

Aggregate Implications of Wealth Redistribution: The Case of Inflation *

Matthias Doepke
UCLA

Martin Schneider
NYU and Federal Reserve Bank of Minneapolis

Abstract

This paper shows that a zero-sum redistribution of wealth within a country can have persistent aggregate effects. Motivated by the case of an unanticipated inflation episode, we consider redistribution shocks that shift resources from old to young households. Aggregate effects arise because there are *asymmetries* in the reaction of winners and losers to changes in wealth. We focus on two sources of asymmetries: differences in the average age of winners and losers, and differences in their labor force status.

JEL: D31, D58, E31, E50.

1 Introduction

Economic shocks or changes in government policy often lead to a substantial redistribution of wealth within the population of a country. An important example of a redistribution shock is an unanticipated inflation episode. When the price level rises faster than previously expected, the real value of all nominal assets declines. As a result, borrowers win—they experience an increase in real wealth—while lenders lose. As documented in Doepke and Schneider (2005a), the amount of this wealth redistribution can be large even for moderate inflation shocks.

In this paper, we assess whether a redistribution of wealth within the population of a country affects key aggregate variables such as output, labor supply,

*Acknowledgements: We have benefited from comments received at the EEA Annual Congress in Amsterdam. Financial support from the National Science Foundation is gratefully acknowledged. Doepke is a CEPR Research Affiliate. Doepke: Department of Economics, University of California, Los Angeles, 405 Hilgard Ave, Los Angeles, CA 90095. Schneider: Research Department, Federal Reserve Bank of Minneapolis, 90 Hennepin Avenue, Minneapolis, MN 55480. E-mail addresses: Doepke <doepke@econ.ucla.edu>, Schneider <ms1927@nyu.edu>. The views expressed herein are those of the authors and not necessarily those of the Federal Reserve Bank of Minneapolis or the Federal Reserve System.

and investment. We provide a simple theoretical framework to show that zero-sum redistribution shocks can have persistent aggregate effects. We then draw on results in Doepke and Schneider (2005a, 2005b) to illustrate the aggregate and welfare effects of redistribution triggered by an inflation episode.

Our results follow from two basic observations. First, zero-sum redistribution shocks have aggregate effects if, and only if, there are *asymmetries* in the reactions of winners and losers to changes in wealth. If the individual response of, say, labor supply to a marginal change in wealth is the same for everybody, then wealth changes that sum to zero generate changes in labor supply that also sum to zero. In contrast, if one group of households (such as those of working age) adjusts labor supply more strongly than another group (such as retirees), then individual responses need not sum to zero, and aggregate labor supply can react to a zero-sum redistribution of wealth.

Second, households spread any changes in individual wealth brought about by redistribution over their lifetimes in order to smooth consumption. Redistribution shocks thus have persistent effects at the individual level. How this persistence translates to the aggregate level depends on how the underlying asymmetry that triggers the aggregate effect is related to age. If the initial response of an aggregate is driven by the individual responses of young households, further moves by the aggregate in the same direction tend to occur. In contrast, if the initial effect is driven by the elderly, it is likely to be reversed, leading to an oscillatory impulse response.

The focus in this paper is on redistribution shocks that shift resources from old to young households. This pattern fits our example of an inflation episode, in which young, working-age borrowers gain at the expense of old, often retired lenders. If retirees adjust labor supply less than workers, such shocks increase aggregate labor supply. Moreover, consumption smoothing induces an increase in aggregate savings and investment. Indeed, while everybody spreads their gain or loss over the rest of their lives, the young have many more periods of life left than the old. They therefore save more, whereas the old absorb a larger fraction of their loss immediately in terms of lower consumption. The increase in savings by the young is thus not offset by the dissaving of the old.

The degree of persistence depends on whether the impact effect on an aggregate variable is driven by the old or the young. For any aggregate variable that moves on impact, a group of households reacts in the same direction at the individual level—the drivers of the aggregate response—while a second group—the mitigators—moves in the opposite direction. If the typical driver is older than the typical mitigator, then, as the drivers die, the influence of the mitigators on the aggregate will grow over time, and the initial aggregate response will be reversed.

In contrast, if the drivers are on average younger than the mitigators, then their effect on aggregates grows as the mitigators die. The response is then characterized by momentum, perhaps even a hump-shaped impulse response, rather than reversal. In the case of an inflation shock, this is what happens to capital and labor supply.

The two dimensions of heterogeneity we emphasize—age and status in the labor force—would lead to asymmetric responses to any redistribution shock that discriminates by age. Our results therefore also apply to the analysis of policy changes such as pension reform, which induces age-dependent redistribution as well. More generally, other patterns of redistribution may combine with other forms of heterogeneity to generate the asymmetric responses required for aggregate effects. An exploration of these issues is left to future research on the macroeconomic implications of redistribution shocks.

In the next section, we present a simple overlapping generations model in which the labor productivity of a worker is linked to age. Section 3 derives the response to redistribution from the old to the young in this setup. Section 4 reviews our quantitative results on the redistribution effects of inflation.

2 A Theoretical Framework

Our theoretical framework emphasizes two dimensions of heterogeneity across people: age and labor productivity. These two dimensions are linked: a common life-cycle efficiency profile of labor supply is shared by all agents. People are identical along all other dimensions.

At any date t , there are $n + 1$ different consumers, identified by their age $a = 0, \dots, n$. Consumers care about leisure l as well as consumption of a numeraire good c . Preferences over future consumption streams for a consumer of age a as of time t are represented by the utility function

$$\sum_{\tau=0}^{n-a} \beta^\tau [\alpha \log c_{t+\tau}^{a+\tau} + (1 - \alpha) \log l_{t+\tau}^{a+\tau}],$$

where subscripts denote the current period, superscripts denote age, and $\beta < 1$. Every consumer is endowed with \bar{l} units of time per period. Consumers divide their time between work and leisure: per unit of time, ϕ^a units of effective labor can be sold in a competitive labor market at the wage rate w_t .

The consumer uses current wealth A_t^a (which includes the value of the time endowment) to buy consumption goods c_t^a , to consume leisure l_t^a , and to save by buying assets s_t^a that pay the gross interest rate R_t . He thus faces a sequence of budget constraints of the form

$$c_t^a + s_t^a + w_t \phi^a l_t^a = A_t^a := R_t s_{t-1}^{a-1} + w_t \phi^a \bar{l}.$$

We focus on aggregate effects of redistribution that derive from consumer behavior. We thus choose a production technology that implies constant factor prices: competitive firms produce the numeraire good using labor and capital according to the production function

$$Y_t = RK_{t-1} + wL_t,$$

where w and $R = \beta^{-1}$ are fixed parameters. The capital stock fully depreciates every period.

An equilibrium of our economy consists of consumption and leisure allocations, as well as wage and interest rates, such that (i) consumers and firms optimize, given prices, and (ii) markets clear, that is:

$$\sum_a c_t^a + \sum_a s_t^a = Y_t; \quad L_t = \sum_a \phi^a (\bar{l} - l_t^a); \quad K_t = \sum_a s_t^a.$$

Since technology is linear, equilibrium prices are determined by the supply side alone: the interest and wage rates must be $R_t = R$ and $w_t = w$, respectively. It follows that quantities are entirely determined by consumer behavior. In particular, output depends on *current* aggregate labor supply and *last period's* aggregate savings. We assume throughout that parameters are such that these quantities are nonnegative.

The consumer's problem in our economy has a closed-form solution. Let W_t^a denote a consumer's total lifetime wealth at time t . It consists of current wealth A_t^a as well as the present value of future labor endowments:

$$W_t^a = A_t^a + \sum_{\tau=1}^{n-a} R^{-\tau} (w\bar{l}\phi^{a+\tau}).$$

We start out by assuming that the optimal choice for leisure is interior, i.e., labor supply is positive in every period. This is necessarily true if the lifetime skill profile is flat, which is the main case considered below. Given Cobb-Douglas preferences, and using $\beta = R^{-1}$, numeraire consumption as a function of wealth can then be written as

$$c_t^a = \alpha \frac{1 - \beta}{1 - \beta^{n-a+1}} W_t^a =: \alpha \mu^a W_t^a.$$

The marginal propensity to consume out of wealth, μ^a , rises with age. This reflects a taste for consumption smoothing over the life cycle. Since old people spread any extra dollar of wealth over fewer periods than young people, their consumption reacts more strongly to changes in wealth. The same is true for

leisure. However, the marginal propensity to consume leisure depends on age also through the skill profile: agents consume less leisure when they are more skilled. The optimal choice for leisure is given by

$$l_t^a = \frac{1 - \alpha}{w\phi^a} \mu^a W_t^a.$$

3 Aggregate Effects of Redistribution Shocks

Suppose the economy is in some initial equilibrium. We are interested in the aggregate effects of a one-time redistribution of wealth. Let $\hat{W}_t^a = \hat{A}_t^a$ denote the transfer to a household of age a . All transfers take place immediately, so that current wealth and lifetime wealth change by the same amount. We also assume that the redistribution is zero-sum among agents alive at t , that is: $\sum_a \hat{W}_t^a = 0$. Since individual consumption and leisure are linear in wealth and technology is also linear, we can directly study the deviations of all aggregates from their values in the initial equilibrium without having to characterize the initial equilibrium itself.

Aggregate effects arise in our framework because of asymmetries in the behavior of winners and losers of redistribution. As a benchmark *without* any asymmetries, consider the case of infinitely-lived agents ($n = \infty$) with constant labor productivity ($\phi^a = \phi$). In this setting, the marginal propensity to consume is equal to $1 - \beta$ for all agents. As a result, the effect of a zero-sum redistribution on aggregate consumption, aggregate leisure, and hence all other aggregate variables of interest is zero.

To focus on life-cycle effects, we assume for now that the skill profile is flat, that is, $\phi^a = \phi$. We also restrict attention to redistribution vectors that shift resources from the old to the young: we assume that there exists a threshold age a^* such that $\hat{W}_t^a > 0$ for all $a < a^*$ and $\hat{W}_t^a < 0$ for all $a \geq a^*$.

Proposition 1 *With a flat skill profile, a redistribution from old to young decreases aggregate consumption and increases savings, aggregate labor supply, and output on impact relative to the benchmark equilibrium.*

Proof The changes in aggregate consumption and leisure are both proportional to $\sum_a \mu^a \hat{W}_t^a$, which is negative because

$$\sum_{a < a^*} \mu^a \hat{W}_t^a < \mu^{a^*} \sum_{a < a^*} \hat{W}_t^a = -\mu^{a^*} \sum_{a \geq a^*} \hat{W}_t^a < -\sum_{a \geq a^*} \mu^a \hat{W}_t^a,$$

where the inequalities follow from the definition of a^* and the fact that μ^a is increasing in a , and the equality holds because the redistribution is zero-sum.

By a similar argument, the change in savings, which is proportional to $\sum_a (1 - \mu^a) \hat{W}_t^a$, is positive, as the savings rate $1 - \mu^a$ is decreasing in age. The change in aggregate labor supply is minus the change in total leisure and is therefore positive. Since capital is fixed in the impact period, output moves in the same direction as aggregate labor supply, and thus increases as well. \square

The only mechanism at work here is consumption smoothing: the decrease in consumption of the old is larger than the increase in consumption of the young, so that aggregate consumption falls and more capital is accumulated.

The next proposition concerns aggregate effects beyond the impact period.

Proposition 2 *With a flat skill profile, at all dates $\tau \in [t + n - a^*, t + n]$ labor is below the benchmark, while capital is still above the benchmark. For $\tau > t + n$, aggregate deviations are zero.*

Proof Aggregates at some date $t + j$ depend only on the actions of agents who are alive at $t + j$. Since prices are fixed, an agent's actions deviate from the benchmark equilibrium only if his wealth deviates. The vector of wealth deviations evolves over time according to

$$\hat{W}_{t+j}^a = \begin{cases} R(1 - \mu^{a-1}) \hat{W}_{t+j-1}^{a-1} & \text{if } a \geq j \\ 0 & \text{if } a < j \end{cases}.$$

Because of consumption smoothing, an individual's wealth deviation does not change signs over the course of the lifetime. For example, since the youngest cohort receives a positive transfer \hat{W}_t^0 at date t , its wealth deviation \hat{W}_{t+j}^j remains positive at all ages $j = 1, \dots, n$. Moreover, only agents who were already alive when the redistribution shock hit at date t are still affected by the shock at $t + j$.

By date $t + n - a^* + 1$, all cohorts who initially received negative transfers have died. As a result, all the terms in $\sum_a \mu^a \hat{W}_{t+n-a^*+1}^a$ are positive. The labor deviation is thus negative. Similarly, the capital deviation $\sum_a (1 - \mu^a) \hat{W}_{t+n-a^*+1}^a$ is positive. The same is true for all later dates up until $t + n$. However, by $t + n + 1$, the last cohort alive during the redistribution shock period has died, and aggregates return to the benchmark. \square

The proposition clarifies the link between the underlying asymmetry that triggers the initial effect of a redistribution shock and the nature of the impulse response to the shock. The initial effect on some aggregate variable arises because, at the individual level, either the old or the young cohorts react more strongly to the shock. If the impact effect is due to a stronger response by the old, then it must be reversed before the economy returns to the benchmark. This is the case

for labor, which increases on impact since the old increase their labor supply by more than the young decrease theirs. After the impact, aggregate labor supply decreases monotonically until, in period $t + n - \bar{a}^* + 1$, all losers from the redistribution have died. Subsequently, aggregate labor supply increases again and approaches the benchmark equilibrium monotonically from below. In contrast, if the impact effect reflects a stronger response by the young, the impulse response will feature momentum in the aggregate variable. This is what happens to capital, here given by the sum of savings out of total wealth. The deviation in the capital stock remains positive until, in period $t + n + 1$, all aggregates return to the benchmark equilibrium.

So far, we have considered life-cycle effects with a flat skill profile. We now examine a particular source of variation in skills: exogenous retirement. For simplicity, we assume that the retirement age coincides with the threshold age \bar{a}^* at which the redistribution transfer becomes negative:

$$\phi^a = \begin{cases} \phi & \text{for } 0 \leq a < \bar{a}^* \\ 0 & \text{for } a \geq \bar{a}^* \end{cases}$$

The redistribution now no longer shifts resources just from old to young cohorts, but also from retired to working cohorts.

Proposition 3 *With exogenous retirement, labor supply falls on impact and remains below the benchmark until it returns to the benchmark in period $t + \bar{a}^*$.*

Proof Since retired agents have no marketable skills, their labor supply is zero both before and after the shock. Any deviations of labor supply from the benchmark are therefore generated by agents younger than \bar{a}^* . All these agents receive positive transfers and work less, so that the deviation is negative both on impact and in all subsequent periods. It is zero once all the agents alive at the time of the shock have retired. \square

With exogenous retirement, the impact effect on labor supply arises because of a stronger response by the young; in fact, it arises because the old do not respond at all. As a result, the impulse response now exhibits momentum, not reversal. The impulse response in capital, in contrast, is still driven by the standard life-cycle effect and retains its previous shape.

4 Aggregate Effects of Inflation

The simple model presented so far has shown that redistribution shocks can have persistent effects on economic aggregates. At the same time, the analysis has abstracted from a number of features that are likely to affect at least the magnitude

of the effects. In this section, we summarize the findings in Doepke and Schneider (2005a, 2005b), where we quantify the potential real effects of a moderate inflation shock in the United States. While the empirical model used in these papers is much richer than the simple framework used here, the main findings derive from the life-cycle and retirement effects emphasized above.

In Doepke and Schneider (2005a), we document the distribution of nominal positions in the United States, combining data from the Survey of Consumer Finances (SCF) and the Flow of Funds Accounts (FFA). We sort nominal instruments into several broad asset classes and then construct, for every household and asset class, a certain *payment stream* that the household expects to receive from the instruments it owns up to maturity. We value the payment stream using the nominal term structure to determine the household *net nominal position* (NNP), defined as the market value of nominal assets less nominal liabilities. Here we include not only direct nominal asset holdings and debt, but also nominal assets held indirectly (such as ownership of shares in a mutual fund that holds nominal bonds) and debt owed indirectly (for example, through ownership of a business that in turn has issued nominal debt).

We calculate individual present value gains and losses from a moderate inflation episode by revaluing each nominal payment stream. Of course, the magnitude of the redistribution depends on how quickly expectations and portfolios adjust during the episode. To address this issue, we posit two scenarios that provide upper and lower bounds for the redistribution effect of inflation. The resulting interval estimates for present value gains and losses are sizeable. We focus on a moderate inflation episode—about 5% extra inflation for ten years—that begins (hypothetically) in 1989. A coalition of relatively old households loses between 7% and 18% of GDP. About three-quarters of the loss are borne by the top 10% of the wealth distribution, who hold a lot of long-term fixed-income securities. However, poor agents who hold most of their savings as deposits are also vulnerable to inflation. Within each wealth category, the largest losses are borne by the oldest households, who are already in retirement. Within the household sector, the main winners are young middle-class households who bought a home and have large fixed-rate mortgages. About half of the total gains in the redistribution accrue to this group. Along the age dimension, the flow of the transfers is therefore broadly similar to the old-to-young distribution discussed in the present paper.

The remainder of the gains go to the government, which also benefits at the expense of the foreign sector; the revaluation of government debt yields between 5% and 14% of GDP. The presence of foreigners is an important difference between the closed economy model of the present paper and the calculations in Doepke and Schneider (2005a, 2005b). Given the growing negative net nominal

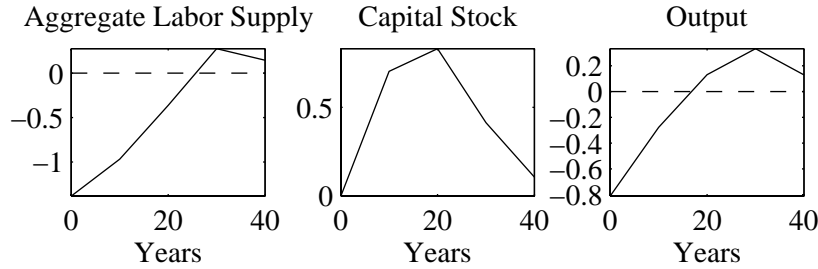


Figure 1: Impact on Aggregates (Percent Deviation from Balanced Growth Path)

position of the United States, redistribution due to inflation is currently not zero-sum among Americans, but rather presents domestic households with a windfall gain. The quantitative analysis in Doepke and Schneider (2005a, 2005b) takes this windfall into account. Also accounted for are the effects of fiscal policy on future generations.

In Doepke and Schneider (2005b), we assume $n = 5$; agents work for four decades and then are retired for another two decades. The model features a standard production technology, allowing for endogenous price movements in response to a redistribution shock. Agents are divided into types with different preferences and skill profiles, to capture wealth heterogeneity observed in the data. Preferences accommodate a taste for durables—modeled as private capital stocks—as well as for bequests. The model is calibrated to cross-sectional data on wealth and earnings from the SCF, as well as to aggregates from the National Income and Product Accounts.

We compute the balanced-growth-path wealth distribution of the model and then shock it with redistribution vectors. We construct the latter by combining the results of the inflation experiment with an assumption about fiscal policy: we have to take a stand on the path of debt as well as taxes, transfers, and spending. Under the benchmark scenario, the government gradually increases the debt/GDP ratio back to steady state and adjusts spending along the way. The resulting impulse responses for selected aggregates are presented in Figure 1. The figure shows results for the upper-bound scenario; results for the lower-bound scenario are qualitatively similar but quantitatively smaller, with an initial change in output about one-third as large as in the upper-bound scenario.

Consider, first, the reaction of aggregate labor supply. In the impact period, we observe a decline in aggregate labor supply by close to 1.5%. As in the theoretical analysis above, this effect is driven by retirement: most losers in the redistribution are past retirement age and therefore do not adjust their labor supply. The young winners, in contrast, consume more leisure, which accounts for

the negative aggregate effect on labor supply. The reaction of the aggregate capital stock closely follows the predictions of the life-cycle model in Section 3. In the impact period, the capital stock is predetermined through savings in the preceding period. Subsequently, the capital stock rises for two decades and then converges back to the balanced growth path. As in the theoretical model presented above, the capital stock is driven by life-cycle effects. The young winners from redistribution increase their savings substantially to smooth consumption over their remaining lives. The old losers, in contrast, are already close to the end of the life cycle and are forced to take most losses in terms of lower consumption, with only a small impact on savings. The capital stock starts to decrease after a number of decades as the young winners from redistribution reach retirement and start to dissave. The overall impact on output is initially driven by labor supply only, resulting in a fall in output of about 0.8% over the first decade. After two decades, the increased capital stock starts to dominate, leading to an increase in output of up to 0.25% relative to the balanced growth path.

The magnitude of these aggregate effects of inflation (which arise from wealth redistribution alone) is similar to what is generated by representative-agent models with standard monetary frictions. The results also demonstrate that aggregate effects arising from wealth redistribution are extremely persistent: they fade out only as the cohorts affected by the redistribution shock die, and therefore persist for decades after the end of the inflation episode.

References

- Doepke, Matthias and Martin Schneider. 2005a. "Inflation and the Redistribution of Nominal Wealth." Unpublished Manuscript, UCLA and NYU.
- . 2005b. "Redistribution of Nominal Wealth and the Welfare Cost of Inflation." Unpublished Manuscript, UCLA and NYU.