Discussion of
"'Inflation's Impact on American Households'
by David Altig, Alan Auerbach, Erin Eidschun, Laurence Kotlikoff, Victor Yifan Ye"

Deborah Lucas
Massachusetts Institute of Technology and NBER

July 17, 2024

Prepared for the 2024 NBER Macroeconomics Annual which was held in Cambridge, MA.

A surge of high and now lingering post-pandemic inflation has rekindled concerns about its costs and distributional impacts. In this timely and ambitious analysis, Altig et. al. explore imperfectly indexed fiscal policy as a potentially important channel for inflation to have real effects. While economists have long recognized the potential for distortions to arise from the interaction of inflation with fiscal rules, this analysis represents a significant advance over past studies along several dimensions: (1) The analysis incorporates most major tax and benefit policies at both the federal and state levels; (2) costs are measured in terms of the percentage consumption change over a household's remaining lifetime rather than at a point in time; and (3) the results are broken out by age and wealth quintile, illuminating distributional effects.

The main hypothetical that the authors examine is, what would happen to discretionary expenditures if inflation were to permanently rise to 10 percent annually, all else equal? That is, what happens if the paths of real wages, real asset returns, current law policy, etc., are held constant. The answer is that on average taxes net of benefits would rise, leaving households in aggregate with lower lifetime consumption expenditures. Notably, the effects are heterogeneous across wealth and age groups, with a few big winners as well as some big losers. The effects are progressive, with the biggest losses borne by very wealthy households, and with significant gains for some poorer households. An important takeaway is that the current fiscal system is far from inflation-neutral, and that modifying policies to correct that could avoid unintended consequences arising from future inflationary episodes.

In the discussion that follows, I briefly recap the modeling approach in order to highlight some of its unconventional features that may cause the magnitude of estimated expenditure reductions to be larger than what would be predicted by a more standard optimizing model. I then focus on the question of whether the finding of progressivity is robust, and suggest several factors that could reverse or significantly weaken that conclusion. Turning to aggregate and policy implications, I consider whether average lifetime spending would still be predicted to fall if an explicit government budget constraint had been incorporated. I conclude with a short discussion of why limited indexation of taxes, although redistributive, might be a second-best policy.

1. Modeling approach

The major advance of this analysis over previous efforts to quantify the effects of inflation via fiscal channels was made possible by use of The Fiscal Analyzer (TFA), an existing tool for retirement planning that several of the authors developed and commercialized, and that has been used in previous academic papers. Most critically for this analysis, the model's code incorporates the rules determining taxes and benefits for hundreds of tax and benefit programs, which in turn affect each individual households' path of earnings and asset holdings over its lifetime. The TFA is a partial equilibrium model that assumes households will smooth deterministic expenditure paths over their remaining lifetime to the maximum extent possible, subject to borrowing and wealth constraints and using actuarially fair life insurance to stabilize the consumption paths of surviving spouses. Any bequests are accidental. The cross-section of households represented and imputations about their characteristics and program participation is based on demographic and economic data from the 2019 SCF and other government data sources.

The authors emphasize several mechanisms that are important for the results. The interactions of inflation and fiscal policies that drive changes in real household expenditures are primarily due to lags in, or the absence of, indexation related to the computation of taxes and benefits; and from the incomplete indexing of thresholds and ceilings that alter benefit program eligibility. The lower wealth households that benefit from higher inflation do so primarily because of expanded eligibility for valuable Medicare and other social insurance benefits. The wealthier households that bear the largest costs are hit with much higher effective tax rates on nominal interest receipts and capital gains.

While the qualitative effects of these driving mechanisms seems intuitively robust, the predicted magnitudes are likely to be quite sensitive to specific modeling choices. The framework used in this paper differs fundamentally from those of standard lifecycle models. The pattern of expenditures and asset allocations are not the solution to a constrained inter-temporal optimization problem. There is no choice about when to retire, how many hours to work, or the size of bequests. There is no possibility for firms to change the structure of their liabilities to avoid the additional taxation of nominal interest or capital gains (e.g., by replacing bonds with dividend-paying stocks). Household portfolio allocations do not respond to rebalancing incentives, for instance, more care taken to hold bonds only in tax-protected retirement accounts.

While the simplifying assumptions in the TFA may be necessary to accommodate the complexity introduced by incorporating realistic fiscal rules, those assumptions tend to all go in the direction of precluding likely behavioral changes that would mitigate reductions in expenditures. For example, people might reduce intentional bequests, work more hours or retire later, or rebalance their portfolios towards assets that are less heavily taxes. This suggests that the reported magnitudes of expenditure reductions, both for different groups of households and in aggregate, are upper bounds rather than expected values. For the households that potentially benefit from inflation's interaction with fiscal policies, behavioral changes would tend to increase those benefits.

It also remains an open question of whether the behavior implied by the TFA model, which was designed to be prescriptive rather than descriptive, is realistic overall and for different wealth and age cohorts. It would be useful to compare the implied expenditure paths in the model with observed patterns for different groups to verify that the assumed behavior is plausible. More ambitious but also very worthwhile, the fiscal apparatus could be incorporated into a more traditional optimizing model and the results compared.

The focus on inflation permanently moving to a rate of 10 percent (or 5 percent) demonstrates that large distortions via fiscal effects are theoretically possible. However, because those scenarios seem unlikely to occur, it would be interesting to also consider the effects of a permanent 3 percent inflation rate. Extrapolating from the reported results for 5 percent inflation, I would expect those effects to be quite small. Nevertheless, that exercise would be interesting because it is germane to the current conversation about the pros and cons of moving to a 3 percent inflation target.

2. Are the findings of progressivity robust?

The authors provocatively claim that inflation is a highly progressive form of net taxation, albeit only in the context of its interactions with the fiscal system. For instance, they note that the 15% median lifetime spending loss for the top 1% is roughly twice that for the bottom quintile (in the 10% inflation scenario). I argue here that there are important omitted considerations, which if taken into account, would weaken or even reverse the conclusion that the effects of inflation via the fiscal system are progressive.

An important omission from the model is the insurance value of benefit programs, particularly for low-income households. Such programs provide valuable protection against catastrophically low consumption realizations. In interpreting their findings, the authors implicitly take average expenditure changes as an adequate proxy for welfare changes. However, for programs with an important insurance component, using average expenditures as a measure of welfare effects can be misleading.

A simple model illustrates the potential magnitude of the omitted costs associated with reductions in social insurance protection caused by higher inflation, for instance because higher nominal income causes a household to lose program eligibility. Consider a low-income household that earns relatively low wages most of the time. Unemployment or health shocks cause it to periodically qualify for government benefit programs. Consumption is hand-to-mouth. Period utility exhibits constant relative risk aversion with curvature parameter γ:

$$U(c) = \frac{c^{(1-\gamma)}}{(1-\gamma)}$$

Income when employed is $35,000 annually, the top of the bottom quartile of income. Own income when unemployed is $5,000. Government benefits when unemployed are $10,000. Hence, total income when unemployed is $15,000 annually.

Expected utility in any period, either (1) with social insurance; or (2) without social insurance is given by:

$$p\ U(35{,}000) + (1\text{-}p)\ U(15{,}000) \quad (1)$$

$$p\ U(35{,}000) + (1\text{-}p)\ U(5{,}000) \quad (2)$$

It is straightforward to calculate how much income the household would willingly give up whenever it is in an employed state for benefit program eligibility when unemployed. The willingness to pay, X, solves:

$$p\ U(35{,}000 - X) + (1\text{-}p)\ U(15{,}000) = p\ U(35{,}000) + (1\text{-}p)\ U(5{,}000) \quad (3)$$

Table 1 shows the willingness to pay for a range of assumptions about risk aversion and the probability of unemployment, and compares the expected premium payment based on willingness to pay with the actuarially fair premium for the social insurance. Willingness to pay far exceeds actuarial cost even with moderate risk aversion. For example, for risk aversion of 1.5

and a 10% chance of being unemployed, the actuarially fair value of the premium is .1 x $10,000 = $1,000. The willingness to pay is found to be $7,300. It is only paid in the employed state, making its expected value .9 x $7,300 = $6,576. The value of the insurance to the household therefore is more than 6.5 times its actuarially fair value. That large multiple underscores that the value of social insurance, particularly for low-wealth recipients facing the risk of very low income, can be many times greater than the expected value of benefits received.

Table 1. Social Insurance: willingness to pay versus actuarial value

- 90% chance employed; γ = 0, .5 or 1.5

Probability employed	0.9
employed income	35000
unemployed income	5000
unemployed income +benefits	15000
CRRA	0
Expected benefit	1000
Voluntary employed premium	1111
Voluntary expected premium	1000

Probability employed	0.9
employed income	35000
unemployed income	5000
unemployed income +benefits	15000
CRRA	0.5
Expected benefit	1000
Voluntary employed premium	2119
Voluntary expected premium	1907

Probability employed	0.9
employed income	35000
unemployed income	5000
unemployed income +benefits	15000
CRRA	1.5
Expected benefit	1000
Voluntary employed premium	7306
Voluntary expected premium	6576

- 50% chance employed, γ = 0, .5 or 1.5

Probability employed	0.5
employed income	35000
unemployed income	5000
unemployed income +benefits	15000
CRRA	0
Expected benefit	5000
Voluntary employed premium	10000
Voluntary expected premium	5000

Probability employed	0.5
employed income	35000
unemployed income	5000
unemployed income +benefits	15000
CRRA	0.5
Expected benefit	5000
Voluntary employed premium	16689
Voluntary expected premium	8344

Probability employed	0.5
employed income	35000
unemployed income	5000
unemployed income +benefits	15000
CRRA	1.5
Expected benefit	5000
Voluntary employed premium	27200
Voluntary expected premium	13600

As the above example illustrates, the economic value of government insurance benefits is likely to greatly exceed the expected payments received for low-income households exposed to catastrophic income risk. The economic value of government insurance programs for young and middle-income households of working age is also likely to exceed the value of expected payouts, even though take-up is rarer than for poor households. By contrast, benefit programs aimed at older households generally have less of an insurance element, and estimates of value based on average benefit receipts is more appropriate. For example, social security payments and Medicare health insurance subsidies provide a level stream of resources rather than episodic cash infusions. For wealthier households, consumption floors provided by government programs are too low and too unlikely to bind to provide significant insurance value.

Insurance benefits are hard to directly quantify using the TFA. The model abstracts from the cost of risk, and the assumption of maximal expenditure smoothing would be likely to cause the estimated benefit of government insurance to be lower than if the limited smoothing by many low-income households were taken into account. A partial remedy could be via discount rate adjustments. Specifically, value flows could be discounted at exogenously determined discount rates that are risk-adjusted. Notable, risk-adjustment could imply negative discount rates for benefits from programs with high insurance value. The authors currently discount all value streams at a single rate based on average real investment returns in the U.S. economy. A move to risk-adjusted discounting would also be helpful for more accurately capturing the costs of tax distortions on different categories of investment income.

Another source of omitted regressivity is that the asset limits in many anti-poverty programs are not indexed for inflation. In fact, eligibility for low-income benefits is subject to severe asset limits for programs like SSI (elderly and disabled poor), Medicaid (low-income health), TANF (welfare), and SNAP (food stamps). Typical limits ae only $2,000 to $3,000 for individuals. Limits vary by state and can be complicated.[1] For example, in North Dakota, a one-person household receives $3,000, two people receive $6,000, and another $25 is allowed for each additional person thereafter. Unlike the one or two-year inflation adjustment lags in most major programs, the asset limit for SSI was last raised in 1989. The price level has increased by factor of 2.5 since that time. What counts as an asset varies by state, but often cash, bank accounts, cars, retirement accounts are included. The affected population is large. In 2021 there were 41 million SNAP participants, and in 2022 there were 7.6 million SSI participants, including a million children with disabilities. Many economists have emphasized the high economic costs of the limits, which tend to lock people into poverty. For example, without assets workers can't afford to spend time searching for better jobs. The effect is to reduce lifetime expenditures and welfare, although it is uncertain by how much.

Turning to the upper end of the income distribution, an omitted effect that makes higher inflation less progressive because it favors the wealthy is the lag in updates to the property values that feed into property tax assessments. Property taxes represent a significant share of total tax revenues, and it is the largest source of state and local revenues. Property taxes collected in 2023 totaled $778 billion. By comparison, total state income tax revenues were $546 billion and total federal tax revenues were $4,439 billion in the same year. Property value typically is reassessed every 1 to 5 years. A back-of-the-envelope calculation shows that with a 3-year lag in valuation updates, and an increase from 0 to 10% inflation, the present value of property taxes over 20 years is 8.3% lower at the 10% inflation rate.

Another effect that is likely to mute the expenditure reductions of wealthier households relative to the predictions of the model are reductions in planned bequests. Recall that the model assumes all bequests are involuntary, but the data suggests otherwise. Presumably, part of the response to higher tax bills would be to cut back on the size of planned inheritances.

Aggregate and policy implications

The aggregate result that the authors emphasize is that a permanent shift in the rate of inflation from zero to 10 percent reduces median lifetime spending by 6.09 percent. That finding seemed at odds with my intuition that in this type of "all else equal" experiment with no changes in real output or real investment returns, the average effect on expenditures should be close to zero. A natural question then is why the model predicts a drop in average or median expenditures?

I think the explanation lies in what the government is implicitly assumed to do outside of the tax and benefit programs explicitly modeled. Inside the model, aggregate tax collections increase relative to aggregate benefit expenditures, leaving people with less money to spend. If other government revenues and expenditures are also held constant, and if there are no changes to the real economy, then the government will run larger surpluses than it did prior to the inflation

[1] Thanks are due to Monroe Gamble who provided me with data on program rules.

increase. If the larger surpluses are saved rather than spent, then government debt will decrease over time. In a Ricardian world, households would expect cuts to future tax rates or some other form of rebate, and expenditures would be higher than predicted in anticipation of lower future taxes. Therefore, if the analysis were to take into account the government's budget constraint and if the usual transversality condition were imposed, as a first approximation no change in aggregate expenditures would be predicted. Presumably, the predictions about the cross-sectional distribution of expenditure effects would be similar to those reported, but with a higher mean and median.

Although the authors do not make policy recommendations, implicit in their analysis, and more explicit in the related literature, is the idea that it would be better to more fully index the fiscal system to inflation to avoid unintentional taxes and transfers. I am fully in agreement with the logic of that idea and the merits of avoiding unnecessary distortions. However, it may be that in a second-best world full indexation would not be optimal. Budget analysts have long warned of ballooning debt and deficits, and the need for some combination of increased taxes and reduced benefits to move to a sustainable path. Of course, such policy changes are politically unpopular, particularly for entitlement programs like Social Security and Medicare that enjoy broad public support. The "failure" to index AGI thresholds for Social Security or to index Medicare's Part-B top premium may be one of few politically feasible channels to limit spending growth and increase revenues. Congress can then reverse unwanted distortionary effects with much less resistance by lowering taxes and increasing benefits. It is notable that full indexing in the private sector is also unusual, and a better understanding of why would help inform the question of optimal indexation in the public sector.