

The Insurance Value of Financial Aid

Kristy Fan¹, Tyler J. Fisher², and Andrew A. Samwick ^{*3}

¹Google, LLC

²Tamara Prep

³Dartmouth College and NBER

December 15, 2020

Abstract

Financial aid programs enable students from families with fewer financial resources to pay less to attend college than other students from families with greater financial resources. When income is uncertain, a means-tested financial aid formula that requires more of an Expected Family Contribution (EFC) when income and assets are high and less of an EFC when income and assets are low provides insurance against that uncertainty. Using a stochastic, life-cycle model of consumption and labor supply, we show that the insurance value of financial aid is substantial. Across a range of parameterizations, we calculate that financial aid would have to increase by enough to reduce the net cost of attendance by 30 to 80 percent to compensate families for the loss of the income- and asset-contingent elements of the current formula. This compensating variation is net of the negative welfare consequences of the disincentives to work and save inherent in the means-testing of financial aid. Replacing just the “financial aid tax” on assets with a lump sum would also reduce welfare.

JEL Classification: D15, I22, G52

*Contact Information: andrew.samwick@dartmouth.edu. 603-646-2893; 6106 Rockefeller Hall, Hanover, NH 03755. We thank Martin Boyer, Scott Carrell, Elizabeth Cascio, Susan Dynarski, Glenn Harrison, Jason Houle, Jonathan Meer, Jonathan Skinner, and seminar participants at the Quantitative Society for Pensions and Savings Summer Workshop, the Boulder Summer Conference on Consumer Financial Decision Making, the Society of Labor Economists annual meeting, the Association for Education Finance and Policy annual conference, the CEAR-RSI Household Finance Workshop, Middlebury College, and Southern Methodist University for helpful comments. We are grateful to Will Zhou, Raymond Chen, Madison Minsk, Sarah Hong, and Amy Hu for research assistance; to John Hudson for computational assistance; and to the James O. Freedman Presidential Scholars program and the Nancy P. Marion Academic Enrichment Fund at Dartmouth College for financial support. Any errors are our own.

1 Introduction

Attending college is an important pathway to higher earnings. Unconditional estimates of the gap in median earnings between year-round, full-time workers with bachelor’s degrees and those with only high school diplomas are about 40 percent for both men and women. Carneiro, Heckman and Vytlačil (2011) estimate the marginal returns to a year of college and find an earnings premium of 8 percent per year of college.¹

The high returns to college have made the financing of college an important topic for both academic research and public policy over the last several decades. High returns to college have been accompanied by high and rising costs of college attendance. Launched as part of the Great Society programs of the 1960s, the federal financial aid system has grown in scale and complexity to help ever more students from low- and middle-income families afford college and the access to higher earnings that college can provide. Most states and institutions of higher education also operate financial aid programs. As shown in recent surveys by Fidelity Investments (2018), how to pay for college is often a savings decision that parents begin before their child enters pre-school.

Financial aid programs exist to enable students from families with fewer financial resources to pay less to attend college than other students from families with greater financial resources. As implemented through both federal and institutional formulas, the amount of financial aid a student receives declines with both the income and assets of his or her family at the time of enrollment and is recalculated for every subsequent school year. The inclusion of family assets in the formulas to determine a student’s “Expected Family Contribution” (EFC) toward college expenses has attracted considerable attention from economists, who have highlighted the resulting disincentive for families to save for college expenses. The concept of the “financial aid tax” dates back to Case and McPherson (1986), and Edlin (1993) provided an early, readable discussion of the financial aid tax. Empirically, Feldstein (1995) estimated a large crowding out of saving due to this tax, spawning a small literature testing the

¹The unconditional estimates are based on the tabulations of person income for 2018 in the Current Population Survey, 2019 Annual Social and Economic Supplement, by U.S. Census Bureau (2019). See Barrow and Malamud (2015) for a recent review of estimates of the returns to college.

robustness of those initial results.²

Omitted from this literature is the recognition that the saving disincentives due to the financial aid tax comprise only the “incentives” side of a standard incentives-insurance tradeoff. In general, providing insurance against risks beyond a family’s control will distort incentives along margins that the family can control. When income is uncertain, a financial aid formula that requires more of an EFC when income and assets are high and less of an EFC when income and assets are low provides insurance against that uncertainty. The incentives-insurance tradeoff is well understood in the literature on optimal redistributive taxation.³ What has not yet been recognized is that, as redistributive mechanisms based on assets and income, an analogous tradeoff is present in financial aid formulas.

The contribution of our paper is to estimate the insurance value of financial aid using a stochastic, life-cycle model of consumption and labor supply in which families save in anticipation of a planned retirement, uncertain income, and the college education of their children.⁴ The main results show that the insurance value of financial aid is substantial. We calculate the insurance value of financial aid by comparing lifetime expected utility under two financial aid systems – one in which there is a stylized version of the current financial aid formula and one in which colleges give aid by simply discounting their tuition, regardless of a family’s income or assets. This comparison is analogous to substituting a revenue-equivalent lump-sum tax for a distortionary tax, which under certainty would be expected to make the family better off. However, given sufficient income uncertainty and risk aversion, the substitution may lower welfare by removing the insurance value of financial aid. Across a range of

²The implicit tax on assets also figures prominently into practitioner guidance on saving for college. See, for example, Onink (2014).

³See Eaton and Rosen (1980) and Varian (1980) for early analyses of this tradeoff in optimal income tax systems. Using variation in tax and transfer systems across U.S. states, Grant et al. (2010) show that state-level measures of redistributive taxation correlate negatively with the standard deviation of the within-state consumption distribution, consistent with an insurance effect of redistributive taxation. Using the Panel Study of Income Dynamics, Hoynes and Luttmer (2011) show that while the redistributive value of state tax-and-transfer programs declines sharply with income, the insurance value of these programs is increasing in income.

⁴In this respect, the analysis is similar to prior papers that have examined the insurance aspects of other tax and expenditure policies. See Hubbard, Skinner and Zeldes (1995) for precautionary saving and social insurance, Engen and Gruber (2001) regarding precautionary saving and the unemployment insurance system, and more recent papers by Rostam-Afschar and Yao (2014) on precautionary saving and progressive taxation, Athreya, Reilly and Simpson (2014) on the insurance value of the Earned Income Tax Credit, and Stepner (2019) on the insurance value of redistributive taxes and transfers, focusing on risks of illness and layoff.

parameterizations, we calculate that financial aid would have to increase by enough to reduce the net cost of attendance by 30 - 80 percent to compensate families for the loss of the income- and asset-contingent elements of the current formula. For parents facing the income process representative of college graduates, a dollar of financial aid delivered through the current formula is worth \$1.44 in lump-sum tuition discounts. Further, this compensating variation is net of the negative welfare consequences of the disincentives to work and save inherent in the means-testing of financial aid.

Similar results obtain when we isolate the incentives-insurance tradeoff that results from the inclusion of assets in the financial aid formula. For the same college-educated family, with the asset tax removed but the treatment of income the same, the average EFC falls by 44 percent. Replacing the contribution from assets with a fixed fee that keeps the family at the same lifetime expected utility requires 15 percent more aid on average, indicating that the insurance value of the including assets in the formula outweighs the welfare loss due to the disincentives to save. That the family needs to be provided with more aid on average when aid is not linked to assets demonstrates that assets provide incremental insurance beyond that of current income. In a life cycle model, having high assets indicates multiple years of fortunate income draws, and families are made better off by a financial aid formula that offers them less aid in these scenarios if that permits more aid to be provided in the opposite scenarios, when a given year of possibly high income has been preceded by multiple years of less fortunate income draws.

The remainder of the paper is organized as follows. Section 2 describes the key features of the financial aid system and briefly reviews the literature on the relationship between financial aid and household saving. Section 3 develops the stochastic life-cycle model of consumption and labor supply that will be used to simulate work and saving decisions and thus measure the insurance value of financial aid. The main results on the insurance value of financial aid are presented in Section 4, along with sensitivity analyses. We consider the elimination of the asset tax in Section 5. Section 6 discusses directions for further research, and Section 7 concludes.

2 The Financial Aid Tax

Most need-based financial aid is governed by either of two formulas: a “Federal Methodology” set by Congress that determines eligibility for federal financial aid,

as well as institutional aid at some colleges, and an “Institutional Methodology” set by the College Scholarship Service that is used by many selective colleges and universities to determine eligibility for institutionally provided aid. While both require information on income and assets, they differ principally in that the Institutional Methodology considers more sources of income, assets (including home equity), and possible deductions.⁵ The key omission from both formulas is assets held in retirement accounts like 401(k) plans and Individual Retirement Accounts. Our analysis abstracts from this portfolio choice, classifying all assets as covered by the formula to avoid the additional complexity. We also analyze an alternative framework in which no assets are covered by the formula.

Aid awarded under the Federal Methodology is based on information reported on the Free Application for Federal Student Aid (FAFSA). The FAFSA combines information on family structure, income, and assets to generate the Expected Family Contribution (EFC), and financial need is calculated by subtracting the EFC from the student’s cost of attendance at a given school (to which we refer as tuition). The components of the formula are presented in the *EFC Guide* published each year and form the basis of the algorithm used in this paper to calculate financial aid.⁶ Students who are unmarried and sufficiently young apply as dependents of their parents. Those who are older, married, veterans, or have dependents of their own apply under the more favorable status of independent students. As our focus is the parents’ labor supply and saving decisions, we consider students as dependents and simplify the model by zeroing out the student’s contributions. We further simplify the modeling of financial aid by using the formula in the Federal Methodology for combining assets and income to calculate the EFC but a fully general measure of assets and income that is more consistent with the Institutional Methodology.⁷

Following the Federal Methodology, the EFC is obtained by considering a family’s “Adjusted Available Income,” (AAI) which is the sum of “Available Income” (AI)

⁵See National Association of Student Financial Aid Administrators (2017) for background on both methodologies and a discussion of their differences.

⁶The methodology for determining the EFC is found in Part F of Title IV of the Higher Education Act of 1965, as amended, and governs awards for federal Pell grants, subsidized Stafford loans, Perkins loans, federal work-study programs, and other opportunities. For the latest and archived *EFC Guide* publications, see: <http://ifap.ed.gov/ifap/byAwardYear.jsp?type=efcformulaguide>.

⁷In the absence of this simplification, each different type of asset in the formula would necessitate both a state variable and a choice variable in the model below. Other alternatives are possible. For example, all assets could be treated as 529 plan assets that accumulate tax free.

and the “Contribution from Assets” (CA), defined as follows:

$$\begin{aligned}
 AI = & \text{ Adjusted Gross Income -} \\
 & \text{ (Federal Income Tax Paid +} \\
 & \text{ State and Other Tax Allowance +} \\
 & \text{ Social Security Tax Allowance +} \\
 & \text{ Income Protection Allowance +} \\
 & \text{ Employment Expense Allowance)}
 \end{aligned}
 \tag{1}$$

$$CA = \text{Max}(0, 0.12 \cdot (\text{Assets} - \text{Asset Protection Allowance}))
 \tag{2}$$

$$AAI = j * CA + k * AI
 \tag{3}$$

Available Income begins with the parents’ adjusted gross income (AGI) from their tax return and subtracts allowances based on other payments that a family would make in order to earn that income. As implemented below, AGI is just the sum of labor income and asset income, and federal income taxes paid are approximated by a simplified version of the federal tax schedule based on that income. Due to the timing of the filing of taxes and financial aid applications, the EFC is based on income from two years prior to the filing. Below, we use the *EFC Guide* for the 2018 - 2019 academic year and thus income tax schedules from 2016. Marginal tax rates under this schedule range from 10 percent at very low levels of income to 39.6 percent at the highest income levels. Other allowances are made for State and Other Taxes, Social Security Taxes, Employment Expenses, and Income Protection. Each of these other allowances is as specified in the *EFC Guide*, with a state tax allowance of 4.5 percent chosen to reflect the middle of the distribution of state tax rates.⁸

The Contribution from Assets is zero if assets do not exceed the Asset Protection Allowance specified in the *EFC Guide* and 12 percent of any excess of assets over that allowance otherwise. These two components are added together in Equation (3) to obtain AAI. (The parameters j and k are equal to 1 in practice but will be altered in

⁸In 2016, payroll taxes of 7.65 percent were levied on pre-tax labor income up to a maximum taxable earnings of \$118,500, with a 1.45 percent tax solely for Medicare on income above that limit. We assume that the earnings come from employment, rather than self-employment, and thus impose only the employee’s share of the payroll tax.

the simulations below when comparing welfare across different financial aid formulas. Likewise, the parameter F in the next equation is zero in practice.) Given AAI, the EFC is calculated (in 2018 - 2019) as:

$$\begin{aligned}
 EFC = & 0.22 \cdot \text{Min}(16,400, \text{Max}(AAI, -3,409)) + \\
 & 0.25 \cdot \text{Max}(0, \text{Min}(AAI, 20,500) - 16,400) + \\
 & 0.29 \cdot \text{Max}(0, \text{Min}(AAI, 24,700) - 20,500) + \\
 & 0.34 \cdot \text{Max}(0, \text{Min}(AAI, 28,900) - 24,700) + \\
 & 0.40 \cdot \text{Max}(0, \text{Min}(AAI, 33,100) - 28,900) + \\
 & 0.47 \cdot \text{Max}(0, AAI - 33,100) + F
 \end{aligned} \tag{4}$$

The EFC is a piecewise-linear spline in AAI with marginal conversion rates that increase progressively from 22 to 47 percent. Two aspects of these formulas are noteworthy. First, the top marginal conversion rate of 0.47 is reached at a fairly low level of AAI, or \$33,100. Second, while the marginal conversion rates of 12 percent in Equation (2) and 22 to 47 percent in Equation (4) have stayed the same over the years, the various nominal amounts in Equation (4) and the dollar values in the allowances in Equations (1) and (2) have increased over time for inflation. The modeling framework below fixes these dollar values in real terms based on the 2018 - 2019 formula.⁹

Figure 1 shows a contour plot of assets and earnings (labor income) that generate given EFCs, ranging from \$1,000 up to \$73,250, which is the full cost of attendance used in the simulations below. For illustrative purposes, these calculations use the allowances for a married couple with a 2:1 ratio of earnings between the spouses, one child in college, and the age of the older parent being 47 years when the child matriculates. The figure shows that the EFC is monotonically increasing in both assets and earnings. The intercepts on the horizontal (vertical) axis show the earnings (asset) levels associated with the contour's EFC assuming that the family's asset (earnings) level is zero. The EFC remains at zero for combinations of earnings and assets toward

⁹The analysis below also incorporates the Simplified Needs Analysis in the Federal Methodology. A dependent student qualifies for this formula if anyone in the household receives means-tested assistance, such as Medicaid or Supplemental Security Income. An asset limit of \$2,000 (\$3,000) is used for single (married) parents in the modeling framework below. For families that meet this criterion, the EFC is zero (bypassing Equation (4)) if income is less than \$25,000, and the contribution from assets is zero (bypassing Equation (2)) if income is between \$25,000 and \$50,000.

the lower left corner of the figure that do not exceed the various allowances. At a college using the Federal Methodology to allocate aid, any EFC that falls below the costs of attendance would make the student eligible for financial aid, potentially up to the difference between those costs and the EFC if the institution committed to meet full demonstrated need. For the highest values of earnings and assets toward the upper right corner of the figure, the EFCs exceed the costs of attendance even at the most expensive colleges, resulting in no financial aid.¹⁰

Figure 2 presents the implied marginal tax rates on earnings inherent in the EFC amounts in Figure 1. The slope of each curve is the incremental change in the EFC for an incremental change in earnings, holding the asset level constant at the value specified in the legend. At low levels of earnings and assets, at which the EFC is zero, the implied marginal tax rates are also zero. Once EFCs become positive, the implied marginal tax rate on earnings can be as high as 40 percent. The key factors driving the implied marginal tax rate are the marginal conversion rates of AAI to EFC in Equation (4), which rise with earnings up to 0.47, and the income sensitivity of the allowances in Equation (1). The latter include state income taxes (assumed to be proportional to income), payroll taxes (which fall from 7.65 to 1.45 percent at the maximum taxable earnings level), and federal income taxes (which are progressive and thus rise with earnings). At higher earnings levels, where the allowances might have a combined sensitivity to income of about 40 percent, the slopes of the curves would be $0.47 \cdot (1 - 0.4) = 0.282$, or about 28 percent. Higher or lower rates may occur at lower earnings levels for which the the income sensitivity of the allowances may be lower but the marginal conversion rates are also lower.¹¹

Figure 3 shows the implied marginal tax rates on assets inherent in the EFC formula. Analogous to Figure 2, each curve in Figure 3 holds earnings constant at the value specified in the legend and then calculates the incremental change in the EFC for an incremental change in assets. Over most of the figure, the implied marginal

¹⁰The EFC is divided by the number of children in college, so with more children in college (and adjusting for the impact of more children on the income allowance) there would be the possibility of financial aid even at these earnings and asset levels.

¹¹These implied marginal tax rates are in addition to the marginal income tax rates from the payroll tax, federal income taxes, and state income taxes, implying potentially high combined tax rates on labor income during the years in which parents have children in college. This possibility is noted but not explored in Feldstein (1995). Using cross-sectional regressions of actual financial aid awards, Dick and Edlin (1997) estimate lower income-sensitivity of financial aid than these theoretical predictions, noting that actual awards are typically not as progressive as the formulas imply.

tax rate on assets is about 6.6 percent. This “asset tax” comes from two sources. The first is in Equations (2) and (4), in which 12 percent of assets are available and are converted at rates up to 47 percent: $0.47 \cdot 0.12 = 0.0564$.¹² The second is in Equations (1) and (4), in which the assets generate income, net of taxes on that income at the federal and state levels, and then are converted at rates up to 47 percent. With a 3 percent rate of return and a 30 percent combined marginal income tax rate on asset income, this would yield an additional $0.47 \cdot 0.03 \cdot (1 - 0.3) = 0.0099$. Combining these two components yields 6.6 percent.¹³

The implied marginal tax rate of 6.6 percent applies in each successive year of college attendance to the remaining assets. Thus, a dollar of assets at the start of college is reduced by $0.066 [1 + (1 - 0.066) + (1 - 0.066)^2 + (1 - 0.066)^3] = 0.239$, or about 24 percent over four years in college. Thus, the financial aid formula levies a substantial tax on assets over a broad range of earnings and asset combinations.¹⁴

These implied marginal tax rates on assets are the impetus for the empirical literature that has estimated whether households respond to the asset tax by saving less. The literature starts with Feldstein (1995), who estimated a reduction of about 50 percent in asset accumulation due to the financial aid tax. His estimation sample was a cross-section of 161 households in the Survey of Consumer Finances 1986. Long (2004) argues that a household’s estimate of the implicit tax on assets that would discourage saving is more complicated than suggested in Feldstein (1995), noting that it depends on factors such as the likelihood of children going to college, the expected cost of college (since the marginal tax rate is zero if the EFC exceeds the cost of attendance), and the possibility that the college does not meet all need, in which case an additional dollar of assets will reduce unmet need rather than financial aid. His methodology generates smaller taxes at the margin and no correlation between those marginal tax rates and asset accumulation. Later studies by Monks (2004) using the National Longitudinal Study of Youth and Reyes (2008) using the Panel Study of Income Dynamics find weak evidence consistent with lower asset accumulation, but

¹²The Institutional Methodology assesses assets at rates between 3 and 5 percent. See National Association of Student Financial Aid Administrators (2017).

¹³That the curves in Figures 2 and 3 are nearly linear and parallel indicates why the contours in Figure 1 are also nearly parallel, with a slope given by the ratio of the slope of the earnings curves to that of the asset curves.

¹⁴These estimates for the asset tax are broadly consistent with those of Dick and Edlin (1997), who estimated marginal asset tax rates of up to 30 percent using cross-sectional data from the 1987 National Postsecondary Student Aid Survey.

at magnitudes much less than Feldstein (1995).¹⁵

The implied marginal tax rates on both earnings and assets in Figures 2 and 3 are also the source of the insurance value of financial aid. As was noted by Eaton and Rosen (1980), even a simple proportional income tax in which the proceeds are redistributed as a lump sum will raise welfare when income is uncertain. As with the prior literature on the disincentive effects of the asset tax, the degree of insurance in the financial aid formula depends on whether the college commits to provide financial aid equal to the difference between the costs of attendance and the EFC. This is assumed in the analysis below and is true of the most well funded colleges and universities, for which the analysis in general is most applicable because these institutions often have the highest cost of attendance and the most generous financial aid programs. However, this issue is not as critical for the insurance value of financial aid as it is for the disincentives of the asset tax. Even if there may be some income or asset ranges over which an institution may not boost financial aid dollar-for-dollar with demonstrated need, generating lower marginal tax rates than in Figures 2 and 3, the insurance provided on the inframarginal need is still present.

3 Stochastic Life-Cycle Model of Consumption and Labor Supply

This section presents a stochastic, life-cycle model of consumption and labor supply in which the traditional retirement motive for saving is augmented by a precautionary motive to save against income uncertainty and a potential need to pre-fund a child's college education. We begin by specifying the model fully, followed by a discussion of the solution method and welfare comparisons, and conclude with a justification for the parameters chosen.

¹⁵The implicit tax on income in the financial aid formula has received less consideration to date as a source of economic distortions. Handwerker (2011) uses the Health and Retirement Study to show that parents delay retirement while paying for a child's college education. She finds little evidence that paying for a child's education has any impact on work intensity for those who are working. In more recent work, Braga and Malkova (2020) show that mothers of college-age children decreased their annual hours of work after the start of a generous merit aid program, while fathers did not adjust their labor supply. The response comes from reductions in hours rather than changes in employment status, and there is no adjustment for fathers or for mothers of children who did not attend college. Note, however, that because this is a merit aid program, the response is an income effect rather than a reaction to the financial aid tax studied here.

3.1 Model Specification

The basic structure of the model is that in each period of life, s , the family chooses values of consumption, C_s , and labor, L_s , as functions of the two state variables in the model, assets, A_s , and labor income from fulltime work, Y_s . The family's value function in period t , $V_t(A_t, Y_t)$, is defined as:

$$V_t(A_t, Y_t) \equiv \max_{\{C_s, L_s\}} E \sum_{s=t}^T \beta^{s-t} (u(C_s) + \theta v(L_s)) \quad (5a)$$

$$u(C) = \frac{C^{1-\gamma}}{1-\gamma} \quad (5b)$$

$$v(L) = \frac{(\bar{L} - L)^{1-\mu}}{1-\mu} \quad (5c)$$

$$\hat{Y}_s = Y_s \left(\frac{L_s}{L^F} \right) \quad (5d)$$

$$X_s = A_s + \hat{Y}_s - h(\hat{Y}_s) - z_s(A_s, \hat{Y}_s) \quad (5e)$$

$$A_{s+1} = (1+r)(X_s - C_s) - g(\hat{Y}_s + r(X_s - C_s)) \quad (5f)$$

$$A_s \geq 0, \forall s \quad (5g)$$

$$L_s^{\min} \leq L_s \leq L_s^{\max}, \forall s \quad (5h)$$

The value function is equal to the sum of the expected utility of consumption and leisure in each period from the current period t to the final period T , discounted by a factor of β each period.¹⁶ The discount factor governs the utility tradeoff across periods – values closer to 1 reflect greater patience. The utility of consumption each period shown in Equation (5b) is assumed to take the Constant Relative Risk Aversion (CRRA) form, where γ is the coefficient of relative risk aversion. With a utility function such as CRRA that has a convex marginal utility function (i.e. $u'''(C) > 0$), there is a precautionary motive for saving, and greater uncertainty in the earnings process will induce greater saving.¹⁷ In Equation (5c), leisure is defined

¹⁶In addition to the additive separability of consumption and leisure, the specification for the value function and the within-period utility makes three simplifying assumptions. The first is that there is no adjustment to the argument of the utility function for the size of the household, even after the child has left for college. The second is that there is no mortality risk and thus no accidental bequests. Third, there is no planned bequest motive.

¹⁷The use of the CRRA utility function is standard in both the empirical and theoretical literature

as the difference between a time endowment, \bar{L} , and the amount of labor supplied, L . The parameter θ governs the relative weight placed on the utilities of consumption and leisure each period. The functional form for the utility of leisure is the same as for consumption, with curvature parameter, μ .

Equation (5d) defines labor income, \hat{Y}_s , as a function of fulltime income, Y_s , and the labor choice, L_s , scaled by an amount, L^F , such that if the family worked exactly, L^F , its labor income would be Y_s . We can think of the ratio $\left(\frac{L_s}{L^F}\right)$ as the fraction of a fulltime year worked or, alternatively, of the ratio $\left(\frac{Y_s}{L^F}\right)$ as an annual wage at which the family is compensated for each unit of labor, L_s . Equation (5e) defines the concept of “cash on hand” that is available to finance consumption and income taxes each period. To obtain cash on hand, X_s , assets are augmented by labor income but reduced by payroll taxes, $h(\hat{Y}_s)$, and costs of college attendance, $z_s(A_s, \hat{Y}_s)$, which may depend on assets and labor income through the financial aid formula described in Equations (1) to (4).¹⁸ The term, $z_s(A_s, \hat{Y}_s)$, can also incorporate the impact on cash on hand of loans taken out to fund educational expenses.¹⁹

Equation (5f) shows how assets accumulate from one period to the next. Cash on hand is used to finance both consumption and the income taxes, $g(\hat{Y}_s + r(X_s - C_s))$, that are due based on capital income and labor income. To avoid the complexity of an additional state and choice variable, the portfolio decision is restricted to a single riskless asset paying a return, r , each period. Thus, the amount of saving is $X_s - C_s$, and capital income is just $r(X_s - C_s)$. The family’s taxes are calculated based on the 2016 tax schedule for a married couple with one child who does not itemize deductions and receives all capital income as interest or dividends rather than capital gains. Payroll taxes are assumed to be paid as the labor income is earned, prior to the consumption decision each period. Since income taxes depend on capital income and

on precautionary saving. CRRA utility means that a consumer remains equally willing to engage in gambles over a constant proportion of current wealth as wealth increases. An alternative, and perhaps more realistic assumption, might be that the consumer will accept larger proportional risks as wealth increases. See Kimball (1990) for a discussion and derivation of the key results for precautionary saving.

¹⁸The payroll tax includes coverage for disability insurance, but the risk of disability is not modeled here. See Chandra and Samwick (2009) for a life cycle model that incorporates the risk of disability.

¹⁹We do not include the income tax deduction for tuition and fees, for which families can claim a deduction of the lesser of tuition and fees or \$4,000 (\$2,000) if their modified AGI is less than \$130,000 (\$160,000), respectively, for those married filing jointly (or half those thresholds for single filers). Hoxby and Bulman (2016) find no evidence that the post-secondary tax deduction affects college-going behavior or other aspects of college financing.

thus the outcome of the consumption decision during the period, they are assumed to be paid at the end of the period.

The last two elements of Equation (5) are the constraints on the optimal choices. Equation (5g) is the liquidity constraint, which requires assets, A_s , to be positive in each period. Other than a student loan which we discuss below, the family cannot borrow against future income to finance current consumption. This is a simplification that nonetheless acknowledges the credit constraints that prevent families from borrowing too heavily against future income outside of a secured or collateralized relationship.²⁰ Equation (5h) imposes minimum, $L_s^{\min} > 0$, and maximum, $L_s^{\max} < \bar{L}$, constraints on labor supply. In retirement, we normalize $L_s = L^F = L_s^{\min} = L_s^{\max}$.

The processes that describe the uncertainty in and evolution of fulltime income are as follows.

Before retirement:

$$\ln(Y_s) = \ln(P_s) + u_s \quad (6a)$$

$$\ln(P_{s+1}) = \nu_s + \ln(P_s) \quad (6b)$$

$$u_{s+1} = \rho \cdot u_s + \varepsilon_{s+1} \quad (6c)$$

$$\varepsilon_{s+1} \sim i.i.d. N(0, \sigma^2) \quad (6d)$$

At retirement:

$$Y_{s+1} = (1 - \chi) \cdot f\left(\frac{1}{35} \sum_{t=s-34}^s \omega_t q(P_t)\right) + \chi \cdot f(q(Y_s)) \quad (7)$$

After retirement:

$$Y_{s+1} = Y_s \quad (8)$$

Prior to retirement, the natural log of fulltime income is equal to the natural log of permanent income, P_s , plus a shock to income, u_s , that follows an AR(1) process. As discussed in Section 3.4.2 below, permanent income follows an estimated age-earnings relationship and may be augmented by an annual rate of ν_s in year s . The innovations to that AR(1) process are assumed to be independently and identically drawn from

²⁰See Hurst and Willen (2007) for an analysis of consumption with a richer modeling of credit constraints.

a normal distribution with mean zero and variance σ^2 .²¹ In this model, the parents retire at a planned date that is known from the beginning of the working life. By assumption, there is also no impact of labor supply, L_s , on any current or future value of fulltime income.

As shown in Equation (7), at retirement, income is calculated as a weighted average of the Social Security benefit, $f(\cdot)$, applied to two earnings measures. The first is an approximation of the average of the last 35 years of permanent income, P_t , where the weights, ω_t , account for the growth in the national average wage, and the function, $q(\cdot)$, accounts for the upper limit of the Social Security maximum taxable earnings. This term is deterministic, transmitting none of the uncertainty in pre-retirement earnings into retirement. The second measure applies the Social Security benefit formula to fulltime income in the year prior to retirement, Y_s . This term transmits the cumulative AR(1) shock, u_s , during the working years into retirement. Income in the year of retirement is a weighted average, with weight χ , of these two terms.²² Employer-provided pensions are assumed not to be defined benefit plans and thus the balances would be included in A_s . After retirement, income is constant in real terms, as shown in Equation (8).

3.2 Overview of Solution Method

The solution method for stochastic optimization problems with multiple state and control variables is discussed in detail in Carroll (2019). As a dynamic programming problem, Equation (5a) can be written recursively for any period t as:

$$V_t(A_t, Y_t) \equiv \max_{\{C_t, L_t\}} u(C_t) + \theta v(L_t) + \beta E_t[V_{t+1}(A_{t+1}, Y_{t+1})] \quad (9)$$

The optimization proceeds backwards through time, from period T to the first period, generating a series of rules for consumption and leisure that determine optimal consumption and leisure as a function of the state variables in that period. In periods

²¹In the simulations, the mean of the shock to the level (not log) of income is normalized to be one in all periods.

²²The estimation of this parameter is described in Section 3.4.2 below. These modeling choices for income are designed to avoid additional state variables, like average income, that would allow a more exact calculation of Social Security benefits at the cost of additional complexity in the model. A richer model could include the risk of involuntary retirement due to health or other reasons and a choice over the retirement age based on economic factors.

before retirement that have both a labor supply choice and a consumption choice, the period-by-period solution to Equation (9) can be found by solving a system of two first order conditions (one for C_t and one for L_t) and the constraints in Equations (5g) to (5h). Thus, it is a system of 5 equations in 5 variables (the two choice variables plus the three Lagrange multipliers). To simplify the solution, we break each within-period problem into two sub-period problems, with the leisure choice occurring in the first sub-period, after the income shock is realized, and the consumption choice occurring in the second sub-period. A full derivation of the solution, including additional simplifications made to the optimization problem, is presented in Appendix A.

The first-order condition for the labor supply choice in the first sub-period is (see Appendix Equation (A.6)):

$$\begin{aligned} \theta v'(L_t) + u'(C_t^*) \left(\frac{Y_t}{L^F} \right) \left(1 - h'(\hat{Y}_t) - z_t^Y(A_t, \hat{Y}_t) - g'(\hat{Y}_t) \right) \\ = \mu^{\max} - \mu^{\min} \end{aligned} \quad (10)$$

The first term is the marginal utility of an additional unit of labor supplied, with $v'(L_t) < 0$. At an interior optimum, this disutility must be equal in magnitude to the gain in utility that occurs in the second sub-period due to the higher consumption made possible by this additional unit of labor supplied. This utility gain has three components – the marginal utility of another dollar of cash on hand to start the second sub-period, $u'(C^*)$, where C^* is the optimal consumption level; the pre-tax “wage” from fulltime work, $\left(\frac{Y_t}{L^F}\right)$; and one minus the marginal tax rates on labor income due to the payroll tax, financial aid formula, and income tax, $\left(1 - h'(\hat{Y}_t) - z_t^Y(A_t, \hat{Y}_t) - g'(\hat{Y}_t)\right)$. Note that the term, $z_t^Y(A_t, \hat{Y}_t)$, indicates that there is a financial aid tax on earning income while the child is in college. The higher is this financial aid tax, the lower the value of L_t . If there is no interior optimum at which the left-hand side of Equation (10) is zero, then labor supply will be at its maximum or minimum and the left hand side will be equal to the non-zero Lagrange multiplier on whichever constraint is binding, μ^{\max} or $-\mu^{\min}$.

The first-order condition for the consumption choice in the second sub-period is (see Appendix Equation (A.8)):

$$u'(C_t) = \beta(1 + r(1 - g'(Y_t))) \cdot \left(E_t \left[u'(C_{t+1}^*) \left(1 - z_{t+1}^A \left(A_{t+1}, L_{t+1}^* \left(\frac{Y_{t+1}}{L^F} \right) \right) \right) \right] + \lambda \right) \quad (11)$$

At an interior optimum (i.e. one in which the liquidity constraint in Equation (5g) does not bind and thus the Lagrange multiplier $\lambda = 0$), the marginal utility of consumption in period t is equal to the discounted expected marginal utility of consumption in period $t + 1$, accounting for the effects of both the after-tax interest rate in period t and the financial aid tax on assets in period $t + 1$, $\left(1 - z_{t+1}^A \left(A_{t+1}, L_{t+1}^* \left(\frac{Y_{t+1}}{L^F} \right) \right) \right)$, where L^* is the optimal labor supply choice. The higher is the financial aid tax, the lower is this term, and thus the higher is the value of C_t at which the first-order condition will hold. If instead the liquidity constraint is binding and $\lambda > 0$, the marginal utility of consumption will be higher, as consumption is constrained by cash-on-hand to be below the value such that its marginal utility is equal to the discounted expected marginal utility of consumption in the next period.

The solution begins in the last period of life, T , when the problem is trivial because the family simply consumes all of its assets and after-tax income, yielding an optimal value for C_T as a function of the state variables, A_T and Y_T . Equations (11) and (10) are then applied in each successively prior sub-period to obtain the optimal consumption and labor supply functions of the relevant state variables until the beginning of the life cycle. Once the optimal consumption and labor supply rules have been obtained, the model can be simulated forward by specifying initial values of the state variables, drawing random shocks to earnings, and applying the labor supply and consumption rules to generate distributions of asset balances, income, EFCs, and other outcomes in each successive period.

3.3 Welfare Comparisons

In the simulations below, the model is evaluated using the distributions generated based on 1,000 independent random draws of the age-earnings profile.²³ The key out-

²³The closest antecedent to this approach in the financial aid literature is the model of Dick, Edlin and Emch (2003), who estimate preference parameters for education and saving to determine the asset reductions due to the financial aid system and simulate the asset and welfare changes that would result from changes to that system. The saving framework in that paper is based on a

come of the model is the expected value of $V_1(A_1, Y_1)$, computed as the average value of this term across the 1,000 age-earnings profiles and a starting asset value of zero at the beginning of the work life. This expected value is a metric by which different financial aid systems can be compared, as those with higher values of $E[V_1(A_1, Y_1)]$ are the ones in which the family is better off.

We consider variations to the financial aid formula based on the parameters $\{j, k, F\}$ in Equations (3) and (4) above:

- j : The fraction of the Contribution from Assets (CA) in the Adjusted Available Income (AAI) formula that is applied.
- k : The fraction of Available Income (AI) in the AAI formula that is applied.
- F : An additional surcharge, in dollars, that is added as a fixed fee to the EFC, unrelated to AAI.

In the current formula, $\{j, k, F\} = \{1, 1, 0\}$. We define the expected (indirect) utility function,

$$W(j, k, F) \equiv E[V_1(A_1, Y_1) | j, k, F] \quad (12)$$

and the expected average costs of attendance as $\bar{z}(j, k, F)$. Note that when $\{j, k, F\} = \{0, 0, \bar{z}(1, 1, 0)\}$, the amount of financial aid is constant at the average amount of financial aid realized under the current formula. We refer to this set of parameters as “Revenue Equivalent” to the current formula.

Our standard comparison sets $j = k = 0$, and we solve for δ such that:

$$W(1, 1, 0) = W(0, 0, \bar{z}(1, 1, 0) - \delta) \quad (13)$$

In this comparison, the income- and asset-contingent components of the financial aid formula are removed and replaced with a lump-sum discount. There is no financial aid tax on saving or labor supply, but neither is there any insurance provided against low assets or income. In this context, δ is a compensating variation, in that adding this amount to the average financial aid under the current formula (i.e. subtracting it from the cost of attending in Equation (13)) restores the family to its expected utility under the current formula. If $\delta > 0$, then the compensating variation is positive – the family needs to be compensated for having $j = k = 0$. This compensation is paid

non-stochastic life-cycle model and thus cannot measure the insurance value of financial aid.

by the college in the form of higher average financial aid costs. We refer to this set of parameters as “Utility Equivalent” to the current formula.

In the analyses below, we also consider variations in the formula in which only CA is eliminated, $j = 0$. This removes the financial aid tax on saving while retaining the formula’s implicit tax on income (including income from assets). We then set either $F > 0$ or $k > 1$ to restore the family’s value of W and compare average financial aid to the current formula. An alternative financial aid formula is preferable if it achieves the same expected utility for the family at a lower average financial aid cost to the college.

3.4 Model Parameterization

In this subsection, we discuss our baseline parameters in Equations (5) - (8), with sensitivity analysis presented in Section 4 below.

3.4.1 Life Cycle Framework and College Costs

We assume that economic life lasts 60 periods, with retirement in the 40th period, corresponding roughly to an adult life of ages 25 to 85 and retirement at age 65. Initial assets, A_1 , are zero. The child is born in the fourth period of economic life and starts college in the 22nd period. (For robustness, we consider a family in which the couple is 10 years older when they become parents.) The costs of attendance, which serve as the maximum value of the EFC, are \$73,250, which approximates the total costs of attendance at highly selective institutions in 2018 - 2019, the year we use for our EFC formula.²⁴ In Equation (3), $j = k = 1$, and in Equation (4), $F = 0$, in the baseline, to implement the EFC formula as specified. In addition to financial aid offered through the EFC formula, the family is assumed to be able to take out a loan of \$10,000 per year of college at the riskfree rate of $r = 3\%$ to be repaid over a period of 20 years.

²⁴Costs of attendance include tuition, fees, room and board, books and supplies, and estimates of travel and other expenses for residential students. Among the Top 10 National Universities as ranked by *US News and World Report*, costs of attendance that year ranged from a low of \$70,010 at Princeton University to a high of \$77,331 at the University of Chicago.

3.4.2 Earnings Process

The earnings process in Equation (6) combines a deterministic profile for permanent (fulltime) income and a persistent stochastic shock. For the deterministic profile, we utilize a cross-section of the American Community Survey (ACS) from 2017 and, following Murphy and Welch (1990), estimate fourth-order polynomials in age for the logarithm of earnings by education group. The advantages of the ACS are that earnings are topcoded at the relatively high value of \$1 million, education categories can be distinguished beyond college completion, and the sample size is large.²⁵ The regression coefficients are shown in Table 1 for five different educational groups – Less than High School, High School Diploma, Some College, (Four-Year) College Degree, and Advanced Degree. The sample is couples with at least one member working fulltime and between the ages of 25 to 64, and the education group pertains to that spouse or the more educated spouse if both are working. If both spouses are working fulltime, between the ages of 25 to 64, and equally educated, the couple’s age is that of the older spouse. The age-earning profiles, corresponding to P_s in Equation (6b), are graphed in Figure 4. To allow for additional heterogeneity in earnings, we add a high-growth version of the Advanced Degree profile that augments earnings growth by an additional 1% per year (i.e., we set ν_s in Equation (6b) to 0.01). In the simulations of the model, results are presented for each of these six groups distinguished by their average earnings profiles.

The stochastic elements of the earnings process are an AR(1) coefficient of $\rho = 0.95$ and a standard deviation of the annual earnings shock of $\sigma = 0.15$. These parameters are standard in the microeconomic literature on consumption and savings and date back to the pioneering work by Hubbard, Skinner and Zeldes (1994). Using a measure of income that includes after-tax labor earnings for both members of a couple plus unemployment insurance, they find $\rho = 0.95$ for the three education groups (Less than 12, 12 to 15, or More than 16 years) in their analysis and roughly $\sigma = 0.15$ for the middle group, with variance decreasing with education. They also include an additional transitory shock, which we exclude here under the conservative assumption that much of the transitory shock is measurement error.

Subsequent literature on empirical earnings processes has used larger datasets and considered more complex stochastic components. An important recent example

²⁵See Gibson and Gibson (2017) for a similar use of the ACS.

is Guvenen et al. (2015), who use a 10% sample of the Social Security Administration’s Master Earnings File to study male earnings from 1978 to 2010. Their estimate of the AR(1) coefficient in the analogous model to Equation (6) is $\rho = 0.962$, with a standard deviation of $\sigma = 0.174$ (their Table III, Column (7)). Another example is DeBacker et al. (2013), who use a panel of tax returns from 1987 to 2009 and can thus examine pre-tax and after-tax household income in addition to male labor earnings. As our modeling framework focuses on couples and deducts payroll and income taxes separately, the closest analogue is pre-tax household income, for which they estimate parameters of $\rho = 0.967$ and $\sigma = 0.164$.²⁶ Thus, the stochastic elements of the AR(1) earnings process are consistent with subsequent literature.²⁷

The final element of the earnings process is the calculation of retirement income. Recall from Equation (7) that income in the year of retirement is a weighted average of two approximations of the Social Security benefit formula. The first applies the formula to the history of permanent income, approximating the benefit level for someone earning the average income at each age. The second applies the formula to the immediate pre-retirement value of fulltime income. The weighting parameter, χ , determines what the share of the final pre-retirement shock is transmitted to retirement income. Since the earnings in the model refer to a couple, we divide the couple’s earnings 2:1 for the higher-earning spouse when calculating retirement benefits, matching the sample average in the ACS data used to generate the profiles. We assume that the national average wage used in the benefit calculation grows at 1 percent per year. The benefit for the lower-earning spouse is constrained to be at least as large as the

²⁶Note that the counterpart to pre-tax household income in our model is \hat{Y}_s , whereas the σ parameter corresponds to fulltime income, Y_s . In the simulated data from the model, discussed below, we observe that the cross-sectional standard deviation of Y_s is higher than that of \hat{Y}_s at all ages. Families consume less leisure when they have low income, partially offsetting negative shocks to fulltime earnings. Thus, using the estimated standard deviation on pre-tax household income likely understates the variation in fulltime income.

²⁷Much of the recent literature in this area has studied the ways in which the standard AR(1) process with normally distributed innovations to log-earnings does not fit the data. For example, Guvenen et al. (2015) show that the distribution of earnings growth rates has negative skewness and, most importantly, substantial kurtosis relative to the normal distribution. Continued work in Guvenen et al. (2019) focuses on the variation across individuals in the fraction of a working lifetime spent non-employed and how to incorporate that variation, and its persistence, into an income process. While these factors are not modeled here, it is important to note that other studies that have considered them, such as Golosov, Troshkin and Tsyvinski (2016), have found that the case for redistributive taxation based on insuring against shocks that display negative skewness and excess kurtosis is *stronger* than under the more restrictive model used here.

dependent spouse benefit of 50% of the benefit of the higher-earning spouse.²⁸ For robustness, we also consider the same total income in a one-earner couple.

The simulations of the model discussed below yield, ex post, full earnings histories and thus the ability to calculate actual Social Security benefits for each of the 1,000 random draws of the age-earnings profile. To determine the value of χ , we regress the couple's benefits applied to that earnings draw on the the benefit formula applied to pre-retirement fulltime income, Y_s (i.e. the second term in Equation (7)). A fixed effects (by education level) regression yields an estimated sensitivity of 0.35, which we use as our estimate of χ . This sensitivity is larger than just $\frac{1}{35}$, any single year's weight in the average indexed earnings calculation, because fulltime income in the last year of work contains information about many prior shocks (and thus many prior years of income) in the AR(1) process. Median replacement rates in the simulated data are declining with the level of the couple's education, from 71 percent for the lowest educational group, to 61 percent for the College Degree group, to 32 percent for the highest educational group.

3.4.3 Utility Function

We choose a baseline value of the coefficient of relative risk aversion, γ , of 3. This value is both common in the microeconomic consumption literature (see, for example, Hubbard, Skinner and Zeldes (1994)) and consistent with a recent meta-analysis of 169 published studies by Havránek (2015) that estimated the elasticity of intertemporal substitution (the inverse of the coefficient of relative risk aversion) to be around 1/3 when correcting for reporting bias. Using income and wealth data from the Panel Study of Income Dynamics (PSID) and the Survey of Consumer Finances (SCF), Cagetti (2003) shows that median wealth holdings by age and education are consistent with risk aversion parameters that are higher than 3, and for some groups, higher than 4. Since a coefficient of relative risk aversion of 3 is higher than estimates found through direct elicitation of the preference parameters from experiments over hypothetical lotteries, we consider lower values for risk aversion in the sensitivity analysis in Section 4.²⁹

We normalize fulltime work to be $L^F = 40$, representing hours per week for one fulltime worker, even when L refers to the couple's labor supply which might include

²⁸The division of income within the couple also plays a small role in the calculation of the EFC.

²⁹For example, Andersen et al. (2008) estimate a value of 0.74.

a second earner. In the model with variable labor supply, we set $L_s^{\min} = 35$ and $L_s^{\max} = 45$ in Equation (5h), allowing for labor supply to change by $\pm 12.5\%$ in each period. The curvature of the utility function for leisure, μ , is set to 3, matching the value for γ . This parameter is related to the Frisch elasticity of labor supply, $\eta = \frac{v'}{L*v''}$. In our formulation, $\eta = \left(\frac{\bar{L}-L}{L}\right) \left(\frac{1}{\mu}\right)$.³⁰ With values of $\bar{L} = 168$ total hours per week and evaluating η at $L = L^F = 40$, a value of $\mu = 3$ corresponds to a value of η of about 1.³¹

The discount factor is a primary determinant of the family's saving choices. We use wealth and income data from the SCF 2016 to calibrate the values of this parameter. Sponsored by the Federal Reserve Board in cooperation with the Department of the Treasury, the SCF is a triennial survey of U.S. families that collects comprehensive data on assets, liabilities, income, and demographic characteristics. We use the SCF data on assets, liabilities, and income by age and education group to construct empirical moments of the wealth-to-income distribution and estimate the values of β that best match simulated moments to their empirical counterparts.

Specifically, we impose the same sample restrictions and distinguish the same five education levels in the SCF as in the ACS. For each education group, we obtain the ratio of average wealth to average labor income (including self-employment income) in four age ranges: 25 to 34, 35 to 44, 45 to 54, and 55 to 64. Following Carroll and Samwick (1997), we define wealth as net worth excluding housing equity and the value of personally owned businesses and construct these variables using the formulas in Bricker et al. (2017). We then simulate our life-cycle model under baseline parameters for a family without college expenses and with labor supply constrained to be $L^F = 40$. We calculate the ratio of average assets to average labor income in the model's 5th, 15th, 25th, and 35th periods, corresponding to the midpoint of the age ranges noted above. For each earnings profile, we find the value of β that minimizes the sum of squared deviations between the four simulated ratios in the model and the four empirical ratios from the SCF.

The results are shown in Table 2 and in Figure 5 for each of the five education groups. In the table, each pair of numbers shows the empirical and simulated ratios

³⁰In the macroeconomic literature on consumption and labor supply, the more typical formulation of the disutility of labor is $v(L) = -\frac{L^{1+\frac{1}{\eta}}}{1+\frac{1}{\eta}}$.

³¹This value is intermediate between the micro- and macro- estimated elasticities found in the literature. See Reichling and Whalen (2012) for a review and Peterman (2016) for a discussion of the differences across estimates.

of average assets to average earnings for a given education group at each age. In the figure, the dots are the empirical moments and the curve shows the simulated moments from the model. The estimated values of β increase monotonically with education, from a low of 0.930 for the households without a high school diploma to 0.992 for the families with advanced degrees. A value of $\beta = (1 + r)^{-1} \approx 0.97$ corresponds to a family who would keep consumption level over time in the absence of precautionary, retirement, or college motives for saving (if income growth were level and taxes were zero).

The result that more educated families exhibit greater patience, and thus a more rapidly rising asset-to-income profile, is standard in the literature. The Goodness of Fit statistic is the sum of squared deviations of the mean ratios at the four age groups being minimized. As is also evident from the figure, this statistic improves markedly with education level, suggesting that the modeling framework is more appropriate for more educated families.³²

We treat θ , which governs the tradeoff between consumption and leisure within a period, as a parameter to be calibrated. For each set of values of the other parameters in the model, we choose θ so that the average value of labor supply during the working years (for a family not facing any college costs) is equal to L^F . For the sixth group, the Advanced Degree group with 1% higher annual permanent income growth, we use the same value of β as for the Advanced Degree group and calibrate a separate value of θ . To get couples with higher earnings profiles to work no more on average than couples with lower earnings profiles, θ falls steadily with education.

4 Model Results

4.1 Consumption and Labor Supply

Figure 6 summarizes the age profiles of average earnings, consumption, assets, and college costs for a family facing the earnings process for a college graduate and the baseline parameters described above. The top panel holds labor supply fixed at L^F , and the bottom panel allows labor supply to be chosen optimally between L_s^{\min} and L_s^{\max} in each pre-retirement period. In both cases, consumption is smoothed from

³²For example, the model's assumption of a life expectancy of 85 years might be more accurate for more educated families.

working years into retirement – consumption is below earnings before retirement and then above retirement benefits thereafter. Asset accumulation makes this possible, as assets rise to a peak of roughly 7 times pre-retirement earnings on the eve of retirement. Assets are spent over the years in which the child is in college and then again, to zero, in the retirement period. Average consumption rises as retirement approaches, due to precautionary saving motives that decline over the working years. Average consumption decreases only slightly during the college years.³³

In the bottom panel, all of the profiles are affected by the ability of the couple to vary its labor supply. The flexibility to increase labor supply later in the working life in response to adverse earnings draws early in the working life allows the family to have higher consumption in those early years. With higher early consumption, the family accumulates fewer assets prior to the college years and prior to retirement. By itself, lower asset accumulation will lower the EFC the family pays for the child to attend college. The EFC is also lowered by the ability of the couple to reduce its labor supply during the years when labor income will be included in the EFC calculation. This reduction in labor supply is evident in the age-earnings profile in the bottom panel of Figure 6 and is presented in more detail in Table 3. The table shows average labor supply, relative to $L^F = 40$, for each age-earnings profile in the periods before, during, and after the college years.

The top panel shows the results using the current financial aid formula, and the bottom panel shows the results using the “Revenue Equivalent” parameters, in which tuition is simply discounted by an amount equal to the average financial aid under the current formula. Note that the earnings draws of fulltime income, $\{Y_s\}$, are the same in both panels. The bottom panel shows that when the cost of college is not linked to income or assets, the couples facing all of the earnings profiles work less during the college years than in the years prior to college. For the more educated families, labor supply continues to decline after the college years, while for the less educated families, labor supply is somewhat higher after the college years but not quite as high as during the pre-college years. Note from Figure 4 that the decline in average earnings near retirement is less pronounced for these three groups.

The top panel shows several differences in average labor supply relative to the

³³The smoothness of consumption around the years of college attendance is consistent with the evidence in Souleles (2000), who shows in the Consumer Expenditure Survey that households’ non-education consumption does not decrease over the academic year in proportion to college expenditures in the fall. That is, at least over short horizons, the household is able to smooth consumption.

bottom panel. First, there is less labor supply during college, when earnings are subject to the financial aid tax on income, with higher earnings profiles generally cutting back the most. The percentage of families constrained by $L^{min} = 35$ during the college years (not shown) varies between 42 to 81 percent across the earnings profiles. Second, labor supply after college is higher than in the bottom panel and higher than during the college years. The changes are again generally larger for higher earnings profiles. Third, there is less labor supplied prior to the college years under the current formula. This is due both to the financial aid tax on assets, which taxes the savings from additional labor prior to college, and to the reduction in the precautionary motive to supply labor prior to college.

4.2 The Financial Aid Tax

Much of the original literature on the financial aid tax focused on issues of horizontal equity. Quoting Edlin (1993), “Two families with identical earnings paths pay dramatically different amounts for college if one saves more than the other.” Table 4 compares average assets and EFCs for our baseline parameters to two alternatives. Across the three parameterizations, the families receive the same draws of fulltime income, Y_s , but choose their labor supply and saving according to their preferences.

The first two columns show the average EFC and assets for the baseline parameters. The first alternative, shown in the next two columns, is a less patient family, with $\beta = 0.93$ matching that of the lowest education group. Such families discount the future more heavily relative to the present, and so have higher consumption and lower labor supply (higher leisure) early in the working life, both resulting in lower asset accumulation by the time the child enters college. Using the family with a College Degree as an example, the Less Patient family accumulates about \$235,000 less in assets and receives about \$18,500 more in financial aid per year.

Though not the explicit focus of the early literature on the financial aid tax, a less industrious family with a higher preference for leisure will also pay less for college. The second alternative is such a family, with θ calibrated to an average labor supply of 38 rather than 40. With 5 percent fewer hours worked, this family has lower consumption both before and after retirement, and thus saves less than the baseline family. Using the family with a College Degree as an example again, the Less Industrious family accumulates about \$81,600 less in assets and receives about

\$5,600 more in financial aid per year.³⁴

4.3 Insurance Value

This subsection calculates the insurance value of financial aid by solving the model described in Section 3 under the current financial aid formula and an alternative in which financial aid does not depend on income or assets. Instead, the college changes the cost of attendance by raising or lowering tuition for all students but giving no other aid. This change effectively converts the potentially distortionary taxes on income and assets in the financial aid formula into revenue-equivalent lump-sum taxes. In the absence of income uncertainty, such a change would make the family better off. However, when the family faces income uncertainty, the welfare losses due to the foregone insurance value of financial aid will counteract and may even outweigh the welfare gains due to the loss of distortions to supply labor and accumulate assets under this alternative formula.

Recall from Equation (13) that we quantify this welfare loss by solving for the amount, δ , such that by adjusting the average college cost, $\bar{z}_s(A_s, Y_s)$, by δ , the family achieves the same expected utility, $E[V_1(A_1, Y_1)]$, with this lump-sum college cost that it obtained under the current formula. Positive (negative) adjustments to the lump sum indicate that the family is better (worse) off under the current formula. A positive (negative) adjustment to the lump sum requires the college to offer more (less) aid on average to keep the family as well off as under the current formula.

Table 5 shows the average EFC, average financial aid, and additional aid required to achieve the same utility as the current formula under our baseline parameters. The top panel shows these outcomes when the family faces earnings uncertainty with a standard deviation of the shock equal to 15 percent. Average financial aid ranges from \$8,070 for the highest education group to \$71,547 for the lowest education group.

For all earnings profiles, the compensating variation is positive, indicating that the family would need to be compensated for receiving aid through lump-sum discounts rather than the current formula. The magnitude of this compensation rises from \$1,386 for the lowest education group to \$15,203 for the College Degree earnings

³⁴In both alternatives, multiplying that additional aid by 4 and dividing by the incremental assets, the implicit tax over four years is approximately 30 percent of the incremental assets. This magnitude is consistent with the calculations underlying Figure 3 but not directly comparable, because the alternatives in Table 4 do not hold income constant.

profile and \$19,063 for the highest education group. These additional aid amounts represent discounts in the cost of college of 29 to 81 percent, with the percentage decreasing with the education level. For the family with the College Degree earnings profile, the compensating variation represents a reduction in the EFC of 39.6 percent and an increase in financial aid of 43.6 percent. Put differently, a dollar of financial aid is worth \$1.44 in lump-sum discounts to tuition, because the financial aid is targeted to realizations of low income and assets when its marginal value is higher.

By way of comparison, the bottom panel provides the analogous outcomes when earnings uncertainty is removed by setting $\sigma = 0$. Note first that the average EFC is lower without earnings uncertainty in all but the top age-earnings profile. With earnings uncertainty and a progressive financial aid formula, fortunate earnings realizations are associated with more incremental EFC than unfortunate earnings realizations lose incremental EFC. Conversely, the average EFC is higher without uncertainty for that top profile because, in this case, there is no financial aid at this average earnings profile. Adding uncertainty does nothing to increase the EFC when earnings realizations are fortunate but does result in lower EFCs when earnings realizations are unfortunate.

With no insurance value of financial aid, the additional aid required to compensate for the loss of the income- and asset-contingent elements of the financial aid formula is either zero or negative. For the family with the College Degree earnings profile, the compensating variation of is now -\$6,330, meaning that the family would be willing to pay an additional 26.1 percent in EFC, or equivalently, forego 12.9 percent of financial aid, to avoid the implicit taxes on income and assets in the current formula. For the top education group, the compensating variation is necessarily zero because the financial aid is zero, and thus implicit taxes on both assets and income are zero at the margin. For the lowest education groups, with high discounting and no earnings uncertainty, pre-college asset accumulation is zero and labor supply is nearly constrained by the lower bound even under the lump-sum discounts, generating a negligible compensating variation.

The positive compensating variations shown in the top panel with earnings uncertainty are *net* of the welfare costs of the disincentives illustrated in the bottom panel. Figure 7 stacks the results from Table 5 to illustrate the difference between the gross and net insurance value of financial aid. The red bars are the zero or negative amounts of additional aid when the families face no uncertainty. The blue bars are

the additional aid amounts of aid when the families face earnings uncertainty. Their height represents the insurance value of financial aid *net* of the welfare loss of behavioral distortions shown by the red bars. The insurance value *gross* of these distortions is approximately the combined height of the blue and red bars together.³⁵

The baseline results in Table 5 and Figure 7 use the change in the amount of financial aid to measure the insurance value. Figure 8 shows the impact of this insurance value on the average consumption profile of the family under the baseline parameters. The green curve shows the same average consumption profile from the bottom panel of Figure 6, re-scaled vertically to highlight its variation over time. As shown in Table 5, financial aid is \$34,876 on average. The orange curve pertains to the alternative in which financial aid is \$34,876 regardless of income and assets. The present value of lifetime resources, and therefore consumption, is the same in this “Revenue Equivalent” alternative. That the latter starts out lower and ends higher is due in part to the need for additional precautionary saving in the absence of the insurance provided by the financial aid formula. With that insurance under the current system, the family can spend more early in life when consumption is relatively low. The effect on asset accumulation is notable: families accumulate about 22 percent less under the current system on the eve of matriculation compared to the “Revenue Equivalent” alternative.³⁶ The lower asset accumulation results in lower consumption later in life. Thus, the current formula better allows the family to smooth consumption over time.

The blue curve in Figure 8 is the average consumption profile that obtains when the additional \$15,203 of financial aid (shown in Table 5) is provided to allow the family to achieve the same lifetime expected utility as under the current financial aid formula. The promise of this additional aid allows the family to raise consumption early in life under this “Utility Equivalent” alternative relative to the “Revenue Equivalent” alternative, but not to the extent as under the current formula. There is

³⁵The partition into gross and net is approximate, because the welfare loss due solely to these distortions when earnings uncertainty is zero will understate the analogous welfare loss when earnings are uncertain, as the family works and saves more due to precautionary reasons. With more labor supply and saving, the distortions are larger, and thus Figure 7 likely understates the gross insurance value.

³⁶This reduction is due to both incentives, in the form of the implicit tax on assets and income, and insurance, with a lessened need to save for precautionary reasons. This figure is the closest analogue, for the earnings profile in question, to the reduction in assets that Feldstein (1995) and the subsequent literature sought to measure.

still a greater need for precautionary saving, and the absence of the insurance from the financial aid formula makes it more costly in terms of financial aid to achieve the same level of expected utility.

A notable feature of the baseline results in Table 5 is that the compensating variation is high relative to the average amount of aid for the top three earnings profiles. For example, for the High Growth, Advanced Degree profile, removing a financial aid formula that provides \$8,070 on average would require *additional* compensation of \$19,063, or 236 percent of the aid. The percentages are also high, at 94.9 and 43.6 percent, for the Advanced Degree and College Degree profiles. The explanation is that the financial aid formula provides very well targeted insurance to these families. To illustrate, we computed the correlation (across the 1,000 draws of the earnings process) between the value of lifetime expected utility in the “Revenue Equivalent” scenario, in which there is only a lump-sum discount, and the amount of financial aid under the current formula. The correlations are negative and increasing in absolute value with the earnings profile, ranging from -0.37 for the lowest earnings profile to -0.89, -0.94, and -0.88 for the top three profiles. As resources increase across earnings profiles, the average amount of aid falls, but the targeting of that aid to the least fortunate realizations improves substantially.

4.4 Robustness

Table 6 presents a sensitivity analysis of the compensating variation as the key baseline parameters are changed one at a time. For each parameter change, we recalibrate the preference parameters, β and θ , (if needed) to ensure fulltime work on average over the working life. In all cases, the earnings profile is for the family with a College Degree. The first row of the table repeats the baseline results from Table 5 for a family facing earnings uncertainty with $\sigma = 0.15$. The next four rows consider changes in parameters that affect the amount of risk in the age-earnings profile. Changing these parameters should have a noticeable impact on the insurance value of financial aid. The compensating variation decreases markedly, from \$15,203 to \$2,446, when the standard deviation of the of the earnings shock falls to $\sigma = 0.10$ and to $-\$4,548$ with a standard deviation of $\sigma = 0.05$. Linearly interpolating between the two, a standard deviation of earnings shocks of about 8.25 percent is required for the insurance value to fully offset the negative welfare consequences of the distortionary taxes on income

and assets in the financial aid formula, conditional on the other parameters. The next two rows change the persistence of the earnings shock, raising it with an AR(1) parameter of $\rho = 0.99$ and lowering it with $\rho = 0.90$. With higher persistence, the compensating variation increases to \$23,896, and with lower persistence, it falls to \$4,827.

The next three rows change preference parameters in the utility function. We can make the family less risk averse and more willing to substitute consumption and leisure intertemporally in response to changes in the budget constraint by reducing the curvature of the utility functions. Setting $\gamma = \mu = 2$ and then $\gamma = \mu = 1$ reduces the compensating variation to \$7,285 and \$208, respectively. As expected, less risk averse families derive less value from the insurance and greater harm from the disincentives in the financial aid formula. We can make the family less responsive to future changes in the budget constraint by reducing the discount factor, β from 0.984 to 0.930. Doing so reduces the compensating variation to \$10,940.

The next two rows change the extent to which the family can vary its labor supply in a given period. Fixing the labor supply at $L^F = 40$ raises the compensating variation to \$22,642. Doubling the range from ± 5 (or 12.5%) to ± 10 (or 25%) lowers the compensating variation to \$10,332. As shown in Table 3, the ability to vary labor supply, even while holding the expected value at $L^F = 40$, is a means of insuring against earnings risk. The more it can be done, the less is the incremental value of the insurance provided by the financial aid formula. Similarly, high initial assets help insure consumption against earnings risk. As shown in the next row, with initial assets of \$300,000, the EFC rises and the compensating variation falls, to \$45,062 and \$10,055, respectively.

The remaining rows of the table change the pass through of earnings uncertainty to retirement income (the parameter χ) or aspects of the financial aid environment like the age of the parents relative to the child, the distribution of earnings within the couple, and the amount of the loans available. When the parents are an additional decade older than the child, they have higher income and assets at matriculation and thus pay more for college. However, the compensating variation is also higher, as the financial aid formula can provide insurance against an additional decade of earnings uncertainty. None of the other parameters have a large effect on the EFC or the compensating variation. Overall, Table 6 shows that the compensating variation is appropriately sensitive to assumptions about the family's risk aversion, the degree of

risk faced, and other sources of insurance but generally robust to other changes in the budget constraint.

5 Removing and Replacing the Asset Tax

The model outcomes in Section 4 compare the current formula to an alternative in which neither income nor assets affect the cost of college. Conceptually, the insurance value of financial aid derives from the implicit taxes on both assets and income. However, the policy interest in the financial aid formula has focused primarily on the implicit tax on assets. In this section, we consider variations in the financial aid formula that eliminate only the Contribution from Assets by setting $j = 0$ in Equation (3), thus making Available Income the only component of Adjusted Available Income.³⁷

We begin by showing in Table 7 how the average EFC, average asset accumulation on the eve of matriculation, and the compensating variation change when the Contribution from Assets is eliminated. For each such outcome, the column headed “ $j = 0$ ” gives the dollar amount in thousands and the column headed “% Baseline” indicates what percentage that dollar amount is of the corresponding dollar amount in the baseline results shown in Table 4 or 5.

The first two columns show the reductions in the average EFC. For the College Degree earnings profile, the new EFC of \$21,600 is 56.3 percent of the baseline value of \$38,374. The residual 43.7 percent of that baseline value is attributable to the Contribution from Assets. For lower earnings profiles, the percentage of the baseline EFC rises, as asset accumulation is lower and thus the Contribution from Assets accounts for less of the average EFC. For higher earnings profiles, the percentage of the baseline EFC also rises, despite higher asset accumulation, as a greater percentage of the families pay the maximum EFC of \$73,250.³⁸

Eliminating the Contribution from Assets has two effects on pre-matriculation asset accumulation, shown in the next two columns. The first is the higher return to

³⁷With $j = 0$, assets only impact the EFC through the inclusion of asset income in Available Income.

³⁸The percentage receiving no aid in the baseline rises from 12.9 percent for the College Degree earnings profile to 41.6 and 70.7 percent for the Advanced Degree and High Growth Advanced Degree profiles, respectively.

saving, which tends to increase asset accumulation.³⁹ The second is the lower cost of college (shown in the first two columns), which would decrease the need for asset accumulation. The first two rows show that for the highest earnings profiles, these two effects roughly offset. For lower earnings profiles, the former outweighs the latter, and asset accumulation rises by 10 to 28 percent relative to the baseline.

The final two columns show the additional aid required to compensate for the loss of the income-contingent components of the financial aid formula. The percentages of baseline are similar to, and generally slightly higher than, the corresponding percentages for the average EFC. For the College Degree earnings profile, financial aid would have to be \$9,500 higher if delivered as a lump sum, reducing the average EFC of \$21,600 by 44 percent. Thus, the insurance value of financial aid is still positive even if the Contribution from Assets were eliminated or, equivalently, if a family managed to save only in types of assets that are not covered by the financial aid formula (e.g. retirement accounts).

The family is better off when “ $j = 0$,” because college costs fall dramatically. To conduct more interesting welfare comparisons regarding the financial aid tax on assets, we remove the Contribution from Assets by setting $j = 0$ and solve for the dollar value of a fixed fee – F in Equation (3) – that restores the family to its expected utility under the current formula. That is, F solves:

$$W(1, 1, 0) = W(0, 1, F) \tag{14}$$

We can then compare $\bar{z}(0, 1, F)$ to $\bar{z}(1, 1, 0)$ to ascertain whether the college must pay more or less on average to achieve this expected utility level for the family. The results are presented in Table 8, with the fixed fees, F , shown in the first column.⁴⁰ The fee rises with the earnings profile before falling somewhat for the highest earnings profile, where the asset tax plays a diminished role due to the frequency with which these families receive no financial aid.

³⁹The asset tax applies with less intensity at the extremes of the distribution, either because the family’s resources are more often sufficiently high that it receives no aid (and thus no marginal reduction in aid for higher assets) or because its resources are so low that the slope of the AAI-to-EFC relationship in Equation (4) is not at its maximum or is identically zero through the Simplified Needs Analysis in the Federal Methodology. Across all earnings profiles, the assumption that the college meets all demonstrated need is relevant here, as relaxing that assumption would mean that the disincentives would apply less consistently and induce less dissaving.

⁴⁰Because the fixed fee is different for each earnings profile, the alternative is not a single formula that could be implemented without knowing the parents’ education level.

The next two columns show that for every earnings profile, the average EFC under the Fixed Fee formula is lower than under the current formula. For the College Degree earnings profile, for example, the family’s expected utility is maintained when the average EFC is \$33,197, or \$5,178 below its average under the current formula. Thus, the college must provide this family with 14.8 percent more aid than under the current formula. That average aid must increase indicates that, conditioning on the net insurance value provided by the implicit tax on income, the net insurance value of the Contribution from Assets is positive. It would lower expected utility to replace the Contribution from Assets with a fixed fee that enabled the college to offer only the same amount of aid on average.

As an alternative to imposing a fixed fee, we again set $j = 0$ but maintain the families’ expected utilities by raising $k > 1$. This alternative formula eliminates the insurance value and disincentives associated with the Contribution from Assets, while increasing the insurance value and disincentives associated with Available Income. We search for the value of k that, for each earnings profile, solves:

$$W(1, 1, 0) = W(0, k, 0) \tag{15}$$

We again compare the average EFC in the two cases. The results are shown in Table 9. The first column shows the values of k that equalize expected utilities. They range from 1.231 to 1.842 and, like the fixed fees in Table 8, rise with the earnings profile before falling back for the highest profile. For the lowest two earnings profiles, the differences in average financial aid are positive but small. For the next three earnings profiles, the additional aid is negative. For example, for the family with the College Degree earnings profile, the same expected utility can be achieved with $j = 0, k = 1.776$ as with $j = 1, k = 1$ with \$1,774 less given in financial aid on average. Starting at the baseline parameters, the family would forsake the net insurance value provided by the Contribution from Assets for additional net insurance value provided by Available Income. This is not true for the family with the highest earnings profile, for whom the disincentives from the asset tax are less relevant.⁴¹

⁴¹As a third alternative, focusing on the highest earnings profile, we also considered eliminating the asset tax while raising tuition. We found that, within the simulated population, even an uncapped tuition would not raise enough revenue to restore (i.e. lower) expected utility back to its level under the baseline parameters. As tuition is raised, the proportion of the population paying the full amount declines too rapidly.

6 Discussion

The early literature on the financial aid tax in Feldstein (1995) and Edlin (1993) highlights the possibility that implicit taxes on labor supply and especially saving would distort behavior and thus lower welfare for families with children anticipating their college years. Our model incorporates those distortions and shows the negative impact on welfare in the absence of earnings uncertainty. When earnings uncertainty is introduced at levels consistent with empirical studies, the compensating variation turns positive. The distortions are still present – their harmful impacts on welfare are just substantially outweighed by the insurance value they provide. The insurance value of financial aid is large and positive, when considering the inclusion of both income and assets in the formula, when considering income by itself, and when considering the incremental effect of assets conditional on income.

This last finding, shown most clearly in Table 8, demonstrates that a family’s asset accumulation on the eve of matriculation contains incremental information about whether the family has been relatively lucky over its pre-college earnings draws, even conditioning on the current draw of earnings during the college years. Knowing this information, and redistributing from lucky to unlucky scenarios, improves welfare by more than enough to offset the welfare loss due to the disincentive effects of implicitly taxing the return to saving.

The source of the welfare gain in our model is that in scenarios in which earnings realizations have been low, the current financial aid formula reduces the cost of attending college whereas the alternatives do not. Such estimates of the insurance value of financial aid will be sensitive to how we model other choices that might alleviate the burden of a high tuition payment in the face of low assets and income. Two choices are already included – increasing labor supply and taking out loans. Another would be to attend a college that costs less but (as must be the case in equilibrium) delivers lower benefits. To the extent that those lower benefits are lower earnings in the future, the timing of the cash flows mimics that of borrowing. Consumption falls less today but income to support future consumption (here thought of as the collective income of all members of the family) is lower.⁴²

⁴²In the model (as in reality), the loan opportunity exists in both the current formula and the alternatives and, perhaps as a result, the sensitivity analysis in Table 6 indicates that it has only a small impact on the compensating variation. If the borrowing opportunity is to better resemble sacrificing future earnings by going to a lower-cost school, it would be available only in the alternative

There are several possible directions for further research, most of which would expand the complexity of the model beyond the framework of two choice variables and two state variables used here. First, we do not consider the real growth in the cost of attending college or the uncertainty surrounding it, even though this cost growth and uncertainty are prominent in policy discussions regarding access to higher education. Incorporating this growth and uncertainty would likely increase the insurance value of financial aid, since the EFC that comes from the Federal Methodology does not depend explicitly on college costs except as a maximum. Indirectly, the growth in the cost of attending college influences the indexing of the dollar values in the financial aid formulas.

Second, as noted above, not all assets are included in the measure of assets used in the financial aid formula. Retirement accounts are excluded from both the Federal and Institutional Methodologies, and home equity is also excluded from the Federal Methodology. A more general model of saving decisions in the presence of financial aid would include a state variable, say M , to represent excluded assets and a choice variable, say m , to reflect net saving in these excluded assets. The family's problem would then be to maximize the same objective function as in Equation (5) by choosing all three of $C(A_t, Y_t, M_t)$, $L(A_t, Y_t, M_t)$, and $m(A_t, Y_t, M_t)$ each period. This is a considerably more complicated problem to solve numerically. Similarly, we do not consider the growing industry of tax-advantaged college saving vehicles, like 529 plans, that make saving for college relatively cheaper than in our model.

Third, we do not consider parents' payments for college in a more general context of intergenerational transfers to children. In such a framework, payments for college could be replaced by direct payments of cash if the value proposition in college becomes less favorable. They could also be replaced by larger bequests, accumulated over a longer period and thus less of a drag on consumption during the working life.

Fourth, the cost to the providers of financial aid of offering this insurance has not been modeled in the analysis, but such costs are likely to be small. Providers like governments and colleges offer financial aid based on this formula to a large population of students – those from families who have been lucky and those who have not. To the extent that the earnings uncertainty these families face is idiosyncratic in nature, the aggregation of aid awards across this population diversifies away the risk. To the extent that there are more systemic shocks to the families' income, the long

formulas.

time horizons for governments and colleges allow them some opportunity to smooth these fluctuations over time.

Finally, as suggested by Dick and Edlin (1997), assets are reflecting lifetime income – information beyond what is available in current income. For a given level of current income, a low level of assets indicates that prior income shocks were sufficiently low that the family found it optimal to consume most of its income. Including assets in the financial aid formula allows the formula to partially insure against those prior shocks as well. Future work can consider how welfare might be improved by introducing a measure of lifetime average earnings into the financial aid formula, allowing the sensitivities to assets and current income to be lessened.

7 Conclusion

Prior literature has conjectured, and provided mixed empirical evidence, that the implicit tax on assets in the financial aid formula creates a distortion in saving behavior. The literature has not considered as extensively that there is also an implicit tax on labor earnings in the formula. Our analysis is the first to recognize that these implicit taxes are merely one component of a standard incentives-insurance tradeoff. Using a stochastic, life-cycle model of consumption and labor supply in which families have precautionary, retirement, and college motives for saving, we show that in a model without earnings uncertainty, the implicit taxes can have modest negative consequences for lifetime expected utility. When families face earnings uncertainty, the insurance value of a financial aid formula based on the current formula is substantial.

Across a range of parameterizations, we calculate that financial aid would have to increase by enough to reduce the net cost of attendance by 30 to 80 percent to compensate families for the loss of the income- and asset-contingent elements of the current formula. For parents facing the earnings process representative of college graduates, a dollar of financial aid delivered through the current formula is worth \$1.44 in lump-sum tuition discounts, due precisely to the targeting of the financial aid to scenarios in which the family has low income or assets and thus a greater marginal value of additional resources.

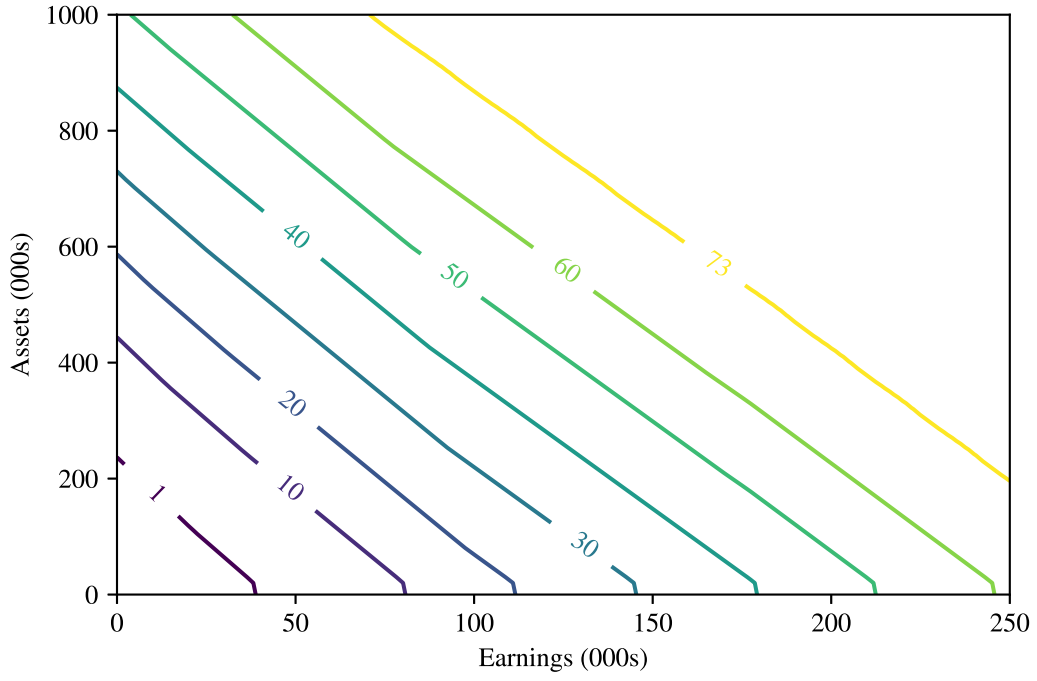
Replacing just the “financial aid tax” on assets with a lump sum would also reduce welfare. To keep the family with college-educated parents as well off with the lump sum, average financial aid would have to increase by about 15 percent. The net

insurance value from this “financial aid tax” is positive at the margin. However, for some families with earnings profiles such that the saving disincentives of the asset tax may be acute, the net insurance value from the implicit tax on assets is less than the net insurance value from the implicit tax on income at the margin. Increasing the sensitivity of the financial aid formula to income and reducing it to assets could increase welfare.

Without considering earnings uncertainty, the welfare losses due to the disincentives in the financial aid tax appear to be the economic costs of the explicitly redistributive financial aid formula. The formula transfers resources from those with (predictably) higher assets and income to those with (predictably) lower assets and income. However, when reasonable amounts of earnings uncertainty are added to the model, the progressive nature of the implicit taxes confers the additional benefits of insurance against that earnings uncertainty. That is, in addition to the justification for the implicit taxes on assets and income based on a desire to redistribute across *ex ante* different groups, there is a justification based on a desire to redistribute within an *ex ante* identical group based on *ex post* realizations of an uncertain earnings process.⁴³ For reasonable parameterizations, the insurance value of means-tested financial aid more than offsets the disincentive costs of means-tested financial aid. Put differently, governments and institutions that provide financial aid according to this formula are able to give less aid than they would have to otherwise in order to keep the family of the college student equally well off.


⁴³Colas, Findeisen and Sachs (2019) also show that optimal financial aid is declining in parental income even without a redistributive motive. Unlike the model here, which focuses on the insurance value of financial aid to the parents, they focus on how the price-sensitivity and the lifetime (positive) fiscal externality of college-going decline with income.

Figure 1: EFC Levels by Assets and Earnings




Source: Calculations based on the EFC Formula in 2018 - 2019.

Figure 2: EFC by Earnings for Specified Asset Levels (000s)



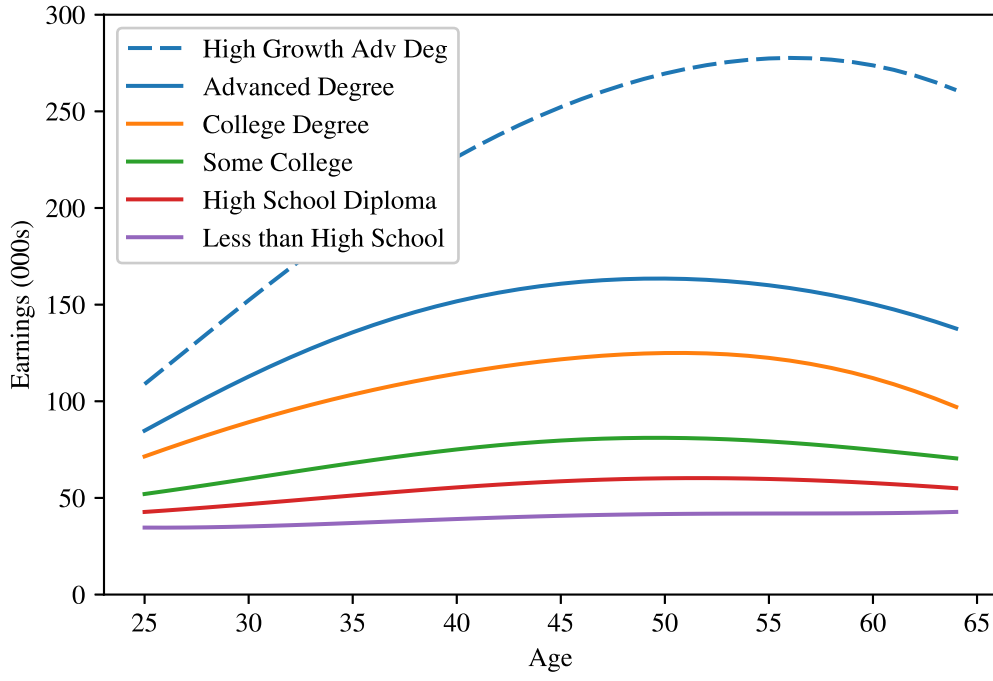
Source: Calculations based on the EFC Formula in 2018 - 2019.

Figure 3: EFC by Assets for Specified Earnings Levels (000s)




Source: Calculations based on the EFC Formula in 2018 - 2019.

Figure 4: Expanded Age-Earnings Profiles by Education: Couples



Source: Estimates based on the American Community Survey 2017.

Figure 5: Simulated and Empirical Wealth-to-Income Moments by Education



Source: Empirical moments (dots) are calculated using the Survey of Consumer Finances 2016. Simulated moments (curves) are from the model under baseline parameters. See Table 2 for the underlying data.

Figure 6: Average Earnings, Consumption, and Assets by Age



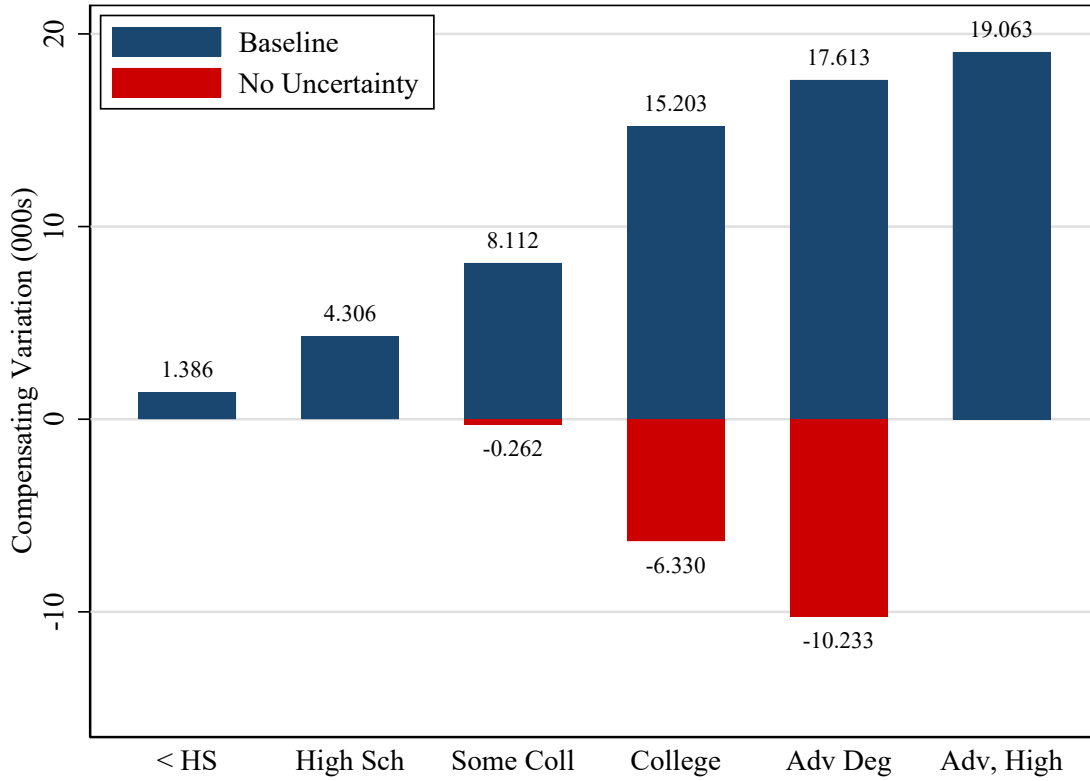



Figure 7: Insurance Value, Gross and Net, by Earnings Profile



Source: Insurance values calculated from the model under baseline parameters. Net insurance value is the height of the blue bars. Gross insurance value is approximated by the combined height of the blue and red bars.

Figure 8: Average Consumption by Financial Aid Formula and Age



Source: Consumption profiles simulated by the model under baseline parameters. “Revenue Equivalent” applies the average financial aid under the Current Formula as a lump sum. “Utility Equivalent” augments that lump sum by enough to restore expected utility to that under the Current Formula.

Table 1: Cross-Sectional Log Earnings Regressions by Education for Couples

	Less than High School	High School Diploma	Some College	College Degree	Advanced Degree
Constant	12.50 (1.274)	11.01 (0.679)	10.83 (0.497)	7.360 (0.539)	6.913 (0.716)
Age	-0.218 (0.122)	-0.0693 (0.0652)	-0.0563 (0.0479)	0.309 (0.0521)	0.334 (0.0678)
Age ² (/10)	0.0812 (0.0427)	0.0339 (0.0228)	0.0375 (0.0169)	-0.0904 (0.0184)	-0.0855 (0.0235)
Age ³ (/100)	-0.0124 (0.00647)	-0.00532 (0.00346)	-0.00668 (0.00257)	0.0130 (0.00281)	0.0104 (0.00355)
Age ⁴ (/1000)	0.000674 (0.000359)	0.000263 (0.000192)	0.000363 (0.000143)	-0.000760 (0.000157)	-0.000519 (0.000196)
Observations	26504	85381	134450	123547	102852
R^2	0.010	0.020	0.034	0.047	0.039

Standard errors in parentheses

Source: 2017 American Community Survey 1-Year PUMS Files

Sample includes all couples with at least one member age 25-64 and working fulltime.

Education and age pertain to the age-eligible member working fulltime with more education.

Table 2: Calibrated Utility Function Parameters

	Less than High School	High School Diploma	Some College	College Degree	Advanced Degree	High Growth Adv Degree
Age 25 - 34						
Empirical	0.398	0.506	0.410	0.793	0.435	—
Simulated	0.402	0.394	0.336	0.512	0.482	—
Age 35 - 44						
Empirical	0.430	1.872	0.833	1.659	2.079	—
Simulated	0.850	1.143	1.259	2.007	2.139	—
Age 45 - 54						
Empirical	0.709	2.183	2.133	4.225	4.274	—
Simulated	1.240	2.057	2.548	3.879	4.292	—
Age 55 - 64						
Empirical	2.329	2.704	4.617	6.295	6.940	—
Simulated	1.612	3.177	4.127	6.434	6.890	—
β	0.930	0.952	0.961	0.984	0.992	0.992
Goodness of Fit	0.973	0.783	0.599	0.339	0.009	
θ	54.646	32.974	21.722	12.099	8.476	4.413

Source: Empirical moments are from the 2016 Survey of Consumer Finances

Sample includes all couples with at least one member age 25-64 and working fulltime.

Education and age pertain to the age-eligible member working fulltime with more education.

Simulated moments are the ratio of average assets to average earnings, by age and education.

Table 3: Labor Supply Variation Over the Life Cycle

Earnings Profile	Years Relative to Child's College		
	Before	During	After
	Current Financial Aid Formula		
High Growth, Adv Degree	42.47	37.96	38.59
Advanced Degree	42.77	36.04	38.53
College Degree	41.77	35.67	39.48
Some College	40.34	36.76	40.88
High School Diploma	40.12	37.28	40.95
Less than High School	40.20	38.93	40.38
	Revenue Equivalent Tuition Discount		
High Growth, Adv Degree	42.82	39.68	37.87
Advanced Degree	43.28	39.24	37.25
College Degree	42.68	39.64	37.74
Some College	41.27	39.73	39.26
High School Diploma	40.60	39.37	40.02
Less than High School	40.27	39.52	40.21

Source: Model simulations under baseline parameters.

Revenue Equivalent Tuition Discount holds financial aid constant while removing the income- and asset-contingent elements of the financial aid formula.

Each cell is average labor supply, where 40 represents fulltime work.

Table 4: The Financial Aid Tax, by Earnings Profile

Earnings Profile	Baseline		Less Patient		Less Industrious	
	EFC	Assets	EFC	Assets	EFC	Assets
	High Growth, Adv Degree	65.2	781.8	49.1	136.8	61.2
Advanced Degree	54.7	655.1	29.5	93.9	49.1	496.0
College Degree	38.4	397.8	19.9	62.8	34.1	316.2
Some College	13.4	114.7	9.3	35.1	13.0	104.2
High School Diploma	6.2	58.5	4.7	22.5	6.2	56.9
Less than High School	1.7	15.3	1.7	15.3	1.9	19.1

Source: Model simulations under baseline parameters and alternatives.

Each cell is average EFC or Assets at matriculation in Thousands of Dollars.

“Less Patient” lowers β to 0.930.

“Less Industrious” calibrates θ to labor supply of 38, not 40.

Table 5: Compensating Variations by Earnings Profile, Baseline Parameters

Parameters	Average EFC	Average Aid	Additional Aid	Percent of EFC	Percent of Aid
Baseline Earnings Uncertainty					
High Growth, Adv Degree	65.180	8.070	19.063	-29.2	236.2
Advanced Degree	54.686	18.564	17.613	-32.2	94.9
College Degree	38.374	34.876	15.203	-39.6	43.6
Some College	13.393	59.857	8.112	-60.6	13.6
High School Diploma	6.199	67.051	4.306	-69.5	6.4
Less than High School	1.703	71.547	1.386	-81.4	1.9
No Earnings Uncertainty					
High Growth, Adv Degree	73.250	0.000	0.000	0.0	NA
Advanced Degree	43.425	29.825	-10.233	23.6	-34.3
College Degree	24.248	49.002	-6.330	26.1	-12.9
Some College	7.089	66.161	-0.262	3.7	-0.4
High School Diploma	3.261	69.989	0.000	0.0	0.0
Less than High School	0.593	72.657	0.000	0.0	0.0

Source: Model simulations for baseline parameters.

Top panel sets $\sigma = 0.15$. Bottom panel sets $\sigma = 0$.

EFC and Aid are in Thousands of Dollars.

Table 6: Compensating Variations, Sensitivity Checks

Parameters	Financial Aid			Percent of EFC	Percent of Aid
	EFC	Aid	Additional Aid		
Baseline Parameters	38.374	34.876	15.203	-39.6	43.6
Decrease σ from 0.15 to 0.10	44.087	29.163	2.446	-5.5	8.4
Decrease σ from 0.15 to 0.05	48.296	24.954	-4.548	9.4	-18.2
Increase ρ from 0.95 to 0.99	35.382	37.868	23.896	-67.5	63.1
Decrease ρ from 0.95 to 0.90	42.318	30.932	4.827	-11.4	15.6
Decrease γ and μ from 3 to 2	36.288	36.962	7.285	-20.1	19.7
Decrease γ and μ from 3 to 1	31.592	41.658	0.208	-0.7	0.5
Decrease β from 0.984 to 0.930	19.816	53.434	10.940	-55.2	20.5
Fix Labor Supply at $L^F = 40$	41.188	32.062	22.642	-55.0	70.6
Widen Labor Supply Range from ± 5 to ± 10	36.178	37.072	10.332	-28.6	27.9
Increase Initial Assets from 0 to 300,000	45.062	28.188	10.055	-22.3	35.7
Increase χ from 0.35 to 0.70	37.409	35.841	15.163	-40.5	42.3
Decrease χ from 0.35 to 0	38.792	34.458	15.447	-39.8	44.8
Become a Parent 10 Years Later	47.704	25.546	18.952	-39.7	74.2
Switch to One-earner Couple	39.057	34.193	16.123	-41.3	47.2
Increase Annual Loan from 10k to 30k	40.580	32.670	15.095	-37.2	46.2

Source: Model simulations for Households with College Degrees.

β and θ are recalibrated as needed to match wealth-to-income and ensure fulltime work for each alternative.

EFC and Aid are in Thousands of Dollars.

Table 7: Results with Asset Tax Removed, by Earnings Profile

Earnings Profile	Average		Average		Compensating	
	EFC		Assets		Variation	
	j = 0	% Baseline	j = 0	% Baseline	j = 0	% Baseline
High Growth, Adv Degree	52.2	80.0	785.6	100.5	16.9	88.8
Advanced Degree	32.5	59.4	659.6	100.7	12.3	70.1
College Degree	21.6	56.3	436.7	109.8	9.5	62.8
Some College	9.0	67.1	147.2	128.4	5.4	66.7
High School Diploma	4.3	68.7	73.5	125.6	2.9	66.4
Less than High School	1.3	77.2	17.5	114.1	1.1	77.0

Source: Model simulations for baseline parameters.

“j = 0” indicates that the Contribution from Assets is set equal to zero.

Each cell is Thousands of Dollars or Percent of the Amount in Table 4 or 5.

Table 8: Replacing Contribution from Assets with Fixed Fee

Earnings Profile	Fixed Fee	Average EFC	Baseline EFC	Additional Aid	Percent of Aid
Adv Degree, High Growth	14.255	61.784	65.180	3.396	42.1
Advanced Degree	17.748	49.260	54.686	5.426	29.2
College Degree	11.527	33.197	38.374	5.178	14.8
Some College	1.847	10.865	13.393	2.529	4.2
High School Diploma	0.520	4.786	6.199	1.413	2.1
Less than High School	0.070	1.385	1.703	0.318	0.4

Source: Model simulations under baseline parameters and alternatives.

EFC and Aid are in Thousands of Dollars.

Table 9: Replacing Contribution from Assets with Higher Income Sensitivity

Earnings Profile	New Income Sensitivity	Average EFC	Baseline EFC	Additional Aid	Percent of Aid
Adv Degree, High Growth	1.466	64.015	65.180	1.165	14.4
Advanced Degree	1.842	55.166	54.686	-0.480	-2.6
College Degree	1.776	40.148	38.374	-1.774	-5.1
Some College	1.390	13.641	13.393	-0.248	-0.4
High School Diploma	1.312	5.961	6.199	0.238	0.4
Less than High School	1.231	1.701	1.703	0.002	0.0

Source: Model simulations under baseline parameters and alternatives.

EFC and Aid are in Thousands of Dollars.

Appendix A Detailed Solution Method

This appendix derives the first-order conditions for the family's optimization problem introduced in Section 3.2. The two sub-period problems in period t are defined as follows. In the first sub-period, the family chooses labor supply, L_t , according to:

$$VL_t(A_t, Y_t) \equiv \max_{L_t} \theta v(L_t) + VC_t(X_t, Y_t) \quad (\text{A.1a})$$

$$\hat{Y}_t = Y_t \left(\frac{L_t}{L^F} \right) \quad (\text{A.1b})$$

$$X_t = A_t + \hat{Y}_t - h(\hat{Y}_t) - z_t(A_t, \hat{Y}_t) - g(\hat{Y}_t) \quad (\text{A.1c})$$

$$L_t^{\min} \leq L_t \leq L_t^{\max} \quad (\text{A.1d})$$

In the second sub-period, the family chooses consumption, C_t , according to:

$$VC_t(X_t, Y_t) \equiv \max_{C_t} u(C_t) + \beta E_t [VL_{t+1}(A_{t+1}, Y_{t+1})] \quad (\text{A.2a})$$

$$A_{t+1} = (1 + r(1 - g'(Y_t)))(X_t - C_t) \quad (\text{A.2b})$$

$$A_{t+1} \geq 0 \quad (\text{A.2c})$$

In Equations (A.1) and (A.2), $VL_t(A_t, Y_t)$ and $VC_t(X_t, Y_t)$ are the value functions for the labor and consumption sub-period problems, respectively. The main change from the original formulation of the problem in Equation (5) is in the way income taxes are calculated, which must be approximated when there are two sub-periods. In the first sub-period, income taxes are collected on labor income assuming that capital income, which is determined in the second sub-period, is zero. This is the term, $g(\hat{Y}_t)$, in Equation (A.1c). The income tax function is progressive in labor income, i.e. both $g' \geq 0$ and $g'' \geq 0$. This approximation means that, with capital income set to zero in the first sub-period, the marginal tax rate on labor income may be understated. In the second sub-period, income taxes are collected on capital income, $r(X_t - C_t)$, at a rate of $g'(Y_t)$, as shown in Equation (A.2b). The income tax on capital income is set equal to the marginal income tax rate based on fulltime

income, Y_t , multiplied by the amount of capital income. The approximations mean that the marginal tax rate is constant at $g'(Y_t)$, rather than progressive, and uses the state variable, Y_t , rather than the prior sub-period's choice variable, \hat{Y}_t , as the base. This latter simplification is required in order to avoid adding \hat{Y}_t as an additional state variable in the second sub-period.

In this new formulation of the family's problem, the first-order condition for the labor supply choice in the first sub-period is:

$$\begin{aligned} \theta v'(L_t) + \frac{\partial VC_t(X_t, Y_t)}{\partial X_t} \left(\frac{Y_t}{L^F} \right) \left(1 - h'(\hat{Y}_t) - z_t^Y(A_t, \hat{Y}_t) - g'(\hat{Y}_t) \right) \\ = \mu^{\max} - \mu^{\min} \end{aligned} \quad (\text{A.3})$$

The first-order condition for the consumption choice in the second sub-period is:

$$u'(C_t) - \beta(1 + r(1 - g'(Y_t))) \left(E_t \left[\frac{\partial VL_{t+1}(A_{t+1}, Y_{t+1})}{\partial A_{t+1}} \right] + \lambda \right) = 0 \quad (\text{A.4})$$

The first term in the first-order condition is the marginal utility of an additional dollar of consumption in period t . The second term is the discounted value of saving that dollar to be used in period $t + 1$. The dollar grows by the after-tax interest rate and has a marginal value of $\frac{\partial VL_{t+1}(A_{t+1}, Y_{t+1})}{\partial A_{t+1}}$ at that time. This marginal value is uncertain because of the shock to income received in period $t + 1$. In this expression, $r \cdot g'(Y_t)$ is the marginal tax on another dollar of saving. The marginal utility of a dollar of assets at time $t + 1$ is discounted back to period t utility by a factor of β . The difference between the marginal utility of consumption and the discounted, expected marginal utility of assets in the next period is zero at the optimal level of consumption.

We can use the Envelope Theorem to obtain analytical expressions for the $\frac{\partial VC_t(X_t, Y_t)}{\partial X_t}$ and $\frac{\partial VL_{t+1}(A_{t+1}, Y_{t+1})}{\partial A_{t+1}}$ terms that appear in these first-order conditions. Applying the Envelope Theorem to Equation (A.2a) yields an expression for $\frac{\partial VC_t(X_t, Y_t)}{\partial X_t}$:

$$\frac{\partial VC_t(X_t, Y_t)}{\partial X_t} = \beta(1 + r(1 - g'(Y_t))) \left(E_t \left[\frac{\partial VL_{t+1}(A_{t+1}, Y_{t+1})}{\partial A_{t+1}} \right] + \lambda \right) \quad (\text{A.5})$$

which is equal to $u'(C_t^*)$ by Equation (A.4). This substitution can be made in Equation (A.3) to get a new first-order condition for the labor supply choice:

$$\begin{aligned} \theta v'(L_t) + u'(C_t^*) \left(\frac{Y_t}{L^F} \right) \left(1 - h'(\hat{Y}_t) - z_t^Y (A_t, \hat{Y}_t) - g'(\hat{Y}_t) \right) \\ = \mu^{\max} - \mu^{\min} \end{aligned} \quad (\text{A.6})$$

This is Equation (10). Applying the Envelope Theorem to Equation (A.1a) yields an expression for $\frac{\partial V L_t(A_t, Y_t)}{\partial A_t}$:

$$\begin{aligned} \frac{\partial V L_t(A_t, Y_t)}{\partial A_t} &= \frac{\partial V C_t(X_t, Y_t)}{\partial X_t} \cdot \frac{\partial X_t}{\partial A_t} \\ &= u'(C_t^*) \left(1 - z_t^A \left(A_t, L_t^* \left(\frac{Y_t}{L^F} \right) \right) \right) \end{aligned} \quad (\text{A.7})$$

Advancing this equation to period $t + 1$ and substituting it into Equation (A.4) yields a new first-order condition for the consumption choice:

$$\begin{aligned} u'(C_t) &= \beta (1 + r (1 - g'(Y_t))) \cdot \\ &\left(E_t \left[u'(C_{t+1}^*) \left(1 - z_{t+1}^A \left(A_{t+1}, L_{t+1}^* \left(\frac{Y_{t+1}}{L^F} \right) \right) \right) \right] + \lambda \right) \end{aligned} \quad (\text{A.8})$$

This is Equation (11). When the family is in retirement, the only choice each period is for optimal consumption, and there is no remaining uncertainty in the income process. For all retirement periods prior to the last period of life, optimal consumption (when the liquidity constraint does not hold with equality) is given by a first-order condition analogous to Equation (A.8):

$$\begin{aligned} u'(C_t) &= \beta (1 + r (1 - g'(Y_t + r(X_t - C_t)))) \cdot \\ &\left(1 - z_{t+1}^A (A_{t+1}, Y_{t+1}) \right) u'(C_{t+1}^* (A_{t+1}, Y_{t+1})) \end{aligned} \quad (\text{A.9})$$

Note that the absence of income uncertainty means that there is no expectations operator around the marginal utility of consumption next period. Further, there is

no need to approximate the income tax function in retirement periods. Finally, in the application of the model considered above, college expenses are assumed to occur before retirement in the model, so the $(1 - z_{t+1}^A(A_{t+1}, Y_{t+1}))$ term is always 1 during retirement years.⁴⁴

⁴⁴When the liquidity constraint that A_{t+1} cannot be negative is binding, then consumption in period t is given by $X_t - \frac{g(Y_t)}{1+r}$.

References

- Andersen, Steffen, Glenn W Harrison, Morten I Lau, and E Elisabet Rutström.** 2008. “Eliciting Risk and Time Preferences.” *Econometrica*, 76(3): 583–618.
- Athreya, Kartik, Deven Reilly, and Nicole Simpson.** 2014. “Young Unskilled Women and the Earned Income Tax Credit: Insurance Without Disincentives?” Federal Reserve Bank of Richmond Working Paper 14-11R.
- Barrow, Lisa, and Ofer Malamud.** 2015. “Is College a Worthwhile Investment?” *Annual Review of Economics*, 7(1): 519–555.
- Braga, Breno, and Olga Malkova.** 2020. “Hope for the Family: The Effects of College Costs on Maternal Labor Supply.” IZA 12958.
- Bricker, Jesse, Lisa J Dettling, Alice Henriques, Joanne W Hsu, Lindsay Jacobs, Kevin B Moore, Sarah Pack, John Sabelhaus, Jeffrey Thompson, and Richard A Windle.** 2017. “Changes in US Family Finances from 2013 to 2016: Evidence from the Survey of Consumer Finances.” *Federal Reserve Bulletin*, 103(3): 1–42.
- Cagetti, Marco.** 2003. “Wealth Accumulation Over the Life Cycle and Precautionary Savings.” *Journal of Business & Economic Statistics*, 21(3): 339–353.
- Carneiro, Pedro, James J Heckman, and Edward J Vytlačil.** 2011. “Estimating Marginal Returns to Education.” *American Economic Review*, 101(6): 2754–81.
- Carroll, Christopher D.** 2019. “Solution Methods for Microeconomic Dynamic Stochastic Optimization Problems.” <http://www.econ2.jhu.edu/people/ccarroll/SolvingMicroDSOPs.pdf>.
- Carroll, Christopher D, and Andrew A Samwick.** 1997. “The Nature of Precautionary Wealth.” *Journal of Monetary Economics*, 40(1): 41–71.
- Case, Karl E, and Michael S McPherson.** 1986. “Does Need-Based Student Aid Discourage Saving for College?”
- Chandra, Amitabh, and Andrew A Samwick.** 2009. “Disability Risk and the Value of Disability Insurance.” In *Health at Older Ages: The Causes and Consequences of Declining Disability among the Elderly*. 295–336. University of Chicago Press.
- Colas, Mark, Sebastian Findeisen, and Dominik Sachs.** 2019. “Optimal Need-Based Financial Aid.” <http://www.dominiksachs.com/downloads/need-based3.pdf>.

- DeBacker, Jason, Bradley Heim, Vasia Panousi, Shanthi Ramnath, and Ivan Vidangos.** 2013. “Rising Inequality: Transitory or Persistent? New Evidence from a Panel of US Tax Returns.” *Brookings Papers on Economic Activity*, 2013(1): 67–142.
- Dick, Andrew W, Aaron S Edlin, and Eric R Emch.** 2003. “The Savings Impact of College Financial Aid.” *Contributions in Economic Analysis & Policy*, 2(1).
- Dick, Andrew W, and Aaron S Edlin.** 1997. “The Implicit Taxes from College Financial Aid.” *Journal of Public Economics*, 65(3): 295–322.
- Eaton, Jonathan, and Harvey S Rosen.** 1980. “Taxation, Human Capital, and Uncertainty.” *The American Economic Review*, 70(4): 705–715.
- Edlin, Aaron S.** 1993. “Is College Financial Aid Equitable and Efficient?” *The Journal of Economic Perspectives*, 7(2): 143–158.
- Engen, Eric M, and Jonathan Gruber.** 2001. “Unemployment Insurance and Precautionary Saving.” *Journal of Monetary Economics*, 47(3): 545–579.
- Feldstein, Martin.** 1995. “College Scholarship Rules and Private Saving.” *American Economic Review*, 85(3): 552–566.
- Fidelity Investments.** 2018. “2018 College Savings Indicator: Executive Summary of Key Findings.” https://www.fidelity.com/bin-public/060_www_fidelity.com/documents/PR/CSI-2018-Executive-Summary.pdf.
- Gibson, David S, and Erin P Gibson.** 2017. “Age-Earning Profiles: Refinement and Applications.” *The Rehabilitation Professional*, 25(1): 13–34.
- Golosov, Mikhail, Maxim Troshkin, and Aleh Tsyvinski.** 2016. “Redistribution and Social Insurance.” *American Economic Review*, 106(2): 359–86.
- Grant, Charles, Christos Koulovatianos, Alexander Michaelides, and Mario Padula.** 2010. “Evidence on the Insurance Effect of Redistributive Taxation.” *The Review of Economics and Statistics*, 92(4): 965–973.
- Guvenen, Fatih, Fatih Karahan, Serdar Ozkan, and Jae Song.** 2015. “What Do Data on Millions of U.S. Workers Reveal about Life-Cycle Earnings Risk?” National Bureau of Economic Research Working Paper 20913.
- Guvenen, Fatih, Fatih Karahan, Serdar Ozkan, and Jae Song.** 2019. “What Do Data on Millions of U.S. Workers Reveal about Life-Cycle Earnings Dynamics?” https://fguvenendotcom.files.wordpress.com/2019/02/gkos_2017_rev3_submit.pdf.

- Handwerker, Elizabeth Weber.** 2011. “Delaying Retirement to Pay for College.” *ILR Review*, 64(5): 921–948.
- Havránek, Tomáš.** 2015. “Measuring Intertemporal Substitution: The Importance of Method Choices and Selective Reporting.” *Journal of the European Economic Association*, 13(6): 1180–1204.
- Hoxby, Caroline M, and George B Bulman.** 2016. “The Effects of the Tax Deduction for Postsecondary Tuition: Implications for Structuring Tax-based Aid.” *Economics of Education Review*, 51: 23–60.
- Hoynes, Hilary W, and Erzo FP Luttmer.** 2011. “The Insurance Value of State Tax-and-Transfer Programs.” *Journal of Public Economics*, 95(11-12): 1466–1484.
- Hubbard, R Glenn, Jonathan Skinner, and Stephen P Zeldes.** 1994. “The Importance of Precautionary Motives in Explaining Individual and Aggregate Saving.” *Carnegie-Rochester Conference Series on Public Policy*, 40: 59–125.
- Hubbard, R Glenn, Jonathan Skinner, and Stephen P Zeldes.** 1995. “Precautionary Saving and Social Insurance.” *Journal of Political Economy*, 103(2): 360–399.
- Hurst, Erik, and Paul Willen.** 2007. “Social Security and Unsecured Debt.” *Journal of Public Economics*, 91(7-8): 1273–1297.
- Kimball, Miles S.** 1990. “Precautionary Saving in the Small and in the Large.” *Econometrica*, 58(1): 53–73.
- Long, Mark.** 2004. “The Impact of Asset-tested College Financial Aid on Household Savings.” *Journal of Public Economics*, 88(1-2): 63–88.
- Monks, James.** 2004. “An Empirical Examination of the Impact of College Financial Aid on Family Savings.” *National Tax Journal*, 57(2): 189–207.
- Murphy, Kevin M., and Finis Welch.** 1990. “Empirical Age-Earnings Profiles.” *Journal of Labor Economics*, 8(2): 202–229.
- National Association of Student Financial Aid Administrators.** 2017. “Need Analysis: Federal and Institutional Methodology.” https://www.nasfaa.org/uploads/documents/NASFAA_Authorized_Training_Need_Analysis_SSG_171D.pdf.
- Onink, Troy.** 2014. “How Assets Hurt College Aid Eligibility on FAFSA and CSS Profile.” <https://www.forbes.com/sites/troyonink/2014/02/14/how-assets-hurt-college-aid-eligibility-on-fafsa-and-css-profile/#5915e97723ec>.

- Peterman, William B.** 2016. “Reconciling Micro and Macro Estimates of the Frisch Labor Supply Elasticity.” *Economic Inquiry*, 54(1): 100–120.
- Reichling, Felix, and Charles Whalen.** 2012. “Review of Estimates of the Frisch Elasticity of Labor Supply.” Congressional Budget Office Working Paper 2012-13.
- Reyes, Jessica Wolpaw.** 2008. “College Financial Aid Rules and the Allocation of Savings.” *Education Economics*, 16(2): 167–189.
- Rostam-Afschar, Davud, and Jiaxiong Yao.** 2014. “Progressive Taxation and Precautionary Saving Over the Life Cycle.” <http://davud.rostam-afschar.de/downloads/RAY2014.pdf>.
- Souleles, Nicholas S.** 2000. “College Tuition and Household Savings and Consumption.” *Journal of Public Economics*, 77(2): 185–207.
- Stepner, Michael.** 2019. “The Insurance Value of Redistributive Taxes and Transfers.” <https://michaelstepner.com/jmp/>.
- U.S. Census Bureau.** 2019. “PINC-03. Educational Attainment-People 25 Years Old and Over, by Total Money Earnings, Work Experience, Age, Race, Hispanic Origin, and Sex.” <https://www.census.gov/data/tables/time-series/demo/income-poverty/cps-pinc/pinc-03.html>.
- Varian, Hal R.** 1980. “Redistributive Taxation as Social Insurance.” *Journal of Public Economics*, 14(1): 49–68.