MPRA
Munich Personal RePEc Archive

The Forced Bid What Privatizing Social Security Would Do to Stock Prices, and Why Recent Developments in the Economics of Market Flows Strengthen the Case Against It

Lane, Edward

7 August 2026

Online at https://mpra.ub.uni-muenchen.de/130383/
MPRA Paper No. 130383, posted 03 Sep 2026 10:25 UTC

The Forced Bid

What Privatizing Social Security Would Do to Stock Prices, and Why Recent Developments in the Economics of Market Flows Strengthen the Case Against It

Edward Lane, ASA, CFP®

July 2026

Executive Summary

Proposals to privatize Social Security, in whole or through a partial carve-out of the payroll tax, rest on one claim: personal accounts invested in the stock market will provide better benefits than the program can pay. This paper examines that claim from three directions, actuarial, financial, and fiscal, and finds that it fails in each.

The actuarial point comes first. Social Security is not an investment account with a disappointing return. It is a package of insurance: retirement income, disability income, and survivor and family benefits, all indexed to inflation for life and built on a progressive benefit formula. The proper private-market benchmark is the cost of replacing that package, an inflation-adjusted life annuity plus disability and survivor coverage, net of decades of fees.

Measured against that benchmark, the comparison fails before any market analysis begins. The international record confirms it. Chile, the one completed experiment, raised its mandatory contribution rate from 10 to 16 percent and still delivered less than half of projected benefits. Sweden achieves its apparent replacement rates only by charging 18.5 percent of wages and dropping inflation protection.

The financial argument is the paper's central contribution. The academic privatization debate closed around 2005 within the classical framework of Diamond and Geanakoplos (2003), in which flows do not move prices because bond issuance induces existing investors to sell stocks to the new accounts. In that framework, the main effect of the policy is a smaller equity risk premium, which removes the 6 to 7 percent real returns assumed in the privatizers' projections.

In 2021, Gabaix and Koijen's inelastic markets hypothesis challenged previous assumptions, revealing that forced flows can increase the total market value by several times the invested dollars. When applied to a 2005-like carve-out, which involved about $330 billion annually in automatic, price-insensitive equity purchases, the analysis in

Appendix A indicates that valuations could rise approximately 28% above the baseline within ten years. This leads to a decrease in forward returns by nearly 1% annually due to yield compression alone and results in around $15 trillion in capital gains accruing mainly to the top tenth of households, who possess about 90% of equities.

The two frameworks disagree on prices but agree on the policy: under either model, the promised returns cannot survive the policy that is supposed to produce them. In addition, every dollar that flows in is scheduled to flow back out. When the large cohorts retire, the structural buying becomes structural selling, concentrated on the people least able to bear it.

The fiscal critique completes the case. As economists in the functional finance and Modern Monetary Theory traditions argued throughout the original debate, the federal government's ability to pay legislated benefits was never in question; the trust fund's ledger is internal bookkeeping, and personal accounts financed by government borrowing create no new national saving. They relabel the deficit while adding fees and market risk. Retirement is paid for in real goods and services produced by the future workforce, and no rearrangement of financial claims can produce those resources.

The conclusion: if broader asset ownership is the goal, the answer is an add-on, not a carve-out. A low-cost national savings vehicle modeled on the federal Thrift Savings Plan, layered on top of an intact Social Security system, achieves what privatizers say they want without the harm. That test is a useful way to evaluate what any particular privatization proposal is actually for.

Table of Contents

Executive Summary 1

Section 1. Introduction: An Old Question and an Outdated Answer 4

Section 2. The Benefits Personal Accounts Would Have to Replace 5

Section 3. The Classical View: Bond Issuance Offsets Stock Purchases 7

Section 4. The Inelastic Markets View: Flows Move Prices 8

Section 5. Both Views Undermine the Promised Returns 9

- 5.1 The classical route: the equity premium shrinks 10
- 5.2 The inelastic route: buying at inflated prices 10
- 5.3 Baker's arithmetic, with a mechanism 11

Section 6. Distributional Effects: A One-Time Transfer to Existing Shareholders 11

Section 7. The Decumulation Phase: The Flows Reverse 12

Section 8. International Experience: Chile, Sweden, and Australia 13

- 8.1 Chile 13
- 8.2 Sweden 13
- 8.3 Australia 14

Section 9. The Functional Finance and MMT View 15

- 9.1 Affordability was never the issue 15
- 9.2 Personal accounts create no new saving 16
- 9.3 Retirement is paid in real resources 17
- 9.4 What each literature adds to the other 18

Section 10. Conclusion 18

Appendix A: A Stylized Calibration of the Forced Bid 20

- A.1 Inputs 20
- A.2 Arithmetic 20
- A.3 Caveats 21

Acknowledgment 23

References 23

Section 1. Introduction: An Old Question and an Outdated Answer

Every proposal to privatize Social Security, whether a full replacement or a partial carve-out of the payroll tax, begins with the same premise: the stock market pays more than Social Security does. The Bush administration proposed a version in 2005 that would have allowed workers to divert 4 percentage points of the 12.4 percent payroll tax into personal accounts. Australia built an entire retirement system on mandatory individual accounts. The idea returns to American politics every few years, and it has once again (Malito 2026). A natural question, asked by every investor and congressional staffer who encounters the idea, is what all that money would do to stock prices. The answer matters, because the projected returns of 6 to 7 percent real[1] are the numbers on which every privatization proposal depends.

The academic literature that grew up around the 2005 debate concluded that prices would barely move. The most rigorous treatment, by Diamond and Geanakoplos (2003), found that the money flowing into stocks would be offset by the government bonds issued to finance the transition from Social Security. That answer settled into the literature, the political moment passed, and the question went dormant.

In 2021, financial economics changed its view of how markets absorb money. Gabaix and Koijen published evidence that the stock market is far less flexible than the classical models assume. In their framework, the inelastic markets hypothesis, a dollar of new money flowing into stocks raises the total value of the market by roughly five dollars. This work is part of the flows revolution in asset pricing.[2] To my knowledge, no one has re-examined the question of privatization using this framework.

The policy debate froze in 2005 with the old assumptions, and the two literatures have never been brought together. I do so in this paper. It extends my earlier Levy Institute paper (Lane 2025), which examined privatization from the actuarial and benefit-design sides, by adding the market side of the argument, and it begins where that paper left off: with the benefits the accounts would actually have to replace.

The classical work on privatization and asset markets, from Auerbach and Kotlikoff's general equilibrium life-cycle models through the Feldstein volumes and Abel's analysis of

[1]The 1994–1996 Advisory Council on Social Security assumed a 7.0 percent real equity return, based on the historical average; the President's Commission to Strengthen Social Security (2001) used the SSA actuaries' assumption of 6.5 percent real. See Baker (1997) for contemporaneous documentation and critique of these assumptions.

[2]The flows revolution in asset pricing is the post-2019 body of work showing that investor demand is inelastic and money flows move prices.

trust fund equity investment to Diamond and Geanakoplos (2003), operates in frameworks where flows have no independent price impact. The inelastic markets literature that follows Gabaix and Koijen (2021) has been applied to quantitative easing, share buybacks, and index flows, but not, so far as I can determine, to Social Security reform. The Modern Monetary Theory and functional finance critiques of privatization, discussed in Section 9, have never been joined to a formal asset-pricing mechanism.

This paper integrates the three perspectives and shows that they reach the same conclusion by different routes. The framework also applies beyond this question: the same mechanics govern sovereign wealth funds, mandatory superannuation, and any other policy that directs persistent, price-insensitive flows into an inelastic market.

Three results follow.

- First, the two frameworks give different answers about stock prices, and the difference determines who gains and who loses.
- Second, both frameworks eliminate the return assumptions on which privatization depends, each by a different route, so the case against carve-outs does not require choosing between them.
- Third, neither framework changes the real economics of retirement: retirees consume food, housing, and care produced by the working population alive at the time, whatever form their financial claims take.

Privatization rearranges the claims. It does not create the resources. Economists in the MMT and functional finance traditions made this point throughout the original debate and were largely ignored. One purpose of this paper is to show that the newest asset-pricing research and the older critique arrive at the same verdict, and that each supplies what the other lacks.

Section 2. The Benefits Personal Accounts Would Have to Replace

The privatization argument compares a stock portfolio's expected return with Social Security's implied return on contributions. The comparison is incorrectly specified, because Social Security is not an investment account. It is a package of insurance benefits, each with a distinct actuarial cost, and the accounts would have to replace the whole package (Lane 2025). Here is a summary:

1. Retirement income benefits. Retirement benefits are based on wage-indexed career average earnings. Reduced benefits are available as early as age 62.

2. Spousal income benefits. Spouses and certain ex-spouses are eligible for up to 50 percent of the worker's retirement or disability benefit.

3. Disability income benefits. Eligible disabled persons may receive lifetime income.

4. Survivors' income benefits. Eligible surviving spouses, children, and dependent parents may receive lifetime income.

5. Children's income benefits. Children of retirees and disabled workers are eligible for income until age 18, or longer if the child is disabled or a student.

6. Family income benefits. Combined family benefits may reach 150 to 180 percent of the worker's own benefit.

7. Inflation adjustment. All of the above are increased annually based on the Consumer Price Index (CPI-W), for life.

8. A progressive formula. The benefit formula replaces a larger share of wages for lower earners.

The proper private-market benchmark is therefore not the historical return of a balanced portfolio. It is the cost of replacing this package at retail: an inflation-adjusted life annuity, plus disability insurance, plus survivor and family coverage, plus the management fees that compound over a forty-year working life.

Priced that way, the comparison fails. Inflation-indexed lifetime income is among the most expensive products the private market sells, when it sells it at all. There is no more efficient mechanism than a defined benefit plan for delivering inflation-indexed retirement, survivor, and disability income because it pools longevity, market, and disability risk across the entire workforce. Unbundling it forces each worker to self-insure longevity or to buy annuities individually at retail prices.

As a former consulting actuary, a consultant to the Office of Personnel Management during the 1980s studying private-sector replacement plans for the Civil Service Retirement System, and a consultant to and observer of many corporate conversions of defined benefit plans to 401(k) and cash balance plans, I have not observed a single instance in which the motivation for the change was to increase benefits for employees. The goal was always to reduce sponsor cost, shift investment risk to employees, or both (Lane 2025). Social Security privatization would be the largest defined-benefit-to-defined-contribution conversion ever undertaken, and there is no reason to expect it to be the first exception.

Two further consequences follow from unbundling, and both fall on the people the current program protects best. First, the progressive formula disappears. Accounts are proportional to contributions at best, so the change removes relatively more protection from lower earners, workers with interrupted careers, and women, who make up about 60 percent of beneficiaries and rely disproportionately on spousal and survivor benefits. Second, any voluntary carve-out invites adverse selection. The highest earners and those without dependents have the most to gain from opting out, thereby draining revenue from the remaining system and worsening its finances for everyone who stays.

The claim to be tested, then, is mis-specified before a single payroll dollar reaches a brokerage account. Sections 3 through 6 take the claim on its own terms, as a contest of investment returns, and show that the returns are not there either. The reason is the policy itself.

Section 3. The Classical View: Bond Issuance Offsets Stock Purchases

Every privatization plan must solve the same transition problem. Social Security is pay-as-you-go: today's payroll taxes pay today's benefits. If workers divert part of the payroll tax into personal accounts, benefits to current retirees do not stop, so the government fills the gap by borrowing. Every dollar that flows into a personal account and then into the stock market is matched, dollar for dollar, by new Treasury bonds sold to the public (Lane 2024).

Now follow both flows. Households, through their new accounts, are buying stocks. Someone must buy the new bonds, and someone must sell the stocks the accounts are buying. In the classical view, prices adjust until willing sellers meet willing buyers, and what matters in the end is the total amount of risk the private sector must hold, not the flow of money. After privatization, the private sector as a whole owns the same companies it owned before, plus a larger amount of government debt.[3] The stock of real assets in the economy is unchanged.

Diamond and Geanakoplos (2003) worked through this logic for the related question of trust fund investment in equities, and their conclusion carries over to personal accounts. In a frictionless market, the main effect of moving Social Security funds into stocks is not a change in stock prices. It is a change in relative returns. Because stock market risk is spread across more households, including workers who previously owned no stocks, the

[3]Held through institutional layers, money funds, banks, pensions, and bond funds, which are ultimately household-owned; Federal Reserve holdings simply substitute one government liability (reserves) for another (bonds) and do not change the aggregate. The segmented identity of the marginal Treasury buyer is taken up in Section 4.

extra return investors demand for holding stocks, the equity risk premium, shrinks. At the same time, the safe interest rate rises, because the private sector must be induced to hold the new government debt.

In short, the classical model says privatization would not produce a stock market boom. It would make stocks a somewhat worse investment and bonds a somewhat better one, permanently. Section 5 returns to what that does to the projected returns.

Note what the model assumes: that when new money arrives to buy stocks, existing holders will sell at close to current prices and move into bonds. In other words, it assumes the demand curve for the aggregate stock market is nearly flat. That assumption went largely unquestioned in the privatization debate because it went largely unquestioned in finance.

Section 4. The Inelastic Markets View: Flows Move Prices

In 2021, Gabaix and Koijen circulated a paper titled “In Search of the Origins of Financial Fluctuations: The Inelastic Markets Hypothesis.” Using detailed data on who holds U.S. equities, they asked a simple question: when new money flows into the stock market, who sells?

The answer is that almost nobody sells without a large price change. The modern stock market is dominated by holders who do not respond to prices as textbooks assume. Index funds hold whatever the index holds. Target-date funds hold a fixed stock-bond mix set by the calendar. Pension funds and insurers operate under mandates that fix their equity allocations within narrow bands. When prices rise, these institutions do not sell; many mechanically buy more. The pool of genuinely price-sensitive investors, the ones who would sell stocks to a wave of new buyers, is small.

The consequence is that flows substantially affect prices. Gabaix and Koijen estimate a multiplier of roughly five: one dollar of net inflow into the aggregate stock market raises total market capitalization by about five dollars, with their own estimates ranging from three to eight across specifications.

The price impact is long-lasting, because a permanent shift in demand produces a permanent shift in the equilibrium price. The figure of five is a central estimate drawn from ordinary variation in aggregate flows. Applying it to a legislated program of the scale considered here is an extrapolation beyond the settings in which it was estimated, which is why Appendix A presents a range of multipliers, and Section A.3 discusses the extrapolation directly.

Now apply this to a Social Security carve-out. The 2005 proposal would have diverted up to 4 percentage points of payroll. Scaled to today's wage base, that is on the order of $400 billion per year of automatic, price-insensitive equity buying, arriving with every payroll cycle for decades. For comparison, the net supply of U.S. equities frequently shrinks in a given year, because companies repurchase more stock than they issue. A forced bid[4] of this size would be among the largest sustained flows in the market's history.

What about the offsetting bond issuance that neutralized everything in the classical model? Under inelasticity, it does not neutralize, because the offset requires the same investors to operate on both sides. In practice, the institutions that absorb new Treasury supply, money market funds, foreign central banks, and bond-mandate investors, are not the institutions that would need to sell stocks to the personal accounts, and most of the potential stock sellers are constrained by mandates from selling at all. The two flows enter different pools, causing the equity pool to reprice.

The state of knowledge is therefore this: the best model of 2003 says privatization barely moves stock prices, and the best-supported model of market flows since 2021 says it moves them substantially. The next section shows that the disagreement does not matter for the policy conclusion.

Section 5. Both Views Undermine the Promised Returns

It is not necessary to decide which model is correct in order to evaluate privatization, because both models eliminate its central claim: stock market returns of 6 to 7 percent per year after inflation, the assumption on which every privatization projection depends. The two frameworks differ as summarized below, and neither is a straw man. Each was the best-supported model of its era, and each fails the proposal on its own terms.

Table 1

	Classical (Diamond-Geanakoplos)	**Inelastic (Gabaix-Koijen)**
Ownership structure	Price-sensitive investors free to rebalance across assets	Index funds, target-date funds, and mandate-bound institutions that hold regardless of price

[4]A bid is an order to buy.

Demand curve for aggregate equities	Nearly flat: sellers appear at close to current prices	Steep: roughly $5 of market value per $1 of net flow
Effect of the carve-out on prices	Minimal: bond issuance offsets equity purchases	Large, persistent valuation increase during accumulation
Route by which promised returns fail	Equity premium shrinks as risk spreads across the workforce	Participants buy at prices their own flows have raised
Distributional effect	Small: prices barely move	One-time capital gain to existing shareholders

5.1 The classical route: the equity premium shrinks

In the Diamond-Geanakoplos framework, spreading stock ownership across the entire workforce reduces the reward for holding stocks. The equity premium exists because stock market risk is unpleasant and concentrated. Spread that risk across roughly 170 million workers, and it is no longer scarce, so the market no longer pays as much to bear it. The safe rate rises, the equity premium falls, and the expected return on the personal accounts lands well below the historical average assumed in the projections.

The projections were built on a market that the policy itself would change. This is not a hostile reading of the classical model. The compression of the premium is Diamond and Geanakoplos's central result, a real cost of the policy established within the privatizers' preferred assumptions, before any question of flow inelasticity arises.

5.2 The inelastic route: buying at inflated prices

In the Gabaix-Koijen framework, the mechanism is different, but the destination is the same. The forced bid raises valuations. A valuation is a price paid today for earnings delivered later, and the higher the price paid, the lower the return earned. The workers whose accounts do the buying are, by construction, buying at the prices their own flows created. Forward returns fall mechanically. The historical returns quoted in privatization materials were earned by investors who bought at historical valuations. The personal-account generation would buy at valuations that history has not produced, because history has not contained a forced bid of this size.

5.3 Baker's arithmetic, with a mechanism

Dean Baker pointed out during the original debate that the privatizers' numbers contradicted one another (Baker 1997). The case for calling Social Security insolvent rested on the Trustees' assumption of slow economic growth. The case for personal accounts rested on real stock returns of 6 to 7 percent. Stock returns ultimately come from corporate earnings, and corporate earnings ultimately grow with the economy. Permanently slow growth and permanently high stock returns can coexist only if price-to-earnings ratios rise without limit, which is impossible.

The inelastic markets hypothesis turns this arithmetic observation into a mechanism. It specifies how valuations are pushed up, who pushes them (the account holders themselves), roughly how large the effect is, and who receives the difference. That last question is the subject of the next section.

Section 6. Distributional Effects: A One-Time Transfer to Existing Shareholders

This section presents the result that holds only under the inelastic model, and it is the one that matters most to anyone concerned with the distribution of wealth. In the classical model, there is no windfall, because prices barely move. Under inelasticity, the accumulation phase produces a large, one-time capital gain, and that gain goes to whoever owns stocks on the day the policy is enacted.

Ownership of U.S. equities is among the most concentrated forms of wealth in the country. The top 10 percent of households own roughly 87 percent of corporate equities and mutual fund shares, and the top 1 percent alone own just over half (Board of Governors, Distributional Financial Accounts, fourth quarter 2025). The bottom half of the country owns about 1 percent.

Now trace the flow of funds. The money doing the buying is payroll tax, deducted from wages, and the payroll tax is capped, so it comes disproportionately from middle and lower earners as a share of income. The price appreciation that money generates accrues to existing shareholders in proportion to their holdings.

The most robust part of this decomposition does not depend on any multiplier: whatever the price appreciation turns out to be, roughly ninety cents of every dollar of it goes to the top tenth of the wealth distribution, because that is who owns the equities. The multiplier only sets the scale. At the central estimate, each dollar of payroll tax diverted into the accounts creates roughly five dollars of market value, of which about four and a half go to the top decile, generated by a tax that applies only to wages and stops at the cap. Appendix

A performs the same arithmetic at market scale, reaching the trillions within the first decade.

Under inelastic markets, then, privatization converts the forced saving of wage earners into capital gains for the wealthiest tenth of households. The transfer operates through market prices rather than through any visible tax or spending provision. It would appear in no budget table and is invisible to the frameworks that dominated the original debate because those frameworks assumed it away.

For readers of my earlier work on the two-track character of the American economy (Lane 2026), this is the distributional core of the case: a policy presented as democratizing capital ownership would deliver its largest single benefit, immediately and in capitalized form, to households at the top, while wage earners pay in over forty years for returns that their own inflows have reduced. The workers do eventually own stocks. They own them at the highest prices the market has ever sustained, purchased out of wages.

Section 7. The Decumulation Phase: The Flows Reverse

Every dollar that flows into a retirement account is scheduled to flow back out from the day it arrives. The demographic structure that strains Social Security's finances, many retirees supported by relatively fewer workers, does not disappear under privatization. It reappears as a flow reversal.

During the decades when the large working cohorts are accumulating, the forced bid supports prices. When those cohorts retire and begin selling to finance consumption, the structural buying becomes structural selling.

In the classical model, this is a non-event, since flows do not move prices in either direction. In inelastic markets, it is a serious problem: the same mechanism that raised valuations on the way in lowers them on the way out, and it does so while the retired cohort is drawing down and is least able to wait out a bad decade.

The symmetry is mechanical. If inflows of roughly half a percent of market value per year raise valuations by 25 to 30 percent across the accumulation phase, comparable outflows exert an equivalent downward pressure, absent offsetting policy or behavioral change. The timing of the reversal is empirically uncertain, and the literature on demography and asset returns suggests it arrives gradually rather than abruptly (Poterba 2001). Gradual, however, is not the same as absent.

Individual timing risk compounds the aggregate problem. Whatever a balance is worth, the retirement income it buys depends on interest rates and market levels on the date it converts to an annuity or begins drawdown. Two workers with identical account balances can end up with very different retirement incomes simply because they retired a few months apart (Lane 2025). Sweden's system, discussed in the next section, runs exactly this lottery at every retirement. A benefit formula does not. It does not matter to a Social Security benefit what the market did in the month a worker turned 67.

The timing of the risk deserves emphasis. Under privatization, the risk of the system is concentrated on people at the end of their working lives, holding assets whose prices are under structural pressure from the simultaneous selling of their own age cohort. Social Security's defined benefit carries no such property, and that is much of the reason to have one.

Section 8. International Experience: Chile, Sweden, and Australia

Privatization is not hypothetical. Versions of it have operated abroad for decades, and the record bears directly on the claim of better benefits at lower cost. Three cases matter most: Chile, which ran the complete experiment; Sweden, the example American privatizers cite most often; and Australia, the largest and most successful version, whose experience also tells against a U.S. carve-out.

8.1 Chile

According to a 2005 analysis by the Social Security Administration (Kritzer 2005), Chile was the first country to implement mandatory individual accounts for retirement income in 1981. The goal at the time was to improve benefits through investment and keep public costs low. As it turned out, emerging benefits amounted to less than half of what was projected, and the mandatory contribution rate has been raised from 10 to 16 percent (Lane 2025). The combined U.S. OASDI employer and employee rate is 12.4 percent.

Argentina, Bolivia, Colombia, El Salvador, Mexico, Peru, and Uruguay followed Chile; as of 2019, their average combined contribution rate was also about 16 percent, and several have partially or wholly reversed their privatizations. Chile itself has added public guarantees back underneath the accounts in successive reforms, because the accounts alone left too many retirees with too little. The completed experiment costs more than the U.S. system, delivers less than it projected, and requires the state to return.

8.2 Sweden

Sweden is the system the Cato Institute holds up as a model, and it is instructive, though not in the way intended. The Swedish public pension credits 16 percent of salary to a notional account that grows with wage rates, plus 2.5 percent to an account the worker invests. At retirement, the balances are converted to an annuity using a unisex life table.

Two features do the quiet work. First, the total pension contribution of 18.5 percent of wages is roughly twice the U.S. OASDI rate. Second, the annuity is fixed, not inflation-adjusted. That is how the apparently higher replacement rates are produced: dropping the inflation adjustment yields a higher initial benefit that then loses real value for the rest of the retiree's life.

The comparison with U.S. Social Security, where every benefit is CPI-indexed for life, is apples-to-oranges, and the Swedish disability, survivor, and dependent benefits are not comparable either (Lane 2025). In effect, the Swedish government has transferred the inflation risk to the retiree. What Sweden demonstrates is the pricing point of Section 2: even a well-run individual account system delivers its headline numbers only by charging materially more and by shifting inflation risk onto participants.

8.3 Australia

Australia is the strongest case for the other side, so it deserves the most careful reading. Since 1992, Australian employers have been required to contribute a percentage of wages, now 12 percent, to superannuation accounts invested largely in markets. The system has accumulated roughly A$4 trillion, more than Australia's annual GDP, and it is genuinely popular. Three observations matter for the American question.

First, superannuation is the closest real-world counterpart to the forced bid modeled in Section 4: a mandated, price-insensitive flow, deducted from every paycheck, running for more than three decades. The outcomes are what the inelastic framework predicts, and the classical framework does not.

The Australian exchange has traded at a persistent valuation premium relative to comparable markets; the asset pool has grown larger than the country's GDP; a small number of very large funds now rank among the biggest owners in the market they are required to keep buying; and corporate behavior has adapted to the permanent flow, with an unusually high dividend payout culture reinforced by franking credits that make dividends especially valuable to domestic super funds. Australia is not a parallel story. It is the accumulation phase of the mechanism described in this paper.

Second, Australia never faced the transition-financing problem, because it never had an American-style pay-as-you-go system to convert. The means-tested Age Pension is paid from general revenue, and superannuation was layered on top as a new savings scheme rather than carved out of an existing program. There was no offsetting flood of bond issuance. The U.S. case would be the asset swap described by the classical model, so Australia cannot show that the swap is harmless. It never ran that experiment.

Third, Australia has not yet reached the deep drawdown phase. Compulsory contributions began in 1992, the contribution rate has risen repeatedly, and net flows remain positive, so the decumulation test of Section 7 is still decades away. Australia is evidence for the first half of the story only.

One further lesson does translate: as the superannuation pool outgrew the domestic market, the funds diversified heavily offshore, so the marginal compulsory dollar increasingly buys global assets rather than Australian ones. A U.S. program would face the same pressure at greater scale and with fewer places for the flow to go, because the U.S. market is where the rest of the world diversifies.

Section 9. The Functional Finance and MMT View

A third body of work addressed the original privatization debate and was largely ignored. Economists in the functional finance and Modern Monetary Theory traditions, Robert Eisner, L. Randall Wray, Stephanie Kelton (then writing as Stephanie Bell), and Dimitri Papadimitriou among them, argued that the debate rested on a false premise. Their case had three parts, and each has held up.

9.1 Affordability was never the issue

The premise of every privatization proposal is that Social Security is going broke. That premise treats the federal government like a household, which it’s not. A government that issues its own currency can always make any payment denominated in that currency. This was not a fringe observation at the time; it was stated in the debate's own record by the sitting Federal Reserve chairman. In March 2005, Representative Paul Ryan asked Alan Greenspan in a House Budget Committee hearing whether personal accounts would improve solvency. Greenspan answered that there is nothing to prevent the federal government from creating as much money as it wants and paying it to somebody; the real question is whether the economy will produce the goods and services those payments are supposed to buy (C-SPAN 2020).

Bell and Wray (2000) had made the same point five years earlier and in more depth: once solvency is recognized as a non-issue for a currency-issuing government, the Social Security "problem" reduces to a demographic question about future real output, which is smaller and more tractable than the financial crisis the privatizers described. Insolvency was never the binding constraint. Real production was, and the privatization debate spent its life arguing about the wrong one.

The trust fund itself, in this view, is internal bookkeeping: one arm of the government holding IOUs from another. Its projected depletion date is an accounting convention that Congress could change by statute, not an economic event. Eisner (1998) made the argument directly: benefits can be paid regardless of what the fund's ledger shows, and the ledger can be credited with whatever amount is needed.

Papadimitriou and Wray (1999), writing at the Levy Economics Institute, showed that the demographic burden itself, measured in real terms, is modest relative to ordinary productivity growth over a comparable period. The federal government cannot prefund future obligations in any real sense; intragovernmental debt in the trust fund merely records a future obligation (Lane 2025).

9.2 Personal accounts create no new saving

The second argument is an accounting point that the privatizers never addressed, and it connects directly to the asset swap in Section 3. Advocates presented personal accounts as a way to prefund retirement and raise national saving. The sectoral balances say otherwise.

When payroll taxes are diverted into accounts, government revenue falls, and the government borrows to keep paying benefits. The deficit grows by the amount flowing into the accounts, and the private sector's new financial assets, the account balances, are matched dollar for dollar by its new holdings of government debt. Net national saving does not increase.

The policy exchanges one existing government liability, the promise to pay benefits, for a pair of new ones, bonds plus market-dependent account balances, and adds management fees and market risk in between. Wray made this point repeatedly during the debate: the accounts are not new wealth; they are relabeled government borrowing.

Accepting this does not require endorsing a broader theoretical framework. Sectoral balances are simply accounting identities valid in any macroeconomic context, so one could agree with the previous points while rejecting all other aspects of MMT. Additionally,

the assertion that deficits are always irrelevant is incorrect; the real-resource constraint discussed below is precisely where deficits matter.

This conclusion isn't limited to heterodox thinkers. Mainstream critiques have also aligned with this view regarding transition borrowing and risk-shifting, from Brookings scholars in 2005 to later analyses by the congressional Joint Economic Committee (Joint Economic Committee 2010). It's also worth noting that this accounting insight aligns with Diamond and Geanakoplos's asset-pricing approach: financial claims on the future cannot replace the need to prepare for it.

9.3 Retirement is paid in real resources

The third argument is the deepest. Retirees do not consume dollars or shares. They consume food, housing, energy, and, above all, the labor of other people, particularly in health care and personal care, during the years of consumption. All of it must be produced by the workforce alive and working at that time. This holds whether a retiree's claim takes the form of a Social Security benefit, an account statement, or cash. The financial arrangement determines the strength of each person's claim on future production. It does not create the production.

Privatization, seen this way, converts a claim backed by the taxing power of the federal government into one backed by future stock prices and charges each worker individually for market risk, timing risk, and fees. The only channel through which it could make future retirees better off in real terms is by causing the economy to build more productive capacity, so that the future workforce produces more per person. In textbook terms, higher stock valuations are supposed to stimulate real investment, because expensive shares make it cheap for companies to raise capital.

The last fifteen years tested that channel, and it failed. American corporations spent the era of historically rich valuations repurchasing shares at a pace that frequently exceeded the combined pace of all new issuance. Corporate investment was not constrained by the cost of equity; companies were retiring equity as fast as they could.

Raising valuations further with a forced Social Security bid would, on this evidence, produce more of the same: higher prices for existing assets, larger buybacks, and no boom in productive capacity. The real burden of supporting the retirees of the 2050s would remain exactly what demography says it is, carried on riskier paper. **If policymakers want to prepare for that burden, the prescription follows directly: invest now in the things that raise future productivity, and stop treating account balances as if they were future goods and services.**

9.4 What each literature adds to the other

For all its clarity about the affordability question and the real-resource constraint, the MMT critique had a gap. It had no formal model of what the forced flows would do inside asset markets, because the finance of the day held that the answer was nothing. MMT economists could correctly say that privatization created no new resources and shifted risk onto workers. They could not say, with a mechanism and a magnitude, what would happen to prices, who would capture the appreciation, and what decumulation would do to the retiring cohort.

The inelastic markets hypothesis provides the missing mechanism: a quantitative explanation of the upward transfer, several dollars of market value per dollar of forced flow, accumulated in the portfolios of households that already hold nearly all the equity. The process goes both ways. The flows literature, which does not take a stance on policy, benefits from the functional finance approach in recognizing that no amount of price appreciation can resolve the real provisioning issue. Therefore, a valuation boom does not constitute a social gain. Each body of literature addresses the other's unresolved question.

Section 10. Conclusion

The early-2000s privatization debate took place within a framework in which market flows were irrelevant. Within this model, the main objections to personal accounts centered on risk shifting, transition costs, and fees. Meanwhile, their primary benefit, higher expected returns, was diminished by the potential for a reduced equity premium, though not completely eliminated.

The inelastic markets hypothesis changes the terrain, and every change strengthens the case against carve-outs. It implies that the accumulation phase would raise the very prices the accounts must pay, mechanically lowering the returns the policy promises. It reveals a large regressive transfer, from payroll to existing portfolios, that the older models could not see. And it exposes a decumulation-phase risk, structural selling into an inelastic market by an entire retiring generation, that falls on exactly the people a retirement system exists to protect.

The classical argument has not gone away. A defender of privatization who rejects the inelastic view and insists that markets are frictionless walks directly into the Diamond-Geanakoplos result that frictionless markets eliminate the equity premium the proposal is selling.

Besides the market arguments, the actuarial point of Section 2 stands on its own: the accounts have always been measured against the wrong benchmark, a raw portfolio return set against an inflation-indexed insurance package covering retirement, disability, survivorship, and dependents, which no individual account can replicate at comparable cost.

There is no current, credible model of asset prices under which diverting payroll taxes into personal accounts delivers the returns its advocates project. Under the best-supported recent model, the proposal would also, in its first decades, function as one of the largest upward transfers of wealth in American history, carried out through the price mechanism.

None of this condemns the goal of broader asset ownership, which is worth pursuing. It has a straightforward policy answer and is an add-on, not a carve-out: a low-cost, well-managed opportunity for workers to set aside additional money for retirement, similar to the federal Thrift Savings Plan, layered on top of an intact Social Security system (Lane 2025).

An add-on creates no transition debt, removes no insurance protection, takes nothing from the workers who rely on the progressive formula, and, because it is voluntary and additive rather than a forced diversion of the payroll tax, introduces no market distortions like those described in this paper. Everything the privatizers say they want survives; only the harm is removed. That is a useful test of what any particular privatization proposal is actually for.

Social Security's actual challenge is real and demographic, as functional finance economists said from the beginning. The financial “insolvency” is a self-imposed accounting convention. The genuine task is sustaining the economy's productivity, capacity, and labor force participation to support mid-century retirees, and if Congress wishes to adjust the program's internal accounting along the way, raising or eliminating the taxable wage cap will do the job without reducing any worker's protections. Those are the debates worth having. The stock market is not a loophole in demography. The functional finance economists knew that in 1998; since 2021, finance has a model that explains, in its own terms, why they were right.

Appendix A: A Stylized Calibration of the Forced Bid

This appendix puts rough magnitudes on the mechanism in Sections 4 and 6. It is a transparent back-of-the-envelope calculation, not a forecast. Every input is stated so that a skeptical reader can substitute different values and rerun the arithmetic. The purpose is to establish orders of magnitude, and the conclusion survives large changes in every assumption.

A.1 Inputs

The policy is a 2005-style carve-out: workers can divert 4 of the 12.4% OASDI payroll tax into personal accounts. With a $10.5 trillion taxable payroll, full participation diverts about $415 billion annually. Assuming 80% of assets, about $330 billion, are allocated to equities, this flow represents roughly 0.5% of the $70 trillion U.S. equity market. The Gabaix-Koijen estimate of the flow-to-value multiplier M is about 5, with a range from 3 to 8; calculations use M values of 2, 5, and 8. The baseline total shareholder yield is approximately 3.5%.

A.2 Arithmetic

The annual valuation increase relative to the no-policy counterfactual is approximately M times the flow share, that is, M × 0.5 percent. Over ten years, holding the flow share constant, this results in the cumulative effects shown in the table. The forward-return drag is the compression of shareholder yield implied by the decade-end valuation level: a market priced at (1 + Δ) above the counterfactual has a yield of 3.5 percent divided by (1 + Δ), before mean reversion or flow reversal effects. Multiplying decade-one flows of about $3.3 trillion by M gives the market value created, with 90 percent allocated to the top decile of households, according to Section 6.

	M = 2 (low)	**M = 5 (central)**	**M = 8 (high)**
Annual valuation lift	1.0%	2.5%	4.0%
Cumulative lift after 10 years	≈10%	≈28%	≈48%
Shareholder yield at decade end	≈3.2%	≈2.7%	≈2.4%

Forward-return drag, yield channel alone	≈0.3 pts/yr	≈0.8 pts/yr	≈1.1 pts/yr
Market value created vs. counterfactual, decade one	≈$7T	≈$16.5T	≈$26T
Of which accruing to top decile (≈90% share)	≈$6T	≈$15T	≈$23T

Even at a low multiplier, the first decade transfers several trillion dollars of capital gains to existing shareholders and reduces future returns promised to accounts. The transfer averages in the mid teens of trillions, and return drag nears a full percentage point annually from yield compression alone, excluding the smaller equity premium or decumulation reversal. The privatization case relies on accounts outperforming Social Security's implicit return by about two percentage points, but this mechanism consumes most of that margin.

A.3 Caveats

Five caveats deserve emphasis, and it is worth noting which direction each cuts.

- First, the Gabaix-Koijen multiplier is estimated largely from unanticipated flows. A legislated program is anticipated, and anticipation invites front-running, which would pull the price appreciation forward to the moment of enactment and concentrate the windfall even more completely on incumbent shareholders, since the new accounts would begin buying only after prices had adjusted.

- Second, the calculation is partial equilibrium and ignores the bond side entirely; Section 4 explains why segmentation weakens the bond offset, but to whatever extent the offset operates, it moves the result toward the classical case, which fails the proposal by the other route.

- Third, holding the flow share of market value constant is conservative in one respect and generous in another: payroll grows with wages, but a market capitalization inflated by the mechanism grows faster, so the flow share, and with it the annual lift, likely declines over time.

- Fourth, none of this is a forecast of market levels, which depend on everything else that happens. It is an estimate of the policy's contribution relative to the counterfactual, which is the correct object of evaluation.

- Fifth, and most fundamentally, the elasticity is estimated from small, high-frequency variation in flows; nothing at the scale or duration of a carve-out exists in the data.

Two forces could reduce the multiplier at scale: price-sensitive sellers may start selling at stretched valuations, and corporations with low cost of equity might issue more to meet forced demand. Both push toward the classical case, countering each other, with the sensitivity range in the table, especially the low-multiplier column, bracketing them. The buyback era and Australian record indicate these dampers are weak at observed valuations but cannot be dismissed at unobserved ones.

Acknowledgment

In preparing this paper, the author used AI assistants (Anthropic's Claude and Perplexity, powered by GPT-5.1) for research, drafting, editing, and formatting support. All analysis, arguments, data selection, and conclusions are the author's own, and the author independently verified all references, figures, and quotations against primary sources.

References

Abel, Andrew B. “The Effects of Investing Social Security Funds in the Stock Market When Fixed Costs Prevent Some Households from Holding Stocks.” American Economic Review 91, no. 1 (2001).

Advisory Council on Social Security. Report of the 1994–1996 Advisory Council on Social Security, Volume I: Findings and Recommendations. Washington, DC, 1997.

Auerbach, Alan J., and Laurence J. Kotlikoff. Dynamic Fiscal Policy. Cambridge University Press, 1987.

Baker, Dean, and Mark Weisbrot. Social Security: The Phony Crisis. University of Chicago Press, 1999.

Baker, Dean. “Saving Social Security with Stocks: The Promises Don’t Add Up.” Economic Policy Institute and Twentieth Century Fund, February 1997.

Bell, Stephanie, and L. Randall Wray. 'Financial Aspects of the Social Security "Problem."' *Journal of Economic Issues* 34, no. 2 (2000).

Bell, Stephanie. “Do Taxes and Bonds Finance Government Spending?” Journal of Economic Issues 34, no. 3 (2000).

Board of Governors of the Federal Reserve System. Distributional Financial Accounts of the United States (equity ownership by wealth percentile), various releases.

C-SPAN. “User Clip: Alan Greenspan Answers Paul Ryan.” November 16, 2020. https://www.c-span.org/clip/house-committee/user-clip-alan-greenspan-answers-paul-ryan/4923562.

Diamond, Peter, and John Geanakoplos. “Social Security Investment in Equities.” American Economic Review 93, no. 4 (2003).

Eisner, Robert. “Save Social Security from Its Saviors.” Journal of Post Keynesian Economics 21, no. 1 (1998).

Feldstein, Martin, ed. Privatizing Social Security. University of Chicago Press, 1998.

Gabaix, Xavier, and Ralph S. J. Koijen. “In Search of the Origins of Financial Fluctuations: The Inelastic Markets Hypothesis.” NBER Working Paper No. 28967, 2021 (SSRN version revised December 2023).

Galbraith, James K., L. Randall Wray, and Warren Mosler. “The Case Against Intergenerational Accounting: The Accounting Campaign Against Social Security and Medicare.” Levy Economics Institute Public Policy Brief No. 98, 2009.

Geanakoplos, John, Olivia S. Mitchell, and Stephen P. Zeldes. “Would a Privatized Social Security System Really Pay a Higher Rate of Return?” In Framing the Social Security Debate, National Academy of Social Insurance, 1998.

Greenspan, Alan. Testimony before the U.S. House Committee on the Budget, March 2, 2005 (exchange with Rep. Paul Ryan on Social Security solvency and personal accounts).

Kelton, Stephanie. The Deficit Myth: Modern Monetary Theory and the Birth of the People’s Economy. PublicAffairs, 2020 (chapter on entitlements).

Kritzer, Barbara E. “Individual Accounts in Other Countries.” Social Security Bulletin 66, no. 1 (2005).

Lane, Edward. “Saving Social Security.” Levy Economics Institute Working Paper No. 1042, February 2024. https://ssrn.com/abstract=5190379.

Lane, Edward. “Protecting Social Security: The Case Against Privatization.” The Levy Economics Institute, July 2025.

Lane, Edward. “The K-Shaped Economy: Divergent Fortunes Across American Counties and the Potential Impact on the 2026 and 2028 Elections.” Unpublished manuscript, 2026. https://ssrn.com/abstract=6377838.

Malito, Alessandra. “Trump Floats Idea of Australia’s Retirement System Again. But What About Social Security?” MarketWatch, July 7, 2026. https://www.marketwatch.com/story/trump-floats-idea-of-australias-retirement-system-again-but-what-about-social-security-de63b8c6.

Papadimitriou, Dimitri B., and L. Randall Wray. “Does Social Security Need Saving? Providing for Retirees Throughout the Twenty-first Century.” Levy Economics Institute Public Policy Brief No. 55, 1999.

Poterba, James M. “Demographic Structure and Asset Returns.” Review of Economics and Statistics 83, no. 4 (2001).

President’s Commission to Strengthen Social Security. Strengthening Social Security and Creating Personal Wealth for All Americans. Final Report, 2001.

Social Security Administration. The 2026 Annual Report of the Board of Trustees of the Federal Old-Age and Survivors Insurance and Federal Disability Insurance Trust Funds.

U.S. Congress, Joint Economic Committee. “Unnecessary Risk: The Perils of Privatizing Social Security.” Washington, DC, 2010.