed effects using an aggregate measure of output. This is used in a number of papers with this JST data, notably Jordà et al. (2020), as the addition of time fixed effects requires a large number of covariates. In this particular context there is also a compelling argument that there are non-uniform trends present in my panel which could in principle bias estimates if a single time trend is used. In particular, the conventional story of a structural change in the money-inflation relationship suggests that changes in the underlying monetary policy framework are important for policy pass through. Since central banks have not adopted uniform policies (and have not done so with uniform timing) across my 16 countries it would appear that time fixed effects are likely inappropriate in this context.

However, given the ubiquitous use of these time trends, I provide them in this appendix to show that their inclusion do not change any of my results qualitatively. Even though they have sizable impacts on the point estimates of the money growth-inflation relationship, decreasing them and increasing the error bands. In Table 11, I show the coefficients for both my working age population and Fair and Dominguez (1991) style polynomial controls with the inclusion of these time fixed effects. The signs of the demographic interactions are identical to those in the paper, and in the preferred specification (with three demographic controls) the effect of lagged money growth remains significant (and sizable) at all horizons. Notably the demographic interactions, though jointly significant at all levels are much weaker in this estimation than in Table 3. Here I use a third order polynomial as it appears to fit slightly better in this specification using the same criteria as in Appendix C, but qualitatively the demographic impact on the country specific KBO decomposition are the same if using fourth order estimations and demographic terms are strongly jointly significant in second and fourth order specifications at all horizons.

Figure 15 plots the same KBO decomposition for the three countries as in the main text of the paper. As before the dramatic shift from 1970 to 1990 is present. For working age population it will take longer for declines to reverse this trend, while the point estimates for 2017 over the full age distribution are similar to those in the main text, though slightly smaller for the United States.

Table 11: *Inflation, Money Growth, and WAP: Impulse Response Horizons*

	WAP Control: Horizon					D1-D3 Controls: Horizon				
	0	1	2	3	4	0	1	2	3	4
m_{t-1}	0.04^{*}	0.09	0.16^{*}	0.16	0.20	0.05^{**}	0.13^{*}	0.27^{**}	0.32^{**}	0.34^{**}
	(0.02)	(0.05)	(0.08)	(0.11)	(0.12)	(0.02)	(0.06)	(0.10)	(0.13)	(0.12)
dmWAP	0.11^{**}	0.47^{***}	0.80^{***}	0.95^{***}	0.87^{***}					
	(0.05)	(0.14)	(0.25)	(0.30)	(0.28)					
$m_{t-1}\times$ dmWAP	-0.02^{*}	-0.09^{***}	-0.15^{***}	-0.20^{***}	-0.21^{***}					
	(0.01)	(0.02)	(0.04)	(0.05)	(0.05)					
D1						0.01	0.14	0.03	-0.15	-0.38
						(0.10)	(0.25)	(0.38)	(0.56)	(0.71)
D2						0.00	-0.14	0.05	0.35	0.72
						(0.16)	(0.41)	(0.64)	(0.94)	(1.20)
D3						-0.01	0.04	-0.04	-0.18	-0.34
						(0.07)	(0.18)	(0.29)	(0.42)	(0.54)
$m_{t-1}\times$ D1						-0.01	-0.06^{*}	-0.09^{**}	-0.13^{**}	-0.15^{**}
						(0.01)	(0.03)	(0.04)	(0.06)	(0.06)
$m_{t-1}\times$ D2						0.02	0.08	0.14^{*}	0.18^{*}	0.21^{**}
						(0.02)	(0.05)	(0.07)	(0.09)	(0.09)
$m_{t-1}\times$ D3						-0.01	-0.03	-0.06^{*}	-0.08^{*}	-0.08^{**}
						(0.01)	(0.02)	(0.03)	(0.04)	(0.04)
Macro $X_{i,t}$	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Finance $X_{i,t}$	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Interact $X_{i,t}$	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Time FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Country FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
R^2	0.693	0.644	0.673	0.676	0.665	0.827	0.799	0.807	0.824	0.851
Observations	1354	1354	1354	1354	1354	1354	1354	1354	1354	1354

Figure 15: *KBO Decomposition: Including Time Fixed Effects*

(a) *Conditional IRF, WAP: United States* **(b)** *Conditional IRF, Full Age Distribution: United States*

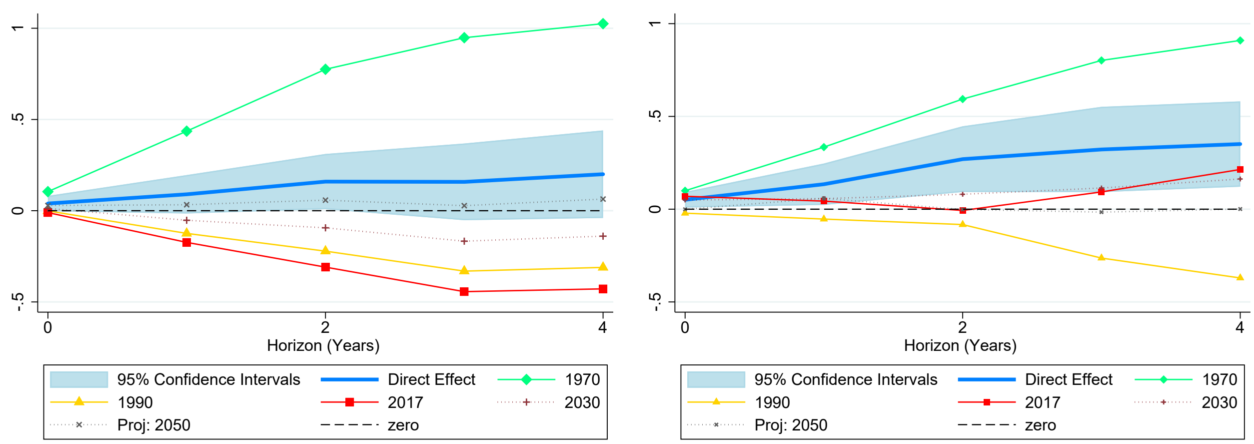

Notes: Direct effects are the average effect of money growth on inflation when population is at long run sample mean. Conditional IRFs are effects of money growth rates conditional on age distribution of each country in a given year.