Measuring the Welfare Cost of a Misperceived Tax: Evidence from College Financial Aid*

Nick Gebbia
Joseph Gray-Hancuch
Paul R. Organ

June 7, 2026

Abstract

College financial aid in the US creates a large implicit tax on parent income. Like many taxes, it is difficult to learn one's tax rate. Using federal tax data, we provide the first estimates of the response of parent income to the college aid implicit tax (elasticity = .04). We show how to combine the elasticity with survey data to identify deadweight loss from the tax, given that individuals misperceive their tax rate. Misperception of the college aid implicit tax increases deadweight loss by 127%, driven by substantial heterogeneity in misperception across individuals. (JEL D90, H21, I22)

*Gebbia: The Hoover Institution at Stanford University. Corresponding author. Email: gebbia@stanford.edu. Gray-Hancuch and Organ: US Treasury Department, Office of Tax Analysis. This paper subsumes an earlier paper by Gebbia titled "Misperception and Income Response to Means-Tested Programs: Evidence from the College Financial Aid Implicit Tax", which was chapter 1 of Gebbia's Ph.D. dissertation at UC Berkeley. This research was conducted while Gray-Hancuch and Organ were employees at the U.S. Department of the Treasury and Gebbia was a contractor established under an Intergovernmental Personnel Act (IPA) agreement with the Internal Revenue Service (IRS) Joint Statistical Research Program. The findings, interpretations, and conclusions expressed in this paper are entirely those of the authors and do not necessarily reflect the views or official positions of the U.S. Department of the Treasury or the IRS. Any taxpayer data used in this research was kept in a secured IRS data repository, and all results have been reviewed to ensure that no confidential information is disclosed. We thank Patrick Kline, Jesse Rothstein, Dmitry Taubinsky, Emmanuel Saez, Anjali Adukia, Hunt Allcott, Alan Auerbach, David Autor, B. Douglas Bernheim, Barbara Biasi, Sandra E. Black, Zachary Bleemer, Shawn Brick, George Bulman, Sydnee Caldwell, Christopher Campos, David Card, Kevin Dano, Stefano DellaVigna, Manasi Deshpande, Mark Duggan, Aaron Edlin, David Figlio, John N. Friedman, Yana Gallen, Joshua D. Gottlieb, Nathaniel Hendren, Caroline M. Hoxby, Hilary Hoynes, Johnny Huynh, Sreeraahul Kancherla, Melissa S. Kearney, Patrick J. Kennedy, Kristy Kim, Henrik Kleven, Camille Landais, Attila Lindner, Felipe Lobel, Benjamin Lockwood, Justin Marion, Carl McPherson, Amanda Pallais, Devin Pope, Joshua D. Rauh, Andrew Schwartz, Andrew Samwick, Caroline Siegel Singh, John D. Singleton, Jonathan Smith, Carolyn Stein, Juan Carlos Suárez Serrato, Jeffrey Smith, Alisa Tazhitdinova, Melanie Wasserman, Evan White, Danny Yagan, and seminar participants at UC Berkeley, the Becker Friedman Institute, US Treasury, the Hoover Institution, University of San Francisco, UC Irvine, and Stanford University for helpful discussions. Gebbia gratefully acknowledges funding support from the California Policy Lab; the Burch Center; the Stone Center; the Institute for Research on Labor and Employment; the Law, Economics, and Politics Center; and the Center for Studies in Higher Education at UC Berkeley.

Means testing of college financial aid in the US creates an implicit tax on parent income so large that it frequently doubles the effective marginal tax rate (MTR) a family would face from income taxes alone. This implicit tax arises due to the phaseout of college aid over parent income: each additional dollar a family earns around the time a child is in college will typically reduce the child's college aid by some fraction of that dollar. Each year, millions of families earning anywhere from \$40k to over \$300k can see their effective MTR increase by 30 percentage points or more due to the college aid implicit tax, dramatically reducing the incentive to earn income. At the same time, the relationship between parent income and college aid is remarkably difficult to learn, leaving substantial scope for individuals to misperceive the tax rate they face. What is the welfare cost that results from the college aid implicit tax, and how can we measure this cost if parents misperceive their tax rate? We answer this question by providing the first estimates of the response of parent income to the college aid implicit tax and developing a new method to measure the welfare consequences of this response when parents misperceive their tax rate.

First, we estimate the elasticity of taxable income (ETI), which in our case is the response of parent income to the college aid implicit tax. To do so, we construct a novel panel dataset from administrative federal tax records, the ideal data source for this setting. The panel of tax data allows us to measure parent income over time and to identify whether, when, and where a dependent child attends college, which would be rare or impossible in other large administrative data sources. We link parents with their children through dependent claiming on the form 1040, and we merge in information from 1098-T forms to measure when and where a child is in college. We then study the trajectory of various sources of parent income reported on form 1040 and W-2 records around the time a child enrolls in college, as well as how these responses relate to household and college characteristics. The rich tax data allows us to investigate which types of households are more and less responsive to college incentives as well as which income sources account for most of the overall adjustments. We focus our analysis on households whose income averaged over two baseline years of 2005 and 2006 is between \$50k and \$150k as this income range faces the greatest extent of college aid phaseout, and we also show results for lower and higher income households.

Our core empirical approach is to estimate an event study of parent income over child age, and we identify responses to college incentives by leveraging sharp changes in incentives faced by parents as a child reaches college age. In the period we study, each year of financial aid depends on the family's income in the preceding year. For a child who begins college at age 18, this means that income earned when the child is 17 suddenly begins to affect the child's college aid, whereas income earned in all preceding years did not affect the child's aid. We use the sharp introduction of the college aid implicit tax when the child is 17 to identify the effect of the college aid implicit tax on parent income. Crucially, college costs are not due at age 17, but first become due when the child is 18. We focus on the response of income at age 17 to isolate the response of parent income to the college aid implicit tax from any change in parent income that may arise from the introduction of college costs. We produce separate estimates of the response of parent income to the introduction of college costs using income changes at age 18.

For families whose child enrolls in college at age 18, we observe a striking pattern in which parent income suddenly falls when the child is age 17 (the first year income

affects aid) and subsequently increases when the child is 18 (the first year college costs are due). Both responses are larger for families whose child attends a college with a high rate of aid phaseout or high cost, whereas the same effects are not observed for families whose child does not attend college at age 18, implying that these patterns are not due to other common age-related phenomena. Focusing on families whose child attends a 4-year college, we find that the age 17 income reduction is almost entirely explained by non-wage income sources like business income, IRA distributions, and capital gains, with no reduction in total wage compensation. Roughly 60% of the age 18 income increase comes from non-wage sources, while 40% comes from increasing total wage compensation.

For families whose child attends college at age 18, parent income falls by 0.7% relative to trend when the child is age 17. The net-of-tax rate falls by 15% at the same time. Scaling the effect on income at age 17 by the change in the net-of-tax rate produces our main elasticity estimate of .04. Families that have a high share of income from non-wage sources at baseline are twice as elastic (.08).

Our elasticity estimates are not driven by selection into college. We show that there is a clear reduction of income at age 17 in the full sample that does not condition on the child's college attendance, which cannot be explained by selection. Moreover, if the age 17 income reduction among college families were driven by selection, we would expect to see a countervailing income increase at age 17 among non-college families, which we do not see. That selection into college does not explain our elasticity estimates is consistent with prior work finding that family income shocks when a child is in high school have little effect on college enrollment (Hilger 2016; Bulman et al. 2021), and the effect is the opposite sign of our result: a positive income shock marginally increases college enrollment, while we document that college families have lower income at age 17.

What do our elasticity estimates mean for the welfare cost involved in means testing college aid? In standard public economics, evaluating the deadweight loss (DWL) produced by a tax requires only one sufficient statistic: the ETI (Feldstein 1999). However, advances in behavioral public economics have shown that this relationship breaks down when individuals do not pay full attention to the tax or when individuals pay attention but misperceive their tax rate (Chetty et al. 2009; Taubinsky and Rees-Jones 2018; Farhi and Gabaix 2020). College financial aid is an exemplary setting in which such behavioral frictions are likely at play, as it is remarkably difficult to learn one's tax rate from aid phaseout. Specifically, while college aid is high-stakes and therefore likely to be *salient* to most families (i.e. they consider the fact that aid depends on income), the aid schedule is difficult to understand and creates much room for families to *misperceive* their tax rate (i.e. the tax rate they think they face differs from their actual tax rate). If families misperceive the tax rate they face from college aid phaseout, then the ETI we estimate does not answer our motivating question regarding the welfare cost of the college aid implicit tax; rather it is only the first of two necessary ingredients.

To identify DWL from the college aid implicit tax under misperception, we must combine the ETI we estimate with the second necessary ingredient: measures of tax rate perceptions. However, there is no preceding paper that provides a method to quantify the welfare cost of a tax that individuals heterogeneously misperceive without imposing structure on the form of misperceptions. Rees-Jones and Taubinsky (2020) consider misperception of the federal income tax schedule by surveying indi-

viduals on their perceived tax rates, but their welfare analysis relies on a particular structural form of misperception: individuals either correctly perceive their marginal tax rate or they incorrectly believe their average tax rate applies at the margin. While this structure is amenable to studying misperception of the federal income tax - individuals know both their income and total tax liability, and might easily compute their average tax rate and assume it applies at the margin - many taxes do not readily admit a structural form for misperceptions. Implicit taxes created by means-tested benefit programs, like college aid, are especially unlikely to produce a simple structural form of misperceptions. The closest example of a paper that does not impose structure on behavioral frictions is Taubinsky and Rees-Jones (2018). However, Taubinsky and Rees-Jones (2018) study salience of the sales tax in a choice experiment, so their methods are not readily translatable to studying misperception of a tax on income. Moreover, Taubinsky and Rees-Jones (2018) are limited to providing a lower bound on the welfare impact from heterogeneous behavioral frictions, and cannot identify the total welfare impact of these frictions. Likewise, Chetty et al. (2009) also study salience, but they do not consider heterogeneity in salience across individuals. Heterogeneity in misperceptions turns out to be a key mechanism in our setting, so the methods of Chetty et al. (2009) are similarly not applicable.

We address this gap by developing a new method to measure the welfare cost of a tax that is misperceived with unrestricted heterogeneity. Our method is easily portable to studying the impact of misperception on the welfare cost of any tax.

Developing the method proceeds in two parts. First, we derive the sufficient statistics for welfare when individuals misperceive their tax rate with nearly unrestricted heterogeneity. In addition to increasing with the ETI, DWL also increases with each of the following: the average misperception; the variance of misperceptions across individuals; and the covariances of squared misperception with income, elasticity, and squared tax rate. The distribution of misperceptions matters: because DWL is convex in tax-induced distortions, heterogeneity in perceived tax rates increases aggregate inefficiency even if the average perception is unbiased. The theoretical result that variance and covariances of behavioral frictions can impact welfare echoes those in Taubinsky and Rees-Jones (2018), Allcott et al. (2019), Farhi and Gabaix (2020), and Allcott et al. (2025).

Second - and our key methodological contribution - we show how to identify each of these components using survey data on perceptions. An important concern with using survey data to measure perceptions is that stated beliefs may be noisy approximations of underlying beliefs. This has first order implications for our setting in which we must measure variance and covariance terms: even if measurement error on stated perceptions is mean zero, simply computing sample variance and covariances will produce biased estimates, because these will combine variation in perceptions with variation in measurement error. We resolve this issue by incorporating a secondary elicitation, which can be thought of as an instrument. Crucially, the secondary elicitation measures the respondent's perception of a tax rate that is slightly different from the focal tax rate. This reduces the likelihood that measurement error is correlated between the two elicitations, but introduces bias terms reflecting how the respondent's perception of the gap between the two rates covaries with the perceived focal rate, income, elasticity, and the true focal rate. We prove that, under a set of conditional independence assumptions, we can use the secondary elicitation to purge the influence of measurement error and identify second moments

of misperceptions, while accounting for identifiable bias correction terms.

To apply our method, we administer a survey to parents via Prolific Academic in which we measure parents' perceptions of the tax rate they face from college aid phaseout. We focus our survey analysis on the 1,279 respondents whose income is between \$50k and \$150k, matching our main tax data sample, and whose child attends or is likely to attend a four-year in-state public college. Our survey elicitations then ask respondents for their belief of average financial aid awarded to an in-state student attending a four-year public college in their state at several points of family income. By aligning our survey questions on aid at four-year in-state public colleges with our sample of respondents whose child attends or is likely to attend such a college, we ensure that the questions are both realistic and relevant for the respondent, under which conditions stated beliefs are more likely to reflect actual beliefs that guide behavior (Mas and Pallais 2017; Fuster et al. 2021; Stantcheva 2021; Fuster and Zafar 2023). Nevertheless, our method is explicitly designed to account for measurement error in reported perceptions through use of the instrument. Additionally, the fact that we ask respondents for their perception of average aid at several points of family income rather than directly asking for their perception of the marginal aid phaseout rate produces a couple benefits in terms of measurement error. First, because average aid is a simpler concept, there is likely to be less measurement error than if we asked directly for perceptions of phaseout rates. Second, because our primary and secondary elicitations ask for perceptions of average aid at several points of income that differ between elicitations, respondents cannot simply recall the tax rate they reported in the first elicitation and repeat it in the second, therefore reducing concerns that measurement error is importantly correlated between the two measures. While we cannot rule out some degree of correlation in measurement error between elicitations, our design minimizes the likelihood that any remaining correlation would drive our results.

Finally, we return to our motivating question, and we measure DWL from the college aid implicit tax by combining the two necessary ingredients: the ETI and the relevant survey moments of perceptions. We focus our DWL calculations among families with income between \$50k and \$150k whose child attends a four-year college. If all families had perfectly accurate perceptions of aid phaseout rates, annual DWL among this population would be \$109 million. However, misperception increases DWL by \$138 million – or 127% – to a total of \$247 million annually. The dramatic increase in DWL from misperception is driven by heterogeneity in misperception across individuals. In fact, we find that parents underestimate their tax rate on average, which alone reduces DWL by 21%. We also find a negative correlation between squared misperceptions and the square of the true tax rate, which alone reduces DWL another 33%. However, we find that the total variance of misperceptions across individuals is substantial enough to increase DWL by 181%.

This paper contributes to two main literatures. First, this paper contributes to the behavioral public economics literature that studies how information frictions affect the welfare consequences of taxation. Existing theory shows that the ETI no longer identifies DWL when there are behavioral frictions, and emphasizes that behavioral frictions can change the welfare consequences of taxation and alter optimal tax policy (Farhi and Gabaix 2020; Craig and Slemrod 2024). A host of empirical evidence has documented misperception of tax rates across various settings, either in direct surveys of perceptions (Fujii and Hawley 1988; Slemrod 2006; Rees-Jones

and Taubinsky 2020) or in reduced form behavior (de Bartolome 1995; Liebman and Zeckhauser 2004; Liebman and Luttmer 2015; Feldman et al. 2016; Kostøl and Myhre 2021).

Missing from the behavioral public economics literature is a method to empirically measure the welfare consequences of tax rate misperception without imposing structure on the form of misperceptions. To this point, while the theory provided by Farhi and Gabaix (2020) similarly shows that heterogeneity in tax rate misperceptions affects the optimal tax rate, they close with the following: "One upshot of this paper is that numerous quantities can in principle have a big impact on optimal policy, but have scarcely or not yet been measured. Measuring those quantities presents an exciting research opportunity."[1]

We show how to measure the welfare cost of a tax that individuals can misperceive heterogeneously, as well as the contribution to DWL due to misperception, by combining a reduced form ETI estimate with individual-level survey measures of perceptions. The approach we introduce can easily be used to measure the welfare consequences due to misperception of other explicit or implicit taxes, including the Social Security earnings test (Liebman and Luttmer 2015), Child Tax Credit (Feldman et al. 2016), or Earned Income Tax Credit (Chetty et al. 2013). Our empirical results are the first to find and quantify a large *increase* in DWL due to tax rate misperception, which is due to substantial heterogeneity in misperceptions. This is counter to earlier results in which structural forms of misperception of the federal income tax tend to reduce DWL (de Bartolome 1995; Rees-Jones and Taubinsky 2020), but it is consistent with the emerging theme in behavioral public finance that variance in behavioral frictions can be more consequential for welfare than the average friction, as recently shown for sales tax salience (Taubinsky and Rees-Jones 2018) and nudges (Allcott et al. 2025). While recent lab evidence suggests that individuals exert costly effort to more accurately learn tax incentives when the stakes are higher (Morrison and Taubinsky 2023), our finding of substantial misperception in the high-stakes setting of college aid demonstrates that misperception may significantly impact welfare even when individuals face strong incentives to overcome behavioral frictions.

Second, this paper builds on and bridges two strands of research related to college aid and prices: those studying the implicit tax created by means testing college aid, and those studying the implications of opaqueness in college aid. It has been observed for decades that college aid is structured as an implicit tax on income and savings (Case and McPherson 1986; Edlin 1993; Dick and Edlin 1997). Several early papers focused on the elasticity of savings with respect to college aid incentives (Feldstein 1995; Long 2004; Monks 2004). Two recent papers have studied whether students change their labor supply during college in response to similar aid incentives placed on student income (Darolia 2017; Park and Scott-Clayton 2018). Despite the long acknowledgement of these incentives, ours is the first paper to estimate the elasticity of parent income with respect to the college aid implicit tax. Because parent income explains the majority of variation in college aid, quasi-experimental estimates of the elasticity of parent income are necessary for evaluating the tradeoff between

[1] Rees-Jones and Taubinsky (2018) close with a similar call: "As the literature progresses, an iterative application of the sufficient statistics approach to welfare, paired with granular measurement of heterogeneous biases in tax settings, appears to be both a conceptually justified and practically implementable approach to the development of empirically informed tax policy."

efficiency and equity or insurance value inherent in means-tested college aid and to build upon existing assessments of the optimal aid schedule (e.g. Colas et al. 2021; Fan et al. 2025). Additionally, our analysis of misperception of the implicit tax in college aid establishes a connection between the literature on college aid incentives and the literature on implications of college aid opaqueness. Several papers have documented that information frictions related to the *level* of college aid are important enough as to have considerable impacts on major life choices such as students' college application and enrollment decisions (Dynarski and Scott-Clayton 2006; Bettinger et al. 2012; Hoxby and Avery 2012; Hoxby and Turner 2013; Dynarski et al. 2021). A recent book makes the same point (Levine 2022). Our analysis studies information frictions related to the *slope* of college aid over income, finding consistent evidence of substantial misperceptions of the aid schedule with important welfare implications.

The remainder of the paper proceeds as follows. Section 1 describes the large implicit tax in college aid. Section 2 details the difficulty of learning one's tax rate from college aid. Section 3 uses federal tax data to estimate the ETI with respect to the college aid implicit tax. Section 4 develops our method to measure the welfare cost of taxation under misperception, and then combines the ETI with survey data to identify DWL. Section 5 concludes.

1 The Large Implicit Tax on Income from College Aid

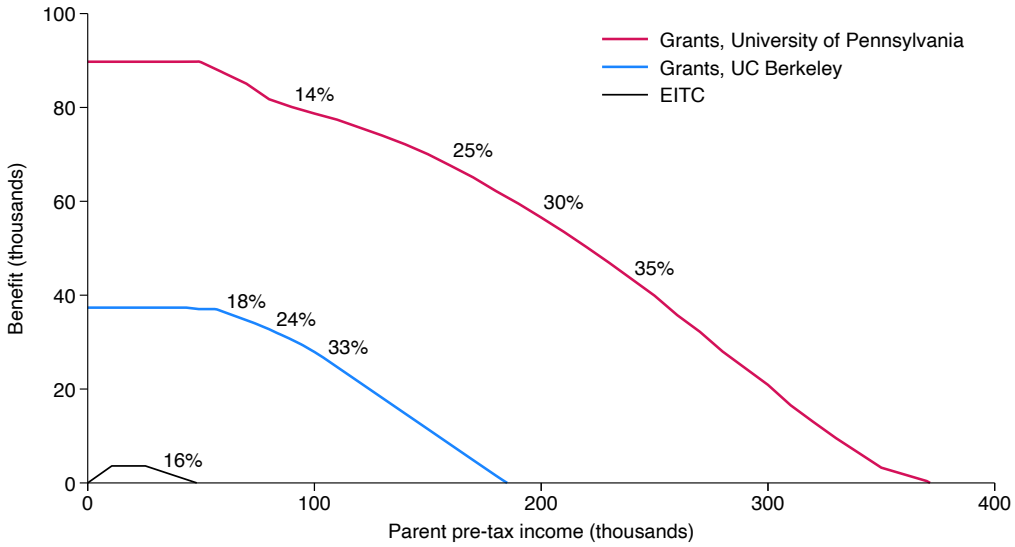
Most students who attend college in the US do not pay the full "sticker cost". A college's sticker cost is the published annual cost to attend that college, which includes tuition and fees, housing, and other important costs like books, supplies, and transportation. Most students receive financial aid that lowers the effective cost they pay. Financial aid comprises grants or scholarships – money that covers college costs and does not need to be paid back – along with loans that are typically offered on favorable terms. Financial aid is available to students through sources including the federal government, state governments, individual colleges, and private entities. In the remainder of this paper, we will use the terms "financial aid" and "aid" to refer only to grants and scholarships offered by the federal government, state governments, and individual colleges, as this is the component that creates high-stakes incentives over income earned by the parents of college students.

The financial aid a student receives is heavily means tested against the income of their parents, so that two students with parents earning different incomes can pay vastly different net costs to attend the same college. Figure 1 displays the total grant aid a student can expect to receive at the University of Pennsylvania (in red), a private four-year university, or at UC Berkeley (in blue), a public four-year university, per year of enrollment based on their parents' pre-tax income. The information in Figure 1 is based on the Net Price Calculator (NPC) page on each university's financial aid website.[2] The annual total cost of attendance at the University of Pennsylvania (UC

[2] All colleges in the US that participate in Title IV financial aid programs – allowing their students to access federal grants, federal student loans, and federal work-study to assist in paying for college – are required to host a NPC website that allows students to enter several aspects of their family background and receive a tailored estimate of the financial aid they will receive and the net cost they will pay. As discussed in Section 2, however, information on the broader aid schedule over

Berkeley) – including tuition and fees, housing, and other expenses – is approximately $94k ($47k). According to the respective NPC websites, students whose parent pre-tax income is below approximately $50k can expect to receive approximately $91k ($37k) in grants per year, cutting their annual cost of attendance down to $3k ($10k). Beginning at incomes slightly above $50k, aid is gradually phased out until it reaches zero for students whose parent income is above roughly $350k ($180k), in which case the family's annual net cost is equal to the full sticker cost of $94k ($47k).

Figure 1: The Size of the Implicit Tax on Income from College Aid



Notes: This figure shows the size of the college aid implicit tax on parent income using the aid schedules at two colleges as examples. The top line in the figure shows the expected annual grant aid for one child enrolled at the University of Pennsylvania as a function of parent pre-tax income. The middle line is the expected annual grant aid for one child enrolled at UC Berkeley. As a reference point, the bottom line is the EITC schedule for a single parent with one child. Marginal implicit tax rates (i.e. the share of a marginal dollar of income that is "lost" to reduced aid) are labeled in the figure. We derive the estimated college aid schedules via the colleges' respective Net Price Calculator websites by producing aid estimates for various points of income, while holding assets fixed at zero and adjusting for estimated income taxes.

Means testing of financial aid over parent income creates a strong incentive for middle- and higher-income families with a child in college to reduce reported income. Figure 1 labels the marginal implicit tax rates on parent income that occur due to the phaseout of college aid. In the case of the University of Pennsylvania (UC Berkeley), for families with income between $50k and $350k ($180k), each additional dollar of reported income reduces the financial aid award (and thus increases the family's net cost) by 14 to 35 cents (18 to 33 cents). This phaseout of aid is akin to implicitly taxing income at the margin during each year of income reported on a student's Free Application for Federal Student Aid (FAFSA), with marginal implicit tax rates ranging from 14% to 35% (18% to 33%). The marginal implicit tax from college aid adds on top of other taxes a family faces such as federal and state income taxes

income like that shown in Figure 1 can only be obtained by repeatedly entering different inputs to the NPC and recording the various aid estimates. Additionally, as discussed in Section 2, college aid varies not only with income but also with family savings, family structure, and other family characteristics. The information in Figure 1 was produced by holding these other inputs fixed.

and FICA, so that middle- and higher-income families with a child in college can face an all-in marginal tax on income reaching up to 70%, dramatically reducing the incentive to earn income (see Figure A.1). While Figure 1 demonstrates the implicit tax using the University of Pennsylvania and UC Berkeley as an example, the college aid implicit tax applies nationwide and is steadily increasing at middle and higher incomes (see Figures A.2 and A.3).[3]

The implicit tax on income from college aid is large relative to other means-tested programs in terms of both the magnitude of marginal implicit tax rates and the wide range of income levels subject to them. For comparison, Figure 1 also shows the Earned Income Tax Credit (EITC) schedule for a family with one child. The EITC phases out at a marginal implicit tax rate of 16%. Phaseout of the EITC begins at incomes just below $30k and extends up to approximately $50k. In comparison, college aid phaseout applies larger marginal implicit tax rates to a range of income that can extend roughly six (UC Berkeley) to fifteen (University of Pennsylvania) times as wide as that affected by EITC phaseout.

Finally, the college aid implicit tax affects millions of families each year, with billions of dollars being allocated under this mechanism. As of fall 2021, there were 8 million students age 24 and under enrolled full-time as undergraduates at degree-granting institutions (National Center for Education Statistics 2024). Of these, 4.8 million attended a 4-year public, 1.9 million attended a 4-year (nonprofit) private, and 1.1 million attended a 2-year public college. As of 2022, a total of $117.2 billion in grant aid was awarded to undergraduates from federal ($38 billion), state ($14.7 billion), and institutional ($64.5 billion) sources (College Board 2024). As a reference point, in tax year 2022, 23 million workers and families received a total $57 billion from the EITC (Internal Revenue Service 2023).

Despite the large implicit tax from college aid that affects millions of families each year and the long literature studying income response to similar incentives in programs like the EITC (see e.g. Eissa and Hoynes 2006; Chetty et al. 2013) and other means-tested benefits, to date there are no empirical estimates of how parent income responds to the college aid implicit tax.

2 The Rate of Aid Phaseout Is Difficult to Learn

While it is simple to understand that college aid is related to parent income, it is remarkably difficult to learn the exact tax rate one faces from college aid phaseout. The difficulties involved in learning one's tax rate from college aid phaseout are well exemplified in the federal Expected Family Contribution (EFC) formula, which determines federal Pell grant eligibility and is an important reference point in state and institutional aid.[4] The EFC formula, and aid generally, is highly nonlinear, including interactions between parent income and other inputs like assets, student income, and more.[5] Figure 2 shows a brief excerpt from the 2015 EFC formula, reflecting the various inputs, allowance amounts calculated from tables not shown in

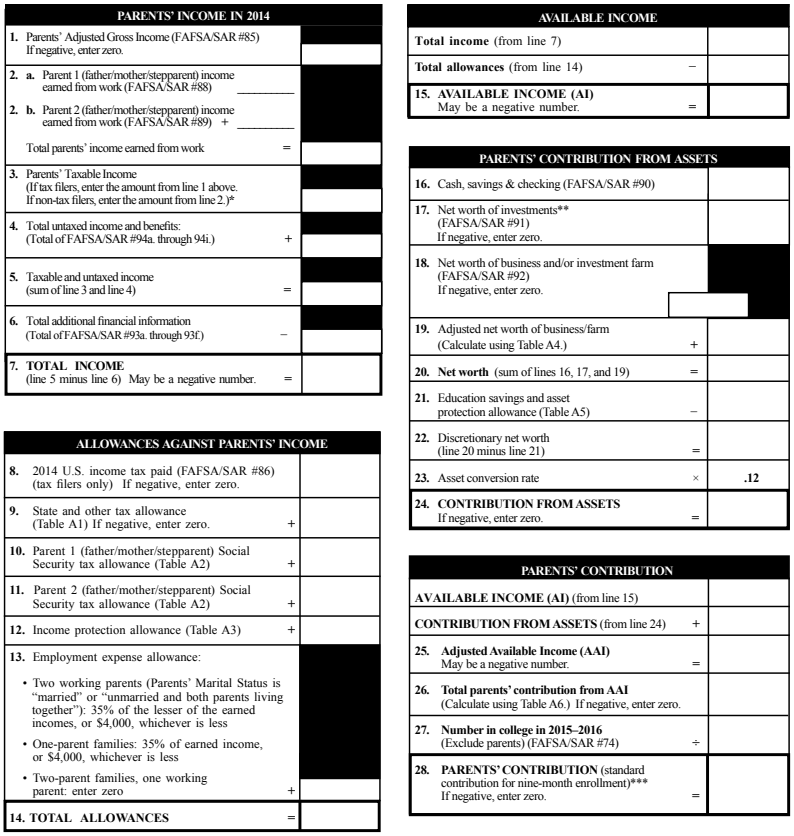
[3] Cook and Turner (2022) similarly document an increase in the progessivity of pricing at US colleges, in a push toward what they call "high-tuition, high-aid" pricing.

[4] Starting with the 2024 academic year, EFC was replaced with the Student Aid Index (SAI). SAI does not significantly alter the relationship between family income and college aid described in this section. We refer to the EFC formula because it was in place during our sample period.

[5] Figure 1 holds other inputs fixed and only varies parent income and associated income taxes.

the figure, conversion rates, and more that ultimately determine EFC and impact aid. Adding to the difficulty of learning one's aid phaseout rate is the fact that state and institutional aid may deviate from the EFC formula in unclear ways and such differences may vary across institutions.

Figure 2: Excerpt of 2015 Expected Family Contribution (EFC) Formula



PARENTS' INCOME IN 2014		
1. Parents' Adjusted Gross Income (FAFSA/SAR #85) If negative, enter zero.		
2. a. Parent 1 (father/mother/stepparent) income earned from work (FAFSA/SAR #88) _____		
2. b. Parent 2 (father/mother/stepparent) income earned from work (FAFSA/SAR #89) + _____		
Total parents' income earned from work	=	
3. Parents' Taxable Income (If tax filers, enter the amount from line 1 above. If non-tax filers, enter the amount from line 2.)*		
4. Total untaxed income and benefits: (Total of FAFSA/SAR #94a. through 94i.)	+	
5. Taxable and untaxed income (sum of line 3 and line 4)	=	
6. Total additional financial information (Total of FAFSA/SAR #93a. through 93f.)	–	
7. TOTAL INCOME (line 5 minus line 6) May be a negative number.	=	

ALLOWANCES AGAINST PARENTS' INCOME		
8. 2014 U.S. income tax paid (FAFSA/SAR #86) (tax filers only) If negative, enter zero.		
9. State and other tax allowance (Table A1) If negative, enter zero.	+	
10. Parent 1 (father/mother/stepparent) Social Security tax allowance (Table A2)	+	
11. Parent 2 (father/mother/stepparent) Social Security tax allowance (Table A2)	+	
12. Income protection allowance (Table A3)	+	
13. Employment expense allowance: • Two working parents (Parents' Marital Status is "married" or "unmarried and both parents living together"): 35% of the lesser of the earned incomes, or $4,000, whichever is less • One-parent families: 35% of earned income, or $4,000, whichever is less • Two-parent families, one working parent: enter zero	+	
14. TOTAL ALLOWANCES	=	

AVAILABLE INCOME		
Total income (from line 7)		
Total allowances (from line 14)	–	
15. AVAILABLE INCOME (AI) May be a negative number.	=	

PARENTS' CONTRIBUTION FROM ASSETS		
16. Cash, savings & checking (FAFSA/SAR #90)		
17. Net worth of investments** (FAFSA/SAR #91) If negative, enter zero.		
18. Net worth of business and/or investment farm (FAFSA/SAR #92) If negative, enter zero.		
19. Adjusted net worth of business/farm (Calculate using Table A4.)	+	
20. Net worth (sum of lines 16, 17, and 19)	=	
21. Education savings and asset protection allowance (Table A5)	–	
22. Discretionary net worth (line 20 minus line 21)	=	
23. Asset conversion rate	×	.12
24. CONTRIBUTION FROM ASSETS If negative, enter zero.	=	

PARENTS' CONTRIBUTION		
AVAILABLE INCOME (AI) (from line 15)		
CONTRIBUTION FROM ASSETS (from line 24)	+	
25. Adjusted Available Income (AAI) May be a negative number.	=	
26. Total parents' contribution from AAI (Calculate using Table A6.) If negative, enter zero.		
27. Number in college in 2015–2016 (Exclude parents) (FAFSA/SAR #74)	÷	
28. PARENTS' CONTRIBUTION (standard contribution for nine-month enrollment)*** If negative, enter zero.	=	

Notes: This figure shows a brief excerpt from the Expected Family Contribution (EFC) formula, a key formula that impacts how college aid relates to parent income and other factors. From The EFC Formula, 2015-2016. Source: U.S. Department of Education, Federal Student Aid.

Information sources available for learning one's MTR from college aid are not transparent and may even add to confusion. One potentially common way for parents to learn their MTR from aid would be through online news articles. Figure A.4 shows excerpts from two online news articles describing how families can maximize their financial aid. Both excerpts make it clear to understand that there is an incentive to temporarily reduce income with section headers reading "Postpone parental income" and "Minimize income in the base year." However, when it comes to describing the size of the incentive, the first piece includes descriptions of the average tax rate (in this context, the percent of income a family is expected to pay toward college, but not at the margin), while the second piece only lists what is roughly the maximum marginal implicit tax rate of 30%. For many families, these numbers will not accurately reflect the MTR they will face from college aid, so articles like these may even exacerbate misperception.

Perhaps the most effective way for parents to learn the MTR they should expect to face from college aid is to use a college's Net Price Calculator (NPC) website, but this requires a degree of savvy. All colleges that receive Title IV federal funds must

host an NPC, where families can enter information on their income, assets, and other aid inputs and receive an estimate of their financial aid and net price. Colleges do not typically post images or information directly conveying the content of the aid schedule over income as in Figure 1 on their financial aid websites; however, if a family is sufficiently savvy, they may think to navigate to a college's NPC, enter one set of income, assets, and other inputs and record the aid estimate that is produced, and then re-enter their inputs with a small change to income and record how predicted aid changes. [6] Even this method can prove difficult, as some colleges use a template for their NPC website that only allows the user to specify one of several income ranges, rather than inputting an exact income amount.

A body of literature documents that families struggle to understand the *level* of aid they will receive and finds that this impacts college decisions (Dynarski and Scott-Clayton 2006; Bettinger et al. 2012; Hoxby and Avery 2012; Hoxby and Turner 2013; Dynarski et al. 2021), and that even existing information tools like Net Price Calculators leave significant uncertainty for families regarding the *level* of aid their child will receive (Levine 2022). It is therefore very unlikely that all families know the exact tax rate they face from aid phaseout, which is the *slope* of aid over income. After estimating the elasticity of parent income due to the college aid implicit tax in Section 3, we return in Section 4 to understanding the implications of this elasticity for welfare if families misperceive their tax rate.

3 Parent Income Responses to College Incentives

3.1 Data

Our analysis of parent income responses to college incentives relies on federal income tax data, which is ideal for our research question. The tax data allows us to link parents with their dependent children; identify the child's age along with whether, when, and where the child enrolls in college; and observe a panel of parent income along with other characteristics of the household. Below we include a summary of our data processing, main sample definition, and descriptive statistics.

3.1.1 Linking Parents and Children

We begin by defining families and their associated children.[7] We define a family as the unique combination of a primary and, if applicable, secondary parent who claimed a child. We identify links between each child and their parent(s) by collating all linkages between primary and secondary filers and the dependents listed on all tax returns. The dataset consists of all children who (1) have a valid Social Security Number (SSN) or Individual Taxpayer Identification Number (ITIN), (2) were claimed by only one family at age 14, (3) were the family's oldest child, and (4) were born between 1985 and 2004. There are approximately 29 million families whose oldest child meets the criteria in the balanced panel. The child is assigned to the family that claimed the child when the child was 14. We exclude any child claimed by multiple families at age

[6] In fact, this is how we produced Figure 1.

[7] Technically, we measure tax units and their claimed dependents. For ease of reading and consistent with prior work using tax data to study filer-dependent relationships (e.g. Chetty et al. (2020)), we use the terms "family", "parent", and "child" throughout this paper.

14, along with all associated families. This restriction lowers the number of families in the sample by 2.8 million, or 4%.

We focus on the family's oldest child, whom we refer to as the focal child. The family's oldest child is identified by selecting the oldest date of birth of any child assigned to the family. We exclude families whose oldest child is one of multiple children with the same date of birth (e.g., twins or triplets) or has a sibling born fewer than 12 months later. This restriction lowers the number of families in the sample by 0.8 million or 1.2%.[8]

3.1.2 College Attendance and College Characteristics

We measure college attendance based on the presence of the information return Form 1098-T, "Tuition Statement". The university files the 1098-T and includes such information as payments received, scholarships and grants, whether the student was at-least-part-time, and whether the student was a graduate student. We only count college attendance based on the presence of a 1098-T that records at-least-part-time undergraduate enrollment. If a student received at-least-part-time undergraduate 1098-Ts from multiple different institutions for the same tax year, we randomly select the institution to assign to the student.

A 1098-T is filed for almost every student, with some notable exceptions including individuals only taking non-credit courses and nonresident aliens. Additionally, schools are not required to file a 1098-T for any student whose tuition is paid entirely by a third-party, paid entirely by scholarships or grants, or was waived. Cronin and Gray-Hancuch (2024) find a majority of universities in 2017 file for every student, ignoring these exceptions. These institutions constitute 68% of all institutions and 65% of the undergraduate population with no clear differentiation between the institutions that file for every student and those that do not on cost-of-attendance, percent of undergraduates with scholarships or grants, average undergraduate scholarships or grants, and average Pell Grants.

Once the student is assigned to an institution for a given year, we use a crosswalk from Cronin and Gray-Hancuch (2024) to match the institution found on the 1098-T with the institution's observation in the Integrated Postsecondary Education Data System (IPEDS) from the National Center for Education Statistics for the 2017-18 academic year, which allows us to merge in college characteristics. From IPEDS, we measure the sector (2-year or 4-year, public or private, not-for-profit or for-profit) and state of the college along with information on total cost of attendance (COA) and average net price and financial aid among several family income ranges. We use sector and state to define enrollment types (for example, four-year in-state public). Regarding costs and financial aid, IPEDS records total COA and average net price and average financial aid for students with family income in each of five income ranges: \$0-30k, \$30-48k, \$48-75k, \$75-110k, and above \$110k. For our main analysis sample that has baseline income averaged over 2005 and 2006 between \$50k and \$150k, we assign a family-specific MTR from college aid for families with baseline income between \$50k and \$75k as the difference in average financial aid across the \$30-48k and \$48-75k income ranges divided by the difference in the median incomes

[8] Due to computational restrictions, families with four or more children of the same age are also excluded from the dataset, even if these are not the oldest children in the family. This restriction lowers the number of families in the sample by approximately 1,000, or 0.002%.

of these ranges, and we assign an MTR for families with baseline income between \$75-150k as the difference in average aid across the \$48-75k and \$75-110k income ranges divided by the corresponding difference in median incomes.[9] We assign a family-specific net price for families with baseline income between \$50k and \$75k using the average net price in the \$48-75k range; for families with baseline income between \$75 and \$110k using the average net price in the \$75-110k range; and for families with baseline income between \$110k and \$150k using the average net price for families with income above \$110k.

There are drawbacks and advantages to using the IPEDS data to measure the marginal tax rates families face from college aid phaseout. One drawback is the lack of granularity, as aid amounts are reported as averages over income ranges as described above, and these bins extend only up to \$110k. Additionally, the IPEDS aid measures are reported as averages among students who were awarded federal grants or loans. Because the aid measures include students who were awarded federal loans, and federal loan eligibility does not depend on family income or assets, even families with high income or high assets are represented. However, we cannot directly assess whether higher or lower income families tend to be represented within a given family income bin. Nevertheless, there are important advantages to measuring aid phaseout using the IPEDS data. The primary alternative way to measure aid phaseout at every college would be to attempt scraping data from each college's Net Price Calculator website, measuring the college's predicted aid for a student with family income at several specified values. Relative to a Net Price Calculator approach, the IPEDS measure has the advantage of being based on actual aid data as reported by each college, rather than aid estimates. IPEDS data is also easily accessible and is presented in a uniform format across all colleges, whereas the feasiblity of scraping all Net Price Calculator websites is unclear, and moreover these websites do not share a common format, meaning it is not possible to construct exactly the same measure of aid phaseout for every college. Because Net Price Calculator estimates depend on income and other inputs to aid like assets that we do not observe, we would need to make assumptions about these other factors as well. Finally, Net Price Calculators can only provide current aid estimates, they cannot provide historical aid schedules. This is important because over time the income threshold below which students tend to receive full aid and therefore face no aid phaseout has increased over time, meaning present-day Net Price Calculators would show zero aid phaseout for a range of income that in reality faced aid phaseout during our sample period. Ultimately, our approach to measuring aid phaseout using IPEDS data is similar to Cook and Turner (2022)'s choice to use the same IPEDS data to document increasing progessivity of college pricing. It is additionally worth noting that the reduced form effects we measure on parent income over child age do not depend at all on the choice of tax rate measurement.

3.1.3 Income and Family Characteristics

We use information from Form 1040 and its related schedules from 2001 to 2022 as our income measures. We supplement this data with measures of wage income

[9] For families with income between \$110-150k, we assign the MTR based on the difference in average aid across the \$48-75k and \$75-110k bins because the highest income bin for which IPEDS reports aid information includes all families with income above \$110k.

from the information return Form W-2. To avoid undue influence from outliers, we winsorize all income measures to their 99.9 and 0.1 percentiles. We also calculate ages for family members other than the oldest child, including parents and other associated children, to provide additional information about the family structure.

Our analysis focuses on parent AGI because AGI is a direct input to college aid formulas – in particular, AGI is a direct input to the EFC formula.

3.1.4 Main Analysis Sample Definition

Our main analysis sample includes families meeting the following criteria: the focal child turns 18 between 2012 and 2016 and the focal child does not attend college at age 17; baseline income, measured as average AGI in years 2005 and 2006, is between \$50k and \$150k, the range that is most exposed to the college aid implicit tax; taxes are filed each year from 2001 to 2022 with positive AGI; and the focal child is claimed from ages 7 through 17.[10] We restrict to persistent tax filers for two reasons. First, because we link parents with children by claiming when the child is 14, our sample necessarily restricts to families that file taxes when the child is 14. Without further restricting to persistent tax filers, the age 14 restriction alone would induce mean reversion that is related to the child's age, confounding our estimates. Second, we wish to avoid mis-measuring income for tax units that dissolve and then file under a different ID – for example, if a tax unit is observed filing in one year but not in subsequent years, this could be because income falls below filing limits, or it could be because two parents separate and begin filing under new tax units. We also restrict to families that claim the focal child from ages 7 through 17 to avoid potential mean reversion that could arise if claiming is endogenous. Importantly, we include robustness checks in which we remove both the persistent filing and claiming restrictions, where we find no evidence of extensive margin effects on filing with positive AGI and we show that the effects we observe on parent income in our main sample remain apparent in the unrestricted sample.

We define two primary types of college enrollment samples for our elasticity estimates. The first, which we refer to as the "Broad aid" college sample, includes families whose child enrolls at age 18 in a 2-year in-state public, 4-year in-state public, or 4-year private not-for-profit college, which has non-missing data on average net price at the college for families in the \$48-75k, \$75-110k, and above \$110k income ranges and on average aid for families in the \$30-48k, \$48-75k, and \$75-110k ranges. Non-missing net price and aid data allows us to construct estimates of the family's net price and tax rate from college aid phaseout based on their income. We focus on in-state public college enrollment because aid at public colleges is typically different for out-of-state students. Our second main college enrollment sample is the "4-year aid" sample, which includes families whose child enrolls at age 18 in a 4-year in-state public or a 4-year private not-for-profit college with the same non-missing net price and aid information.

Table A.1 shows descriptive statistics for our main analysis sample at the family level. 1.3 million families are included in our main analysis sample. Average income in the baseline years 2005 and 2006 is \$87k, of which W-2 wage income accounts

[10] We begin the claiming restriction at age 7 because this is the age of our youngest cohort (age 18 in 2012) in the first year of our data (2001). We end the claiming restriction at age 17 as claiming rates tend to be lower at ages 18 and older.

for $80k. The median age of the primary filer is 47 when the oldest child is 17. 51% of families' focal child enrolls in any college at age 18. 43% enroll at a "Broad aid" college, and 37% enroll at a "4-year aid" college. Among those in the 4-year aid sample, the average tax rate from aid phaseout is 13%, the average cost of attendance is $37k, and the average net price is $22k. The average tax rate from aid phaseout, cost of attendance, and net price are marginally lower in the Broad aid sample due to the inclusion of 2-year colleges, which tend to have lower cost and correspondingly less aid phaseout.

2-year colleges account for a minor share of enrollment in the Broad aid sample, consistent with the fact that 2-year colleges are more likely to have students whose grant aid exceeds total costs and therefore are not issued a 1098-T, as documented in Cronin and Gray-Hancuch (2024). While we produce elasticity estimates for the Broad aid and 4-year aid samples, we devote more attention to the 4-year aid sample, both for this reason and because 4-year colleges are the more important source of college aid phaseout.

3.2 Research Design

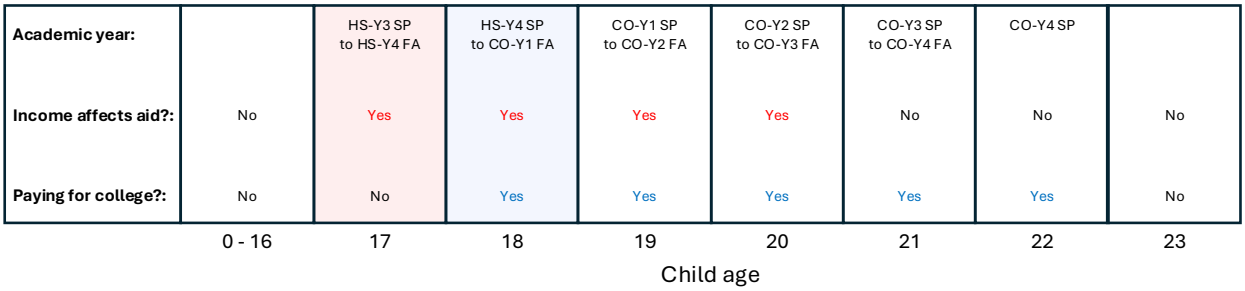
Our approach exploits sharp year-to-year changes in the college-related incentives a family faces based on their child's age, which are detailed in Figure 3. Across calendar years corresponding to the child's age as of September 1, Figure 3 shows whether parent income earned in that year will affect the child's college aid and whether the family is paying for college, assuming the child begins college immediately after high school and enrolls for four academic years. In our sample period, each year of financial aid is based on the family's income in the preceding year. From child ages 0 through 16, the family is not exposed to any sharp changes in college-related incentives. When the child is 17, parent income newly begins to affect the child's college aid.[11] However, as this year corresponds to the spring of the child's third year of high school and the fall of their fourth year of high school, the family is not yet paying for college.[12] We have shaded the period at child age 17 in red to indicate the sharp introduction of the incentive to reduce parent income as it now affects the child's college aid (substitution effect). When the child is 18, the family newly begins paying for college. Because each year of college aid references a separate year of parent

[11] Beginning with the 2017-18 academic year, the reference year of income used to determine college aid has shifted back by one year, a policy change known as "prior-prior". This means that the first year in which parent income affects a child's college aid, for a child who enrolls in college immediately after high school, is now the year in which the child is age 16. Our empirical analysis focuses on cohorts before prior-prior was introduced. This is for two reasons. First, earlier cohorts allow us to study the dynamics of parent income until the child is age 23 and likely out of college, whereas later cohorts would restrict our ability to study income dynamics during and after college. Second, our survey data described later in the paper includes a question asking respondents when they believe income begins to affect college aid. Despite the survey being fielded several years after prior-prior began, only 5% of respondents correctly identify the year when the child is 16 as the first year income affects aid. 74% of respondents still believe the first year income affects aid is the child's age 17 year, while 20% think it is the child's age 18 year. This evidence indicates that parent responses to the college aid implicit tax are unlikely to be very different after prior-prior compared with before.

[12] Moreover, the child has not yet moved out of the house to attend college, which could independently impact parent income.

income, parent income earned when the child is 18 continues to affect aid.[13] We have shaded the period at child age 18 in blue to indicate the sharp introduction of the incentive to increase income to help pay for college (income effect). Our empirical design exploits the sharp and independent introduction of the incentive to reduce income when it newly affects a child's aid at age 17 and of the incentive to increase income to help pay for college at age 18 in order to separately identify the response of parent income to these two competing incentives.[14] While both effects are likely still operating in later years as the incentives persist, we focus on age 17 and age 18 to independently measure each effect.

Figure 3: Timeline of College-related Incentives



Notes: This figure shows, over child age as of September 1, the academic terms corresponding to the matched calendar year, whether income earned in that year will affect the child's college aid (for a child who begins college immediately after high school and is enrolled for four years), and whether college bills are due. At age 17, parent income newly affects college aid, but college costs are not due. At age 18, college costs are newly due.

Nonparametric Event Study.—Motivated by the sharp timing-based changes in college-related incentives shown in Figure 3, we use an event study approach to examine whether there are accompanying sharp changes in parent income at child ages 17 and 18 for families with a college-going child. This is reflected in Equation (1). In Equation (1), i indexes families (tax units) and t indexes calendar years. Our outcome is a first difference of a given measure of income Y, with $\Delta Y_{it} = Y_{it} - Y_{i,t-1}$. Specifying our model in first differences allows for immediate separation of responses that occur at age 17 and age 18, though we include levels specifications in the appendix. β_k is the differential change in parent income at child age k relative to the omitted age 16, which is just prior to the introduction of college incentives. Because observations in which the child is older will tend to mean the parents are similarly older, we include a full set of parent age indicators to nonparametrically

[13] Students must re-apply for financial aid each year, with each year of aid being determined by a subsequent year of income. In Figure 3, income earned when the child is 17 affects aid in the first year of college, income earned when the child is 18 affects aid in the second year of college, income earned when the child is 19 affects aid in the third year of college, and income earned when the child is 20 affects aid in the fourth year of college.

[14] In theory, one could also exploit variation that arises when these incentives are removed toward the end of college. However, if there are important dynamics to the substitution or income effects, this will muddy comparisons in later years. As a result, we focus on variation around the beginning of college for our main estimates, though our results also show evidence of effects being reversed toward the end of college when incentives are removed.

control for the lifecycle pattern of income over parent age.[15] We also include controls for each calendar year to account for changes in macroeconomic factors. Because we define baseline income samples using income in specific calendar years rather than in event time, the calendar year fixed effects also control out mean reversion patterns.

$$\Delta Y_{it} = \sum_{k \neq 16} \beta_k \cdot \mathbf{1}[\text{child_age}_{it} = k] + \sum_{p} \alpha_p \cdot \mathbf{1}[\text{parent_age}_{it} = p] + \sum_{q} \gamma_q \cdot \mathbf{1}[t = q] + \epsilon_{it} \tag{1}$$

Equation (1) mirrors the specification used in the recent child penalty literature, which studies the effect of childbirth on a mother's earnings (Kleven et al. 2019; Kleven et al. 2025). A natural appeal of employing a similar specification as in the child penalty literature comes from the analogous structure of the questions: the child penalty literature and this paper both study how parent income responds to factors introduced at specific child ages, with the child penalty literature focused on childbirth (child age 0) and this paper focused on child ages 17 and 18. The variation that identifies the child age coefficients β_k in Equation (1) is therefore the same as the variation used in the child penalty literature: conditional on parent age and year, the child age coefficients β_k are identified by variation in the age at which parents had their child.

Families who are older or younger when they have their child tend to have different incomes, giving rise to pretrends; for example, Kleven et al. (2019) observe trends reflecting that individuals who are younger when they have their child tend to have lower occupational rank. One added benefit of specifying Equation (1) in first differences is that this effectively controls out any such trends that are linear in levels over parent age at childbirth. To the extent that trends remain in the first differenced specification, this reflects a relationship between income growth and parent age at child birth. The patterns we observe below are consistent with parents who are younger at child birth having not only lower income levels, but also slower income growth, which in turn is driven by slower growth of wage income. Because this relationship between parent age at child birth and income growth appears steady, identification of the effects of college-related incentives on parent income relies on sharp changes in income relative to trend when college incentives are introduced at child ages 17 and 18. We follow Dobkin et al. (2018) and refer to Equation (1) as our nonparametric event study in which trends can be assessed visually and flexibly, and we form our main estimates of the effects of college aid incentives from a parametric event study that controls for a linear trend (Dobkin et al. 2018; Freyaldenhoven et al. 2025), reflecting a linear relationship between parent age at child birth and income growth.

Parametric Event Study.—Following visual evidence of slight linear pretrends in our nonparametric event studies using our main outcome variable of first-differenced log income – reflecting a steady linear relationship between income growth and parent age at child birth – we estimate the effects of college incentives on parent income in the following parametric event study:

[15] We control for the age of the primary filer. We bin parent ages at 18 and below and at 80 and above, and we include a full set of indicators for parent age at these two binned endpoints and each age from 19 to 79.

$$\Delta Y_{it} = \sum_{k \notin [8,16]} \beta_k \cdot \mathbf{1}[\text{child_age}_{it} = k] + \lambda \cdot \text{child_age}_{it} + \sum_{p} \alpha_p \cdot \mathbf{1}[\text{parent_age}_{it} = p] + \sum_{q} \gamma_q \cdot \mathbf{1}[t = q] + \nu_{it} \quad (2)$$

Equation (2) is identical to the nonparametric Equation (1), except that here we omit indicators for child ages 8 through 16 and add a linear control in child age to account for the relationship between income growth and parent age at child birth. Conditional on parent age and year, this approach fits a linear trend in ΔY_{it} over child age using child ages 8 through 16 and identifies subsequent child age effects relative to the trend.

While this approach can be sensitive to the specification of the trend, several factors provide reassurance that our main estimates are robust. First, trend extrapolation will be more sensitive with a shorter pre-treatment period and when effects are measured over a longer horizon. Our long pre-treatment period and focus on immediate effects leave less scope for trend sensitivity. Second, our nonparametric event study results show visually clear evidence of a slight but very stable linear pretrend, followed by stark trend breaks immediately at child ages 17 and 18. We also estimate a version of our nonparametric specification where the outcome variable is the second difference of log income. In line with the first-differenced pretrend being stable and linear, we find a flat pretrend in the second-differenced specification, with identical main effect estimates at age 17. Third, consistent with our long pre-treatment period and focus on immediate effects, we find that our main results are unchanged when we allow for a quadratic trend in child age. Finally, we show that our estimated effects scale with the size of the incentives: we do not see the same effects in placebo families whose child who does not attend college, and effects among the treated families are larger at colleges that have more aid phaseout and higher total costs.

Our timing-based approach offers important advantages compared with other designs we might have pursued, particularly for estimating the substitution response of parent income to college aid phaseout. As discussed in Section 2, it is difficult for families to know the full college aid schedule, though it is simpler to understand broadly that college aid is tied to parent income. As a result, the elasticity of parent income to college aid phaseout as a whole may be different from the elasticity of parent income to changes in more nuanced aspects of the aid schedule, where the latter is likely to be smaller. Our timing-based strategy identifies the former response of parent income to aid phaseout as a whole, whereas an approach based on bunching at kinks in the aid schedule or one leveraging more specific variation in aid formulas would identify the latter response of parent income to these aspects. As a result, the elasticity we estimate is relative to a counterfactual in which aid is not means-tested against parent income.

Our design aims to isolate the substitution response to college aid phaseout from any income effect due to college costs based on the sharp and independent introduction of aid phaseout incentives at child age 17 and college costs at child age 18. It is possible that parents could anticipate college costs and begin increasing income earlier in order to smooth the shock of college costs. Because our main estimates are in first differences and are relative to income changes at earlier child ages, if income

is already shifted by a constant level in earlier years or if it is growing at a constant rate due to anticipating college costs, neither of these will confound our estimate of the substitution response at age 17. Our estimate of the substitution response could be confounded by an income effect if this causes a sudden increase in the growth rate of income at age 17 relative to prior years. To the extent that this occurs, it would cause us to underestimate the magnitude of the substitution elasticity. Our pattern of results by income type – non-wage income falls suddenly at age 17 and increases at age 18, while wage income does not respond at age 17 and increases at age 18 – is consistent with a sudden income effect primarily occurring at age 18, but we cannot rule out some effect at age 17. Because families may increase income in earlier years in order to save for college, our estimate of the income effect at age 18 is not intended to capture the entire lifecycle response to anticipated college costs. Rather, it is intended as illustrative evidence that immediate income effect responses appear to occur primarily at age 18 rather than age 17 which is our focal year for the substitution response, and as a parameter that is of interest in its own right similar to recent work studying how parents use concurrent debt to help finance a child's college (Bhargava et al. 2025).

We address further potential identification concerns after presenting our main results below.

When estimating our non-parametric and parametric specifications, we balance our data in calendar years from 2001 through 2022. Because we use variation across several cohorts, the earliest (latest) child age effects are only represented among the latest (earliest) born cohorts. We estimate our regressions with endcaps that bin the unbalanced child ages on each end (Miller 2023), and below we only report effects among balanced child age ranges. The balanced age ranges are 7 through 23 in levels specifications, 8 through 23 in first differences, and 9 through 23 in second differences. In all regressions, we cluster standard errors at the tax unit level.

3.3 Results

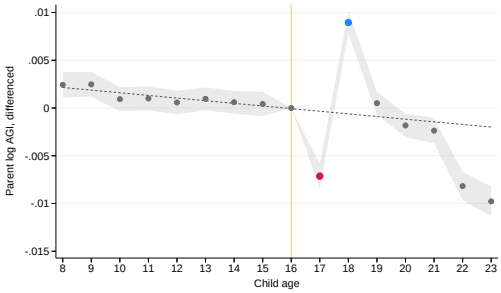
3.3.1 The Effects of College on Parent Income

Figure 4 shows our main estimates that reflect how parent income responds to college incentives. Each panel of Figure 4 reports estimates from our nonparametric regression specification over child age as of September 1, where the outcome variable is first differences of log AGI. Because the outcome is in first differences, the cumulative effect at any point comes from summing effects up to that point. We overlay a linear trend in each figure that is fitted to child age coefficients from ages 8 through 16, so that deviations from this trend at ages 17 and 18 reflect the estimates from our parametric specification. We highlight the point at age 17 in red and the point at age 18 in blue to indicate the introduction of the college aid implicit tax and of college costs, respectively, corresponding to the shading in Figure 3. The vertical yellow line at age 16 separates the pre-treatment period (ages 8-16) from the period when college-related incentives are in place (ages 17+). The six panels of Figure 4 present estimates for different subpopulations that face varying degrees of college-related incentives.

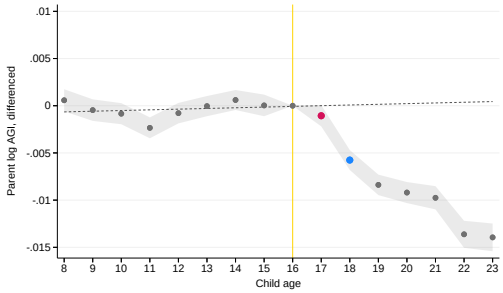
The patterns presented in Figure 4 are striking in their consistency with the incentives outlined in Figure 3. Panel A shows estimates for families whose child

attends a 4-year college. As discussed in Section 3.2, because our variation comes from comparing families who had their child at different parent ages conditional on the year and parent age at the time of each income observation, we observe slight pretrends. For the 4-year college sample, there is a slight downward pretrend in first differenced log income, reflecting that parents who had their child at a younger age tend to have slightly lower income growth at the time of measurement. This trend is very steady over child ages 8 to 16, with income growth varying by about 0.2% over this 8 year range. Then, immediately when parent income begins to affect the child's college aid at age 17, income falls by 0.7% relative to trend. At age 18, when the child begins college and college costs first come due, parent income increases by 0.9% relative to trend. Income growth reverts back near trend for the following three years. Then, when the child is age 22 and would typically be finishing their college studies, parent income sharply falls again by nearly the same amount as it increased at age 18. This pattern is consistent with a simple model in which parent income responds to both types of incentives outlined in Figure 3: parents reduce income due to the substitution incentive created by the college aid implicit tax, which is isolated in the age 17 effect; and parents increase income precisely during the years when their child incurs college costs, reflected in the spike at age 18 and the dip at age 22.

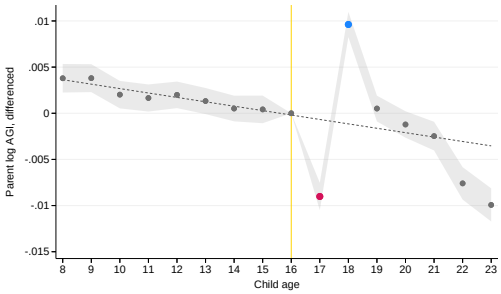
Figure 4: Parent Income Responses to College Incentives



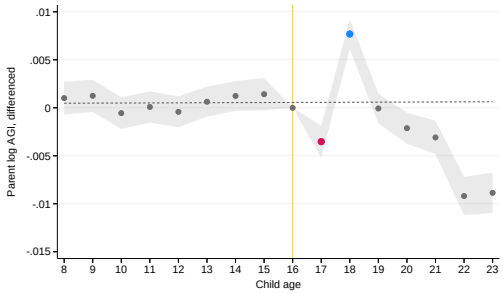
(A) 4-year college



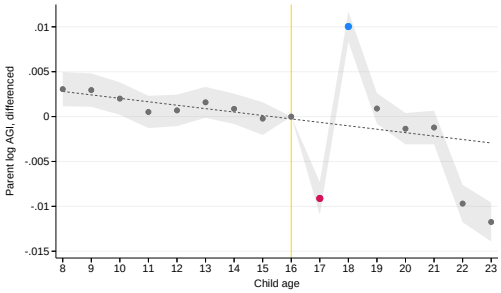
(B) No college



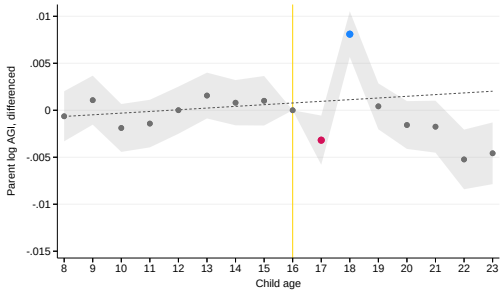
(C) High aid phaseout college



(D) Low aid phaseout college



(E) High cost college



(F) Low cost college

Notes: This figure shows estimates of child age effects from the nonparametric event study regression in Equation (1), where the outcome is first-differenced log AGI. Because our variation comes from parents who are different ages when they had their child, trends arise if families with different income trajectories have children at different ages. We overlay a linear trend fitted to the age 8 through 16 effects. Differences from this trend reflect our parametric event study estimates that control for the trend. All panels restrict to households that have baseline income averaged over 2005 and 2006 between \$50k to \$150k, filed taxes with positive AGI for the years 2001 through 2022, claimed the relevant child as a dependent when the child was age 7 through 17, and whose child turned 18 between 2012 and 2016. Each panel shows estimates for a separate subsample based on the child's college attendance at age 18: Panel (A) is a four-year in-state public college or four-year private college that has IPEDS data on average aid by family income bin; Panel (B) is no college; Panel (C) is for children who receive at least one 1098-T at age 18 from a college with an aid phaseout rate above 10%, defined by the change in average aid between families in income bins \$30-48k and \$75-110k divided by the difference in the midpoints of each range; Panel (D) is for children who receive a 1098-T at age 18, but any 1098-T they receive is from a college with an aid phaseout rate below 10%; Panel (E) is for children who receive at least one 1098-T at age 18 from a college with a total Cost of Attendance above \$30k; Panel (F) is for children who receive a 1098-T at age 18, but any 1098-T they receive is from a college with total Cost of Attendance below \$20k. Standard errors are clustered at the tax unit level.

Without further evidence, the sharp drop of income at age 17 and the sharp increase at age 18 presented in Panel A could reflect other factors that vary with a child's age. To address this, Panels B through F show estimates for samples that vary in their exposure to college incentives. Effects concentrate among groups facing greater incentives, indicating that the age 17 and age 18 effects in Panel A are not due to other common age-related phenomena.

Panel B shows placebo estimates for families whose child does not attend college at age 18, as measured by the child not receiving a 1098-T. We observe a slight and nearly negligible upward pretrend, reflecting that for this population, parents who had their child at a younger age may have slightly higher wage growth at the time of measurement. Contrasted with the sharp break in trend at age 17 observed in Panel A, Panel B does not show a similar sharp trend break but instead shows a very minor dip relative to the extrapolated linear trend. The dip relative to the linear trend in Panel B is only 0.1%, compared with the 0.7% drop in Panel A. Two factors are worth observing regarding the fact that there is any deviation from trend at age 17 in Panel B. First, as the child's age 17 year spans from the spring of their third year in high school through the fall of their fourth year in high school, some of the families in Panel B may hold beliefs with nonzero probability that their child will attend college, causing some income reduction. Second, for some of the families in Panel B, our data might classify the child as not attending college despite the child actually enrolling. This is because our measure of college attendance depends on observing a 1098-T, and a number of colleges do not file a 1098-T for every student (Cronin and Gray-Hancuch 2024). Nevertheless, the clearest evidence regarding the age 17 effect loading onto college families comes from comparing Panel B against the sharp trend break in Panel A. At age 18, we observe a stark drop of income for non-college families, which suggests similarly that the sharp increase of income at age 18 in Panel A would not have occurred in the absence of college.

Panels C and D show estimates separately for families whose child attends a college with a high rate of aid phaseout (10%+) or a low rate of aid phaseout (<10%), respectively. Relative to trend, parent income falls by twice as much at age 17 for the high phaseout group (0.8%) compared with the low phaseout group (0.4%). The increase of income at age 18 relative to trend is also greater for the high phaseout group (1.1%) than for the low phaseout group (0.7%), consistent with the fact that high phaseout colleges also tend to be high cost colleges, as a higher baseline cost means more baseline aid to phase out. Panels E and F show estimates for families whose child attends a college with a high cost of attendance ($30k+) or a low cost (<$20k). The patterns at ages 17 and 18 are nearly identical between Panels E and C and between Panels F and D, again reflecting that high cost colleges also tend to be high aid phaseout colleges. One distinguishing feature is the income reduction at age 22. While this reduction is similar between the high and low aid phaseout groups, the reduction is much more stark for the high cost than for the low cost group, consistent with families maintaining higher income precisely while they face high college costs. The fact that the age 17 income reduction is so similar between high phaseout and high cost colleges and between low phaseout and low cost colleges also suggests one way that families might be able to differentiate high from low phaseout colleges even if they do not know the exact aid schedules: they can simply reference cost, as high (low) cost colleges tend to be high (low) phaseout colleges.

Figure A.5 shows the same results using levels of log AGI as the outcome instead

of first differences. One benefit of the levels specification is that it is straightforward to see the cumulative effect at any point. Here, consistent with evidence in Kleven et al. (2019) and as discussed above in Section 3.2, we clearly see trends reflecting that parents who were younger when they had their child also tend to have lower income. Because we observe slight linear trends in our first differences specifications, we overlay quadratic trends in the levels specifications shown in Figure A.5. Relative to these trends, we observe the same effect patterns as Figure 4 shows in first differences: college families reduce income at age 17 and increase income at age 18, whereas non-college families see no effect at age 17 and begin reducing income at age 18; responses are larger among college families whose child attends a college with high aid phaseout or high cost versus families whose child attends a college with low aid phaseout or low cost.

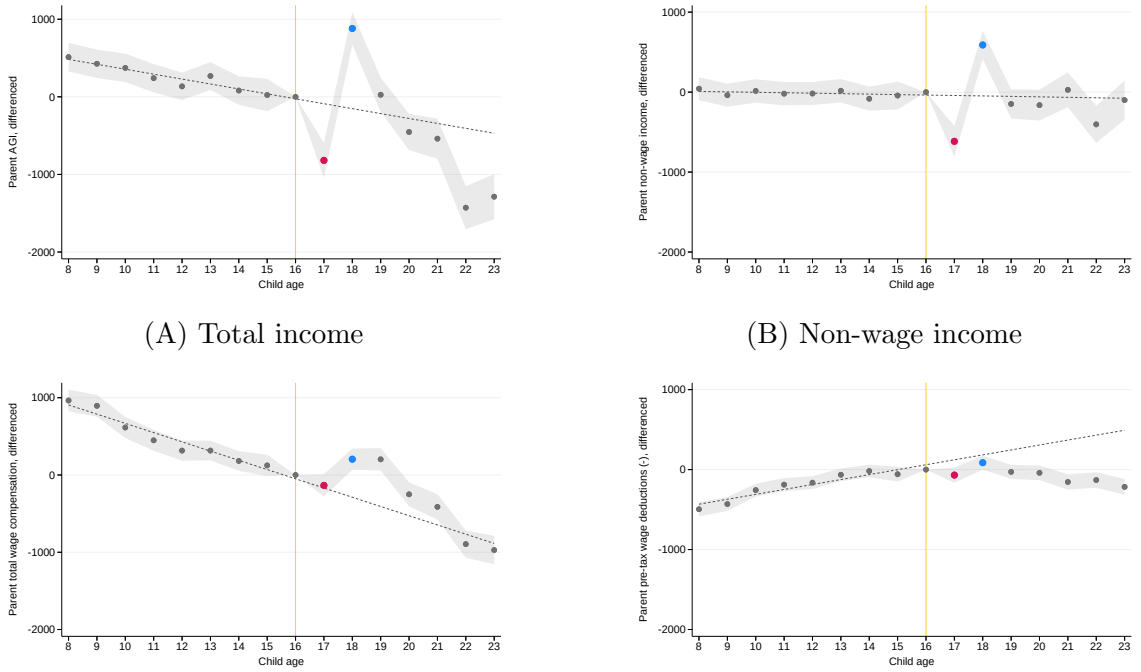
Figure A.6 shows the same set of results using second differences of log AGI as the outcome. We only show estimates through age 17 in Figure A.6, because effects at ages 18 and older now mix college effects across two adjacent years. These results show flat pretrends and the same effect pattern occurring at age 17 as in our parametric first differences specification, reinforcing that our main estimates are not an artifact of trend specification.

3.3.2 Mechanisms: How Parents Adjust Income

How do parents adjust income in response to college incentives? To address this question, we leverage the richness of the tax data and estimate our nonparametric event study for several different types of income in the 4-year college sample. Figure 5 shows the results. Panel A shows estimates with first differences of AGI as the outcome. Relative to trend, AGI falls by \$739 at age 17 and increases by \$1,027 at age 18. Panels B through D decompose the effects on AGI. Panel B shows non-wage income, defined as AGI minus total W-2 income. Non-wage income accounts for nearly all of the age 17 income reduction (\$580), and it also accounts for a majority of the age 18 income increase (\$628). Panel C shows total wage compensation, defined as W-2 income that is subject to the Medicare tax (Box 5). This is our cleanest measure of a pure labor supply response, as it includes wage income prior to any pre-tax deductions resulting from contributions to an employer-sponsored retirement plan. The downward trend in Panel C reveals that the trends we observe in Figure 4 – that income growth is slower among parents who are younger at child birth – are driven by slower growth in wage income among parents who are younger at child birth. Total wage compensation does not respond at age 17, after which it increases by \$456 at age 18, effectively accounting for the remainder of the age 18 AGI effect. Finally, Panel D shows the impact on AGI from pre-tax wage deductions. This is defined as total W-2 income (Box 1) minus Medicare wages (Box 5), and it primarily reflects the amount of total wage compensation that is excluded from AGI due to pre-tax contributions to an employer-sponsored retirement plan such as a 401(k), 403(b), or 457(b). Because Panels B through D are defined to add up to AGI, our measure in Panel D is the negative of pre-tax wage deductions: when this measure is lower, it reflects more pre-tax retirement contributions that in turn lower AGI. We see evidence of small adjustments at ages 17 and 18, reflecting that parents marginally increase retirement contributions at age 17 to reduce AGI and then reduce these contributions at age 18. Because the trend for pre-tax deductions appears quadratic,

we view this result as suggestive. In sum, nearly all of the age 17 income reduction is due to non-wage income, while the age 18 income increase comes roughly 60% from non-wage sources and 40% from wage income.

Figure 5: Which Types of Income Parents Adjust



(A) Total income

(B) Non-wage income

(C) Total wage compensation

(D) Pre-tax wage deductions (-)

Notes: This figure shows estimates of child age effects from the nonparametric event study regression in Equation (1). The outcome in Panel (A) is first-differenced AGI, in Panel (B) is the first difference of AGI minus W-2 Box 1, in Panel (C) is the first difference of W-2 Box 5, and in Panel (D) is the first difference of W-2 Box 1 minus Box 5. Because our variation comes from parents who are different ages when they had their child, trends arise if families with different income trajectories have children at different ages. We overlay a linear trend fitted to the age 8 through 16 effects. Differences from this trend reflect our parametric event study estimates that control for the trend. All panels restrict to households that have baseline income averaged over 2005 and 2006 between \$50k to \$150k, filed taxes with positive AGI for the years 2001 through 2022, claimed the relevant child as a dependent when the child was age 7 through 17, whose child turned 18 between 2012 and 2016, and whose child attended a four-year in-state public college or four-year private college that has IPEDS data on average aid by family income bin at age 18. Standard errors are clustered at the tax unit level.

Figure A.7 further breaks down the responses of non-wage income into subcomponents. Business income accounts for substantial portions of the age 17 reduction and the age 18 increase of non-wage income. Together, business income, taxable IRA distributions, and capital gains account for just over half of the age 17 reduction of non-wage income and two-thirds of the age 18 increase of non-wage income. Interestingly, there is evidence of short-term anticipatory capital gains shifting, reflected in the increase of capital gains at age 16. However, the age 17 reduction of income does not appear to be driven by short-term anticipatory income shifting overall, as we do not see a strong increase of total income or of total non-wage income at age 16.

Motivated by the high sensitivity of non-wage income at age 17 and the respon-

siveness of both wage and non-wage income at age 18, Figure A.8 shows responses of first differenced log AGI separately for families with a high or a low share of non-wage income. The age 17 effect is twice as large for high-non-wage families (1.3%) as for low-non-wage families (0.6%). Consistent with the increase of business income and other non-wage factors at age 18, the age 18 effect is also twice as large for high-non-wage families (1.7%) as for low-non-wage families (0.8%). Informed by these results, we separate our estimates of the age 17 elasticity of income in the following section by high and low non-wage families.

3.3.3 The Elasticity of Parent Income to the College Aid Implicit Tax

Finally, we turn to computing the key reduced form parameter of interest: the elasticity of parent income with respect to the college aid implicit tax. Table 1 shows our results. Each column of Table 1 reflects a different sample. The first row shows estimates of the reduced form effect on parent log AGI, which is the age 17 effect in our parametric regression specification with first differenced log AGI as the outcome. The second row shows the first stage effect on the log net-of-tax rate, which is the age 17 effect in our parametric regression specification with first differenced log net-of-tax rate as the outcome. The net-of-tax rate is one minus the tax rate from college aid, reflecting the share of each marginal dollar of income that the family "keeps" after accounting for the corresponding reduction of aid. The third row shows the ETI, which is the ratio of the reduced form to the first stage.[16] This approach isolates the response of parent income to variation in the log net-of-tax rate that is induced by the introduction of the college aid implicit tax when the child is age 17, conditional on parent age and year fixed effects and relative to a linear trend in child age.

The first two columns of Table 1 show estimates for the Broad aid and 4-year aid samples. In both, the log net-of-tax rate falls by about .15, with a slightly greater reduction in the 4-year aid sample. In response, income falls by 0.7% in both samples, with a slightly larger response in the 4-year aid sample. This produces nearly identical elasticities across the two samples at .04. The third and fourth columns of Table 1 show estimates among the 4-year aid sample split into families with a low or high baseline share of non-wage income. Families with a high share of non-wage income are twice as elastic (.08) as families with a low share of non-wage income (.04). Table A.2 shows results based on a version of our parametric regression specification that includes a quadratic trend in child age. The results in Table A.2 are essentially unchanged relative to the results in Table 1.

3.3.4 Robustness and Extensions

Selection does not explain the age 17 effect.—As noted above, the evidence shown in Figure 4 indicates that the age 17 effect among college families is not due to other common age-related phenomena, as evidenced by a lack of such effect among non-college families and the scaling of effects for families whose child attends a college with higher aid phaseout or higher cost. Another type of identification concern is that the choice to enroll in college might be endogenous to preceding household income shocks, creating selection into treatment. We can directly assess in the tax data

[16] To get standard errors, we estimate the ETI by 2SLS using the child age 17 indicator as an instrument for the first differenced log net-of-tax rate. The first stage F-statistic is well above 10,000 in all reported specifications, so we omit reporting this statistic in our results.

Table 1: ETI due to the College Aid Implicit Tax

	By college type		By non-wage income, among 4-year aid sample	
	Broad aid sample	4-year aid sample	Low non-wage	High non-wage
Reduced form	-0.0065	-0.0069	-0.0062	-0.0127
	(0.0005)	(0.0006)	(0.0005)	(0.0021)
First stage	-0.1477	-0.1569	-0.1565	-0.1597
	(0.0003)	(0.0003)	(0.0003)	(0.0008)
ETI (2SLS)	0.0442	0.0439	0.0396	0.0794
	(0.0034)	(0.0035)	(0.0035)	(0.0133)
N	11,119,582	9,477,378	8,165,556	1,311,822

Notes: This table shows estimates of age 17 effects from the parametric event study regression in Equation (2). The first row shows the reduced form effect on the outcome of first-differenced log AGI. The second row shows the first stage effect on the first difference of the log net of tax rate, which is one minus the marginal tax rate from college aid phaseout. The third row shows the ETI, which is the ratio of the reduced form to the first stage. All panels restrict to households that have baseline income averaged over 2005 and 2006 between \$50k to \$150k, filed taxes with positive AGI for the years 2001 through 2022, claimed the relevant child as a dependent when the child was age 7 through 17, and whose child turned 18 between 2012 and 2016. The first column also restricts to households whose child attends a "Broad aid" college at age 18, and the second column restricts to households whose child attends a "4-year aid" college at age 18. Among households whose child attends a "4-year aid" college at age 18, the third column restricts to households whose W-2 income accounts for more than 80% of their AGI on average over 2005 and 2006, and the fourth column restricts to households whose W-2 income accounts for less than 80% of their AGI on average over 2005 and 2006. Standard errors are clustered at the tax unit level.

whether this type of selection is important. Suppose the age 17 effect were entirely due to selection – if a family experiences a negative income shock when the child is 17, the child becomes more likely to attend college. If this were the case, we would expect two patterns in the data. First, since the age 17 "effect" is only due to selected sorting into college, there would be no effect on income in the aggregate sample that combines all families irrespective of their child's college attendance. Figure A.9 reports estimates from this combined sample of all families with a child turning 18 between 2012 and 2016, regardless of college attendance. Figure A.9 shows a clear reduction of income at age 17 in the aggregate sample, which is inconsistent with the selection story. Second, if selection into college were important, then we would observe an upward spike in income for the non-college group: as families that face a negative income shock at age 17 sort into the college group, they leave the non-college group, leaving average income higher at age 17 among those who remain in the non-college group. Figure 4 Panel B shows no upward spike of income at age 17 in the non-college group, again pointing toward selection being unimportant. The age 17 effect we observe in the 4-year aid and Broad aid samples also does not appear to be shaped by selection into specific college types, as Figure A.10 shows similar effects among families whose child receives a 1098-T from any college at age 18. That selection into college does not explain the age 17 effect is consistent with prior work studying the effect of shocks to parent income – through parent layoffs (Hilger 2016) and lottery wins (Bulman et al. 2021) – when a child is in high school on the child's college enrollment. These papers find that the effects of such shocks are small, and moreover the direction of these effects is inconsistent with selection explaining our

effects: a positive income shock marginally increases college enrollment, which cannot explain our finding of income reduction among college families.

The age 18 income effect.—While the primary focus of this paper is on the substitution response to the college aid implicit tax at age 17, we briefly offer interpretation of the observed increase of income at age 18 among college families, as seen in Figure 4. This increase is not due to common age-related phenomena, evidenced by the reduction of income at age 18 among non-college families. While the increase among college families and decrease among non-college families could reflect selection into college based on anticipated income shocks, a few pieces of evidence suggest that such selection does not fully explain the observed age 18 effect. First, Figure A.9 shows a slight increase of income at age 18 on net in the full sample that does not condition on college status, meaning at least some of the increase among college families is not due to selection. Second, Figure 4 shows a drop of income at age 22 that is more pronounced among families that face greater college costs. This is consistent with parents increasing income precisely while college costs are due, but is not rationalizable through children selecting into college at age 18 based on anticipated negative income shocks at age 22. Finally, while non-wage sources account for roughly 60% of the age 18 increase for college families, there is no change in non-wage income at age 18 for non-college families, shown in Figure A.11. This suggests that none of the non-wage income increase among college families results from selection. To the extent that the age 18 income increase for college families is caused by college itself, one might argue this results from parents no longer having a child at home, rather than from college costs. Figure A.12 rejects this explanation, as college families with a younger sibling (defined as age 1-14 when the focal child is 18, so there are no immediate competing effects from the younger child's college enrollment) in the household actually increase income by more than families without a younger sibling.

This all points toward a real increase of income among college families that is tied to college costs. Such an effect is difficult to rationalize in a simple lifecycle model with perfect foresight, complete credit markets, and standard optimization, as such a model would predict income smoothing at ages 18 and 22. Possible explanations for the observed income effect include imperfect foresight, as parents might face ex ante uncertainty over whether and where their child will attend college or they might systematically underestimate costs; incomplete credit markets and binding liquidity constraints; or nonstandard optimization such as mental accounting (Hastings and Shapiro 2018; Chorniy et al. 2025).

Removing filing and claiming restrictions.—Our preceding analysis restricts the sample to families that file taxes with positive AGI in all years from 2001 to 2022 and that claim the focal child from ages 7 to 17. As discussed in Section 3.1, this sample restriction is to avoid confounding that could otherwise arise due to mean reversion patterns and to ensure we do not mis-measure income for tax units that dissolve and then file as new tax units. Nevertheless, the main effects we document above in our main analysis sample can also be seen clearly when we remove the filing and claiming restrictions. Figure A.13 shows estimates for our 4-year aid college sample, but with the persistent filing and claiming restrictions removed. The only such restriction is that the family filed taxes and claimed the focal child when the child was age 14, which is necessary to link parents and children. Panel A shows results where the

outcome variable is an indicator for filing taxes with positive AGI. The coefficient is elevated at age 14 with some decrease at younger and older ages, reflecting mean reversion. The decrease is greater at older ages, likely reflecting tax units that dissolve over time or stop filing taxes for other reasons. The important takeaway from Panel A is that the pattern is smooth through ages 17 and 18, suggesting that there is no sizable extensive margin effect from college-related incentives on filing with positive AGI that would contaminate our intensive margin estimates. Consistent with the lack of extensive margin effect, Panel B shows that our main effects remain apparent in the unrestricted sample. Panel B shows estimates in which the outcome is first differenced AGI. We record AGI as zero in the case of non-filing. There is a large drop at age 15, the first age past the age-14 restriction after which point tax units can dissolve or stop filing. Beginning from age 15, the coefficients at ages 15 and 16 are not statistically different from one another, consistent with a roughly constant rate of tax unit dissolution beginning from age 15 onward. However, the age 17 coefficient is -$655 and the age 18 coefficient is $1,062. These coefficients are similar to the effects on AGI in our main restricted sample shown in Figure 5 Panel A. Altogether, Figure A.13 strongly suggests that the extensive margin of filing with positive AGI does not meaningfully respond to college incentives at ages 17 and 18, which is an interesting result independently and implies that our main estimates on the intensive margin are not driven by changing composition among persistent filing and claiming families.

Estimates for low- and high-income families.—Our main estimates focus on families with income in the $50k to $150k range because this is where most aid phaseout occurs. Figure A.14 shows estimates for low-income families who are unlikely to be exposed to significant aid phaseout. We define low-income families as those whose average income over years 2009 and 2010 is below $25k. We use income from 2009 and 2010 so these families are still likely to have low income when their child begins college. Because college aid also depends on assets, we further restrict this group to families that have no interest income in 2009 and 2010, reflecting that they have no interest-bearing assets. Consistent with this group facing little aid phaseout, we see no clear evidence of income reduction at age 17, as Figure A.14 Panels A (college families) and B (non-college families) show nearly identical patterns of income through age 17. Figure A.15 shows estimates for high-income families. We define high-income families as those whose average income over 2009 and 2010 is between $150k and $300k, where we similarly use income from 2009 and 2010 so these families likely still have high income when their child begins college. Panel A shows that high-income college families reduce income by 0.6% when their child is 17, whereas Panel B shows that high-income non-college families see no deviation in income at age 17. High-income families might reduce income in part because their child attends a college with a high cost of attendance, meaning the aid phaseout range is much wider – for example, Figure 1 shows that families with income upward of $300k still face aid phaseout if their child attends the University of Pennsylvania. Consistent with this, we see that the high-income response is largely driven by families whose child attends a very high cost of attendance college. Panel C shows that income falls by 1.1% relative to trend for high-income families whose child attends a college with total cost above $50k, while Panel D shows that income only falls by 0.3% for high-income families whose child attends a college with total cost below $50k.[17]

[17] High-income families may still face some aid phaseout at these colleges, and some high-income

4 Measuring Welfare under Tax Rate Misperception

4.1 Theory

In public economics, the standard motivation for estimating the ETI with respect to a particular tax is that the ETI identifies the DWL produced by the tax (Feldstein 1999). However, given the difficulties involved in learning one's exact tax rate from college aid phaseout as documented in Section 2, does the ETI estimated in Section 3 still identify DWL from the college aid implicit tax?

To answer this question, we begin by deriving the following sufficient statistics expression for DWL under tax rate misperception:

$$DWL \approx \underbrace{N \cdot \frac{1}{2} \cdot \Bigg\{ E\left[Z_{i,1} e(X_i) \tau_i^2\right] \cdot \Big(E[\theta_i]^2 + Var(\theta_i) \Big)}_{\substack{\text{Correct Perceptions} \\ \text{(standard Harberger triangle)}}} + E\left[e(X_i)\right] E\left[\tau_i^2\right] \cdot Cov\left(\theta_i^2, Z_{i,1}\right)$$
$$+ E\left[Z_{i,1}\right] E\left[\tau_i^2\right] \cdot Cov\left(\theta_i^2, e(X_i)\right) + E\left[Z_{i,1}\right] E\left[e(X_i)\right] \cdot Cov\left(\theta_i^2, \tau_i^2\right) \Bigg\} \tag{3}$$

The full model setup and derivation of Equation (3) can be found in Appendix B.1.[18] In the model, individuals face a tax rate τ_i but perceive a tax rate $\tilde{\tau}_i$. We define the individual-level misperception measure $\theta_i \equiv \frac{\tilde{\tau}_i}{\tau_i}$. $Z_{i,1}$ refers to counterfactual income in the taxed period if there were no tax.

How does DWL relate to the ETI under misperception? The first observation from Equation (3) is that the ETI appears only implicitly. In Equation (3), $e(X_i)$ is a parameter of the utility function that represents the structural elasticity with respect to the perceived tax $\tilde{\tau}_i$, which we allow to depend on covariates X_i. Under misperception, this structural elasticity may not be equal to the reduced form elasticity economists typically estimate; however, as we prove in Appendix B.2, we can identify this structural elasticity as a function of the standard ETI and information on actual and perceived tax rates:

$$e(X_i) = \beta^{ETI}(X_i) \cdot \frac{E[\ln(1-\tau_i) \,|\, X_i]}{E[\ln(1-\tilde{\tau}_i) \,|\, X_i]} \tag{4}$$

where $\beta^{ETI}(X_i)$ is the standard ETI, which can depend on covariates X_i. The intuition behind Equation (4) is as follows. Suppose individuals uniformly underestimate their tax rate, $\tilde{\tau}_i < \tau_i$. Then, individuals will respond less to the tax than the utility function parameters would dictate if they had perceived the tax correctly. In turn, the standard ETI observed in the data – which captures the response of income to the true tax rate – will understate the structural elasticity parameter. Equation (4) shows that, if one knows the average perception of the log net-of-tax rate, it is

families may reduce income enough to go from getting no aid to getting some aid.

[18] The expression is an approximation due to the standard use of a second order Taylor expansion.

possible to convert between the ETI and the structural elasticity. Appendix B.2 contains further discussion of this point with a graphical representation. Nevertheless, it is the structural elasticity that appears in Equation (3) because it governs utility; the standard ETI only appears implicitly through its relationship to the structural elasticity shown in Equation (4).

As a baseline case, when perceptions are accurate (i.e. $\tilde{\tau}_i = \tau_i$ for all i), then we have that $\theta_i = 1$ and $\beta^{ETI}(X_i) = e(X_i)$, and Equation (3) reduces to the standard Harberger triangle. In other words, Equation (3) nests the result from Feldstein (1999) that under correct perceptions the ETI identifies DWL.

However, when individuals misperceive their tax rate, Equation (3) shows that the ETI is no longer sufficient to identify DWL. Instead, under misperception, the sufficient statistics for welfare include the ETI and several features of the joint distribution of misperceptions, income, elasticity, and tax rate. Specifically, DWL increases with the average misperception, the variance of misperceptions across individuals, and the covariances of squared misperceptions with income, elasticity, and squared tax rate.

It is noteworthy that DWL depends on moments of the distribution of misperceptions beyond the average misperception. DWL increases with the average misperception because this increases the observed ETI, conditional on the structural elasticity. However, conditional on the average misperception, DWL also increases with second moments. DWL increases with the variance of misperceptions due to the standard assumption of increasing marginal disutility of producing income, which means that larger reductions of income in response to a tax bear disproportionately greater impacts on DWL than do smaller reductions. As a consequence, in utility terms, it is worse to have some individuals reduce income by a large amount and others reduce income by a small amount than for all individuals to reduce income by the same amount, even if the average reduction is equal in both cases.[19] This result is directly related to the standard public finance result that the DWL induced by a tax is proportional to the square of the tax rate, and it is consistent with a few recent results that variance of behavioral frictions can have first order effects on welfare (Taubinsky and Rees-Jones 2018; Farhi and Gabaix 2020; Allcott et al. 2025). DWL also increases with the covariances of squared misperceptions – which accounts for both the average and variance channels discussed above – with income, elasticity, and squared tax rate, because a greater value of any one of these factors produces greater baseline DWL that is then scaled by misperceptions.

In the context of this paper, the key observation from the theory is the following: If at least some families misperceive the exact tax rate they face from college aid phaseout, which seems likely from Section 2, then the ETI estimated in Section 3 does not directly identify DWL from the college aid implicit tax. However, it may still be possible to identify DWL by combining the ETI with information on the distribution of tax rate misperceptions, which we turn to next.

4.2 Identifying Misperception Moments using Survey Data

To implement Equation (3) and measure DWL under tax rate misperception, one must combine a well-identified ETI estimate with well-identified measures of several features of the distribution of misperceptions. To estimate distributional features of

[19] Supplementary Appendix C explains this with an example.

misperceptions, one might naturally consider surveying individuals on their perception of tax rates, as has been used in previous work (Fujii and Hawley 1988; Slemrod 2006; Rees-Jones and Taubinsky 2020). However, a challenge arises: asking a survey respondent for their perceived tax rate may produce a noisy measure of their underlying perception. If survey measures of tax rate perceptions are noisy but the noise is mean zero, then we can easily identify average perceptions. But, even in this case, a noisy measure of perceptions will bias measures of second moments, which appear in Equation (3) and have been shown to be first order in other contexts including tax salience or the effects of nudges (Taubinsky and Rees-Jones 2018; Allcott et al. 2025). In this section, we show how to identify the moments that involve misperceptions in Equation (3) using survey data even if responses are noisy measures of underlying perceptions, which allows for empirical application of Equation (3) and which is the key methodological contribution of this paper.[20] Intuitively, our solution is to also measure the respondent's perception of a slightly different tax rate, and to use this measure as an instrument to purge measurement error from second moments while correcting for any influence due to the secondary rate being slightly different from the focal rate.

In the survey, respondents provide a noisy signal of their perceived tax rate, $\hat{\tilde{\tau}}_{1i} = \tilde{\tau}_i + \epsilon_{1i}$. $\tilde{\tau}_i$ is the respondent's underlying perception of the tax rate, and the error term ϵ_{1i} reflects measurement error. We will assume that the error term ϵ_{1i} has conditional mean zero and is conditionally uncorrelated with the perceived tax rate, with income, and with elasticity: $E[\epsilon_{1i}|X_i, \tilde{\tau}_i] = E[\epsilon_{1i}|Z_{i,1}, \tau_i] = E[\epsilon_{1i}|e(X_i), \tau_i] = 0$ and $Cov(\tilde{\tau}_i, \epsilon_{1i}|Z_{i,1}, \tau_i) = Cov(\tilde{\tau}_i, \epsilon_{1i}|e(X_i), \tau_i) = Cov(\epsilon_{1i}, Z_{i,1}|\tau_i) = Cov(\epsilon_{1i}, e(X_i)|\tau_i) = 0$. Note these assumptions imply also that $E[\epsilon_{1i}|\tau_i] = 0$ and $Cov(\tilde{\tau}_i, \epsilon_{1i}|\tau_i) = 0$. A simpler way to conceptualize these assumptions is through the more restrictive condition that the measurement error ϵ_{1i} is jointly independent of one's perceived tax rate, income, covariates X_i and corresponding elasticity, and true tax rate. This joint independence would generate our set of required conditions, though our conditions are less restrictive than the full joint independence condition. Likewise, if the measurement error were pure random noise with mean zero, this would also satisfy our conditions, though again our conditions are less restrictive. The assumptions on ϵ_{1i} reflect that it is simply measurement error and thus unlikely to be related to various other factors.

Crucially, respondents also provide a second signal that is related to their perceived tax rate, which can be thought of as an instrument. An important feature for our implementation is that the elicitation for the instrument will ask the respondent to provide their perception of a tax rate that is slightly different from the tax τ_i, which we denote as $\tau_i + a_i$. Then we denote the secondary measure from the survey, the respondent's reported perception of the tax rate $\tau_i + a_i$, as $\hat{\tilde{\tau}}_{2i} = \tilde{\tau}_i + \tilde{a}_i + \xi_i = \tilde{\tau}_i + \epsilon_{2i}$. Here, $\tilde{a}_i$ is the respondent's actual perception of a_i, the difference between the secondary tax rate $\tau_i + a_i$ and the focal tax rate τ_i. ξ_i is measurement error. Then we define $\epsilon_{2i} = \tilde{a}_i + \xi_i$, which clarifies that the difference between the secondary survey measure $\hat{\tilde{\tau}}_{2i}$ and the respondent's underlying perception of the focal tax $\tilde{\tau}_i$ combines two factors: the difference in actual perceptions of the two tax rates $\tilde{a}_i$ and measurement error ξ_i.

[20] Indeed, Taubinsky and Rees-Jones (2018) can only provide a lower bound on the welfare impact of heterogeneous tax salience, and follow up work in Brusco and Glass (2020) provides theory demonstrating the challenges in fully identifying DWL with heterogeneous salience.

We make the following assumptions regarding the error term ϵ_{2i}:

Error exogeneity: $Cov(\tilde{\tau}_i, \epsilon_{2i}|Z_i, \tau_i) = Cov(\tilde{\tau}_i, \epsilon_{2i}|e(X_i), \tau_i) = 0$.

Error independence: $Cov(\epsilon_{1i}, \epsilon_{2i}|Z_{i,1}, \tau_i) = Cov(\epsilon_{1i}, \epsilon_{2i}|e(X_i), \tau_i) = 0$.

Together, the exogeneity and independence conditions ensure that the only reason $\hat{\tilde{\tau}}_{2i}$ and $\hat{\tilde{\tau}}_{1i}$ are systematically related to one another is due to variation in $\tilde{\tau}_i$. Exogeneity says that, conditional on the true tax rate and covariates, any relationship between $\hat{\tilde{\tau}}_{2i}$ and $\hat{\tilde{\tau}}_{1i}$ is not influenced by correlation between the error term ϵ_{2i} and the underlying perception $\tilde{\tau}_i$. An example of a violation of exogeneity would be if, conditional on the true tax rate, respondents who perceive a higher tax rate also systematically underreport their perception while respondents who perceive a lower tax rate report their perception accurately on average, which seems unlikely. Because the error term ϵ_{2i} includes both measurement error ξ_i and the respondent's perception of the gap between the focal and the instrument tax rates, $\tilde{a}_i$, exogeneity also means that the respondent's perception of the gap $\tilde{a}_i$ is not systematically related to their perception of the focal tax rate $\tilde{\tau}_i$, conditional on the true tax rate τ_i. As an example, suppose the true secondary rate is a constant reference rate of 15% that does not vary across respondents. Among respondents whose true focal rate τ_i is 25%, exogeneity implies that those who overestimate their focal rate also tend to perceive a higher secondary rate compared to those who underestimate their focal tax rate. This effectively implies a relevance condition as well: conditional on the true tax rate and covariates, those who perceive a higher secondary rate also perceive a higher focal rate.

Independence says that, conditional on the true tax rate and covariates, any relationship between $\hat{\tilde{\tau}}_{2i}$ and $\hat{\tilde{\tau}}_{1i}$ is not influenced by correlation between the error term ϵ_{2i} and the error term ϵ_{1i}. An example of a violation of independence would be if, conditional on the true tax rate, respondents who overreport (underreport) their perception of the focal rate systematically overreport (underreport) their perception of the secondary rate. Such a pattern could be a concern, for example if the focal and the secondary rates referenced the same underlying tax rate and some respondents simply remember the estimate they provided in the first measure and repeat that number for the second measure. Our empirical implementation takes two important steps to alleviate such a concern. (1) The underlying tax rate we use for our secondary measure is slightly different from the underlying focal tax rate. This means we are not simply asking the respondent to repeat the same perception twice, rather we are asking for new information in our secondary measure. (2) Instead of asking the respondent to directly provide their perceptions of tax *rates*, we ask for their best guess of *levels* at various points of income, which are asked on different screens in the survey, and then we construct marginal tax rates based on the slope of these points. The income points we ask about also are different between the focal and the secondary rate questions. This makes it even more difficult for respondents to simply repeat the same tax rate between the two measures. Instead, by asking the respondent for their perception of tax levels at various points of income which differ between the focal and secondary elicitations, our approach makes it less likely that measurement error in the slope over these points would be systematically related between the two measures. A further benefit to asking the respondent for their perception of tax

levels at various income points rather than directly asking for their perception of the marginal tax rate is that tax level is a simpler concept to understand, so our method is not affected if respondents misunderstand the concept of a marginal tax rate.

In addition to error exogeneity and independence, we also assume that ϵ_{2i} is conditionally mean-independent of income and elasticity: $E[\epsilon_{2i}|Z_{i,1}, \tau_i] = E[\epsilon_{2i}|e(X_i), \tau_i] = E[\epsilon_{2i}|\tau_i]$. These assumptions simplify estimation. We note as well that, as written above, error exogeneity and independence also imply exogeneity and independence conditional on τ_i alone: $Cov(\tilde{\tau}_i, \epsilon_{2i}|\tau_i) = Cov(\epsilon_{1i}, \epsilon_{2i}|\tau_i) = 0$, which we use in our proofs.

Under these assumptions, and defining the survey-based measures of misperception $\hat{\theta}_{1i} = \frac{\hat{\tilde{\tau}}_{1i}}{\tau_i}$ and $\hat{\theta}_{2i} = \frac{\hat{\tilde{\tau}}_{2i}}{\tau_i}$, we have the following identification results:

1. $E[\ln(1-\tilde{\tau}_i)|X_i] \approx E[\ln(1-\hat{\tilde{\tau}}_{1i})|X_i]$
2. $E[\theta_i] = E[\hat{\theta}_{1i}]$
3. $Var(\theta_i) = Cov(\hat{\theta}_{1i}, \hat{\theta}_{2i}) - Cov\left(\frac{E[\hat{\tilde{\tau}}_{1i}|\tau_i]}{\tau_i}, \frac{E[\hat{\tilde{\tau}}_{2i}-\hat{\tilde{\tau}}_{1i}|\tau_i]}{\tau_i}\right)$
4. $Cov(\theta_i^2, Z_{i,1}) = Cov(\hat{\theta}_{1i}\hat{\theta}_{2i}, Z_{i,1}) - Cov\left(\frac{E[\hat{\tilde{\tau}}_{2i}-\hat{\tilde{\tau}}_{1i}|\tau_i]}{\tau_i^2}\hat{\tilde{\tau}}_{1i}, Z_{i,1}\right)$
5. $Cov(\theta_i^2, e(X_i)) = Cov(\hat{\theta}_{1i}\hat{\theta}_{2i}, e(X_i)) - Cov\left(\frac{E[\hat{\tilde{\tau}}_{2i}-\hat{\tilde{\tau}}_{1i}|\tau_i]}{\tau_i^2}\hat{\tilde{\tau}}_{1i}, e(X_i)\right)$
6. $Cov(\theta_i^2, \tau_i^2) = Cov(\hat{\theta}_{1i}\hat{\theta}_{2i}, \tau_i^2) - Cov\left(\frac{E[\hat{\tilde{\tau}}_{2i}-\hat{\tilde{\tau}}_{1i}|\tau_i]}{\tau_i^2}\hat{\tilde{\tau}}_{1i}, \tau_i^2\right)$

The proofs for these identification results can be found in Appendix B.3.

The identification results warrant a couple remarks. First, the instrument measure $\hat{\theta}_{2i}$ is necessary to identify variance and covariances. Without the instrument, one would be unable to distinguish variance of misperceptions from variance of the measurement error ϵ_{1i}. Similarly, one would be unable to distinguish correlation between squared misperceptions and income, elasticity, or squared tax rate from correlation between the square of the measurement error ϵ_{1i}^2 and the same, which amounts to an inability to distinguish correlation between the conditional variance of misperceptions and income, elasticity, or squared tax rate from correlation between the conditional variance of the error term ϵ_{1i} and the same. Second, the variance and covariance results all contain a bias correction term. The bias correction terms result from the fact that the underlying secondary tax rate differs slightly from the focal tax rate. Because the rates differ, the relationship between $\hat{\theta}_{1i}$ and $\hat{\theta}_{2i}$ is influenced partly by how the respondent's perception of the gap between the two tax rates, $\tilde{a}_i$, covaries with the perceived focal tax rate, income, elasticity, or true focal tax rate. The bias correction terms measure and account for this.

These identification results allow us to combine a reduced form elasticity estimate with survey-based measures of misperceptions to measure DWL under tax rate misperception as in Equation (3).

4.3 Survey Data

To identify DWL under misperception of tax rates from college aid phaseout, we combine our estimates of the ETI from Section 3 with an online survey we administered to directly measure parents' perceptions of the college aid implicit tax schedule.

The survey was administered through Prolific Academic, a website that connects survey administrators with survey takers, between September 26, 2023 and October 19, 2023. The survey targeted parents with a child born between the years 2001 and 2023, so that the child would be between age 0 and 22 at the time the survey. As we discuss further below, our main estimates pool data from respondents with children from age 0 to 22, and we show that our results are similar if we restrict only to respondents whose child is near college age. The survey instructed the respondent to consider one child in particular and to answer all questions with regard to that child. Additional screeners for eligible participants included US nationality and fluency in English. The sample includes respondents from across the US and is not probability weighted. In total, 3,158 respondents completed the survey. Each respondent was paid between $3 and $4. All survey responses are anonymous.

In addition to asking respondents to provide their perceptions of the college aid schedule, the survey measures information including the respondent's state, previous year income, share of income from nonlabor sources, child age, and assessments of the child's likelihood of attending a four-year college or specifically a four-year college in the respondent's state. A full description of the survey can be found in Supplementary Appendix D.

The survey includes a primary and a secondary elicitation of the respondent's perception of college aid phaseout rates. For the focal rate, the respondent is asked to provide their best guess of average aid provided to in-state students attending a four-year public college in the respondent's state, with family income in each of the following bins: $0-30k, $30-48k, $48-75k, $75-110k, and above $110k. These income bins correspond to the bins for which colleges report publicly available statistics on average aid through IPEDS – the same data used in Section 3 to measure the MTR from college aid for individuals in the tax data – so that we can compare respondent's guesses against the true aid schedule in their state. We follow the same approach described in Section 3.1 to match respondents to their nearest income bin and measure their corresponding perceived tax rate from aid phaseout, $\hat{\tilde{\tau}}_{1i}$, and their true tax rate from aid phaseout, τ_i. Following Section 4.2, we then define the individual's misperception parameter as $\hat{\theta}_{1i} = \frac{\hat{\tilde{\tau}}_{1i}}{\tau_i}$.

For the secondary elicitation, the respondent is asked to provide their best guess of aid provided to an in-state student attending a four-year public college in the respondent's state, whose total parent income is $45k, $55k, $65k, or $75k. This elicitation occurs on a separate screen from the primary elicitation. We define our secondary tax rate perception measure $\hat{\tilde{\tau}}_{2i}$ by taking the difference between the respondent's guesses of aid for a student with parent income of $75k versus $45k and dividing by the income difference of $30k. We similarly define the individual's instrument-based misperception $\hat{\theta}_{2i} = \frac{\hat{\tilde{\tau}}_{2i}}{\tau_i}$.

To improve the precision of our estimates, we winsorize both $\hat{\tilde{\tau}}_{1i}$ and $\hat{\tilde{\tau}}_{2i}$ to their respective 2.5 and 97.5 percentiles. We also winsorize $\hat{\theta}_{1i}$ and $\hat{\theta}_{2i}$ to their respective 2.5 and 97.5 percentiles after constructing these initially using the un-winsorized values of $\hat{\tilde{\tau}}_{1i}$ and $\hat{\tilde{\tau}}_{2i}$.

We make four restrictions for our analysis sample. First, we restrict to respondents whose household income from the previous year is in the $50k to $150k range, to match the income range in our analysis of the elasticity of parent income in Section 3. Because our survey did not screen on income, this reduces our sample to 1,723. Second, we restrict to respondents who identify one child who is age 0 to 22 and

is at least 15 years younger than the respondent, which reduces our sample size to 1,672. Third, because our measures of perceptions of college aid phaseout rates, described below, relate to four-year public colleges in the respondent's state, we restrict to respondents who answer that their child (i) has not yet attended college and is "Somewhat likely" or "Extremely likely" to attend a four-year public college in the respondent's state,[21] or (ii) currently attends or previously attended a four-year public college in the respondent's state. This reduces our sample size to 1,294. Finally, to ensure data quality, we restrict to respondents who pass an attention check question at the end of the survey and who have a valid Prolific ID, which produces our final analysis sample with 1,279 respondents. Data quality is high, reflected in the fact that only 1% of respondents are screened out due to failing the attention check or not having a valid Prolific ID, and seen similarly in the median survey completion time of 13 minutes[22].

Finally, we weight all the following analysis to be representative of the US college population based on the respondent's state and income. To do this, we merge in total enrollment counts of full-time, first-time, degree-seeking students at four-year public colleges by state and income bin using IPEDS data. We weight individuals in the survey who have income in the \$50-\$75k range based on enrollment counts in the \$48k-\$75k range, those with income in the \$75k-\$110k range based on enrollment counts in the \$75k-\$110k range, and those with income in the \$110k-\$150k range based on enrollment counts above \$110k. Our survey data successfully represents individuals across nearly all state and income bin combinations: of the 153 possible combinations of the 50 US states plus D.C. interacted with the three income bins, 135 are represented in our analysis sample; D.C. is not represented, but each of the 50 states is represented.[23]

Table A.3 reports descriptive statistics for our analysis sample. Average income is \$98k, and 5% of respondents report that at least 20% of their income is from nonlabor sources. 63% of respondents have completed a bachelor's degree. The average respondent is 40 years old. The last three rows of Table A.3 preview aspects of the following DWL estimates. On average, respondents slightly underestimate their MTR from aid. The average MTR from aid is 14%, and the average respondent's guess of their MTR from aid is 11%. The average misperception parameter – the ratio of perceived MTR to true MTR – is 0.89. However, there is high (non-instrumented) variation in misperceptions. The standard deviation of MTR guesses is over five times as large as the standard deviation of true MTRs, at 0.23 compared with 0.04.

[21] Specifically, we restrict to respondents who answer "Somewhat likely" or "Extremely likely" to each of the following questions: "How likely is it that this child will attend a four-year college as a full-time student at some point?", and "If this child does eventually attend a four-year college, how likely are they to attend a public four-year college in [state]?", where the second question fills in the respondent's state. When the respondent is asked the second question, they are simultaneously shown a list of several examples of four-year public colleges in their state.

[22] The survey contains several sections in addition to those used in this paper. The sections are as follows: survey introduction and background questions about family and child; baseline guesses of college costs and aid schedules; information treatment; post-treatment guesses of aid schedules; demographics, attention check, and survey feedback. Supplementary Appendix D contains the full survey instrument.

[23] The combinations not found in our analysis sample include the \$48-75k income bin for D.C., North Dakota, and Rhode Island; the \$75-110k income bin for Alaska, D.C., Hawaii, New Mexico, Rhode Island, and Vermont; and the \$110k plus bin for Alaska, D.C., Hawaii, Montana, North Dakota, New Hampshire, Vermont, West Virginia, and Wyoming.

The standard deviation of individual misperceptions is 1.64, which is similarly large compared with a mean of 0.89.

4.4 Results

4.4.1 Misperception Moments and Deadweight Loss

We begin by documenting the first stage relationship between the respondent's perception of their focal tax rate from aid phaseout $\hat{\tilde{\tau}}_{1i}$ and the survey instrument $\hat{\tilde{\tau}}_{2i}$. Figure A.16 plots, for 20 bins of the instrument, the average value of the instrument and the average value of the respondent's best guess of their own MTR within that bin. The figure includes a quadratic fit. There is a clear positive relationship with a negative quadratic term. Regressing $\hat{\tilde{\tau}}_{1i}$ on linear and quadratic terms in $\hat{\tilde{\tau}}_{2i}$ at the individual level confirms this, as both terms are statistically significant at the 1% level. The F-statistic for the individual-level regression is 78, and the R^2 is 0.31. We conclude that the instrument captures important variation in $\hat{\tilde{\tau}}_{1i}$, and as a result the instrument will allow us to consistently estimate the components of Equation (3) under the assumptions described in Section 4.2.

We now proceed to estimating each of the parameters involving misperceptions in Equation (3) using the survey data, following the identification results in Section 4.2. The results are shown in Table 2. The estimates in Column 1 do not use the instrument and they exclude bias correction terms – for example, $Var(\theta_i)$ is estimated with $\hat{Var}(\hat{\theta}_{1i})$; the estimates in Column 2 make use of the instrument but still exclude bias correction terms – for example, $Var(\theta_i)$ is estimated with $\hat{Cov}(\hat{\theta}_{1i}, \hat{\theta}_{2i})$; and, finally, the estimates in Column 3 – our preferred estimates – make use of the instrument and include bias correction terms as in the identification results of Section 4.2. To implement the bias correction terms, we measure $E[\hat{\tilde{\tau}}_{1i}|\tau_i]$ and $E[\hat{\tilde{\tau}}_{2i} - \hat{\tilde{\tau}}_{1i}|\tau_i]$ by regressing $\hat{\tilde{\tau}}_{1i}$ and $\hat{\tilde{\tau}}_{2i}$ linearly on τ_i and taking the fitted values. All standard errors in Table 2 are bootstrapped with 1,000 repetitions.

Table 2 shows that misperception affects DWL through several channels, and it also highlights the utility of our instrument-based approach with corresponding bias correction. The first two rows are unchanged across columns because these do not rely on the instrument, and as a result they also do not require bias correction. The first row shows that the ratio that converts between the structural elasticity $e(X_i)$ and the reduced form ETI $\beta^{ETI}(X_i)$, which is $\frac{E[\ln(1-\tau_i)|X_i]}{E[\ln(1-\tilde{\tau}_i)|X_i]}$, is very close to and not statistically different from one.[24] The second row reports estimates of Bias, which we define as $E[\theta]^2 - 1$. Consistent with the summary statistics in Table A.3, we see that on average parents underestimate their tax rate from aid phaseout, which reduces DWL. The third row shows estimates of the variance of misperceptions, which differ importantly across the three columns. While the raw variance of our misperception measure is 2.69, restricting to instrumented variance reduces this to 1.89, and incorporating the bias correction term further reduces this to 1.81. The non-instrumented number significantly over-estimates variance because it also incorporates variation that is due to measurement error, which the instrument purges, demonstrating the utility of our instrument-based approach. Our preferred estimate that incorporates the instrument and bias correction remains large at 1.81 and indicates that variance of

[24] We also do not detect significant differences in the conversion ratio when we split the sample into those with high or low non-labor income, which is our proxy for elasticity differences.

Table 2: Misperception Moments

	Raw	Instrumented	Instrumented and bias corrected
	(1)	(2)	(3)
Elasticity conversion ratio	1.01	1.01	1.01
	(0.05)	(0.05)	(0.05)
Bias	-0.21	-0.21	-0.21
	(0.09)	(0.09)	(0.09)
Variance	2.69	1.89	1.81
	(0.21)	(0.20)	(0.20)
Cov with income (thousands)	6.82	12.10	10.26
	(6.72)	(6.92)	(6.86)
Cov with elasticity (Non-labor income share >= 20%)	0.07	0.02	0.11
	(0.06)	(0.05)	(0.08)
Cov with tax rate	-0.014	-0.019	-0.013
	(0.002)	(0.003)	(0.002)
N	1,279	1,279	1,279

Notes: This table shows estimates of the components of Equation (3) involving tax rate perceptions, which includes the moments involving individual misperception θ_i and the conversion ratio that translates between the reduced form ETI and the structural elasticity as shown in Equation (4). We define Bias as $E[\theta_i]^2 - 1$. The first column does not use the instrument. The second column uses the instrument but ignores bias correction terms. The third column uses the instrument and includes bias correction terms. Standard errors are bootstrapped with 1,000 repetitions. All values are weighted to match the US college population based on state and family income range using IPEDS data.

misperceptions across individuals increases DWL. The next two rows show estimates of the covariance of squared misperceptions with income and with elasticity. We proxy for elasticity using an indicator for whether non-labor sources account for at least 20% of the respondent's household income,[25] following our finding in Section 3 that households with a high share of non-wage income reduce income by a greater extent in response to the college aid implicit tax. Across all three columns, we do not detect a statistically signficant relationship between squared misperceptions and either income or elasticity. Finally, the last row shows estimates of the covariance of squared misperceptions with the square of the true tax rate. We see differences in the magnitude across columns, but in all cases we observe a negative covariance between squared misperceptions and squared tax rate, which reduces DWL.

Finally, following Equation (3), we combine our estimates of misperception moments from Table 2 with our ETI estimates from Table 1 in Section 3 to identify DWL from the college aid implicit tax and the impact of misperception. Because we did not detect statistical evidence that the structural elasticity differs from the reduced form ETI or that the covariances between squared misperceptions and income or elasticity are different from zero in our setting, we omit the elasticity conversion factor and these two covariances from our DWL calculations. Note that because point estimates are positive for the two omitted covariances and above one for the elasticity

[25] Specifically, the survey question for a respondent with a partner reads: "Roughly, what percent of your total household income comes from sources **other than** work (wages, salaries, tips, etc.), self-employment, or business ownership? Examples of other income sources include capital gains, interest, dividends, rental income, retirement income, or transfer payments. Consider your income and your partner's income." Respondents select one of the following options: 0% - 5%, 5% - 10%, 10% - 20%, 20% - 40%, 40% - 60%, or "More than 60%". We aim to target one dimension along which elasticity importantly varies, though is it possible that misperception may covary with elasticity along other dimensions, which is an important avenue for future work to consider.

conversion ratio, our estimates of DWL and the impact from misperception would only be higher if we included the point estimates of these terms in our DWL calculations, so our estimates are conservative in this regard. To estimate the components $E[Z_{i,1}e(X_i)\tau_i^2]$ and $E[Z_{i,1}]E[e(X_i)]$ that appear in Equation (3), we simply compute the corresponding sample averages directly in the tax data.[26] Consistent with our ETI estimates, we restrict to households with baseline income averaged over 2005 and 2006 between \$50-150k. We also focus on households with a child who is enrolled at a four-year in-state public college or a four-year private college in 2016. We count all households with a child enrolled in a single year so our estimates correspond to total annual DWL among all middle-income households with a child in college. This attributes the ETI estimate in Section 3 to all years a family's income affects a child's college aid; because other factors like college costs also influence income when the child is age 18+, we use age 17 to isolate the response to the college aid implicit tax, and we assume this substitution response is active in all years that income affects aid. The total DWL from the college aid implicit tax among all families is almost surely greater than our estimates among middle-income families, as for example we detect income reduction among higher-income families in Section 3.

Table 3: Annual Deadweight Loss from the College Aid Implicit Tax And the Impact of Tax Rate Misperception

	Per family	Total (in millions) Oldest child only	Total (in millions) All children	Percent change due to misperception
Total	\$244	\$157	\$247	+127%
without misperception	\$107	\$69	\$109	-
due to misperception	\$137	\$88	\$138	+127%
bias	-\$23	-\$15	-\$23	-21%
variance	\$194	\$125	\$197	+181%
covariance with tax rate	-\$35	-\$23	-\$36	-33%
N		645,076	1,011,502	

Notes: This table shows estimates of annual DWL due to the college aid implicit tax. The table reports total DWL, the share that would occur without misperception, and the share that occurs due to misperception and its sub-channels. Estimates are based on Equations (3) and (4), and they rely on both the tax data and the survey data. We plug in the ETI estimates from Table 1 in Section 3 and moments involving income, elasticity, and tax rate from the tax data. We plug in estimates of moments involving tax rate perceptions from Table 2 based on the survey data and the identification results in Section 4.2. The first column shows estimates per family. The second column shows estimates (in millions) summed across the 645,076 households in the tax data with baseline income averaged over 2005 and 2006 between \$50k and \$150k, whose oldest child turned 18 between 2012 and 2016 and was enrolled in 2016 in a four-year in-state public college or a four-year private college that has data in IPEDS on average aid by family income bin. The third column shows estimates (in millions) summed across the 1,011,502 households in the tax data with baseline income averaged over 2005 and 2006 between \$50k and \$150k, who have a child who in 2016 is between the ages of 17 and 25 and is enrolled in a four-year in-state public college or a four-year private college. The fourth column shows the percent change in DWL due to misperception and its sub-channels.

The DWL results are shown in Table 3. The first column of Table 3 reports annual

[26] We allow the elasticity to vary by high or low non-wage income share. $Z_{i,1}$ conceptually measures the individual's income in the taxed period if the tax were not in place. Therefore, to measure $Z_{i,1}$, we start with each family's observed income and then subtract the effect on income due to the college aid implicit tax. In the small number of cases in which we measure a (small) negative MTR from IPEDS data, we replace τ_i^2 as zero.

average DWL at the household level.[27] The second column scales the per-household effects to report the total annual DWL summed over all households whose oldest child attends a four-year in-state public college or a four-year private college in 2016. The third column, our best estimates of total annual DWL, scales the per-household effects to report the total annual DWL summed over all households with any child attending a four-year in-state public college or a four-year private college in 2016. In total, the college aid implicit tax creates $247 million in DWL annually among our target population – middle-income households with a child attending a four-year in-state public or four-year private college. If perceptions were accurate across all households, DWL would instead be $109 million. Interestingly, because in our setting the structural and reduced form elasticities are roughly equal, this number also corresponds to the total DWL a naive researcher would compute from the ETI alone without accounting for misperception, meaning the naive researcher would miss more than half of the true DWL. In general, when the structural and reduced form elasticities differ, the naive researcher's estimate of DWL that only relies on the ETI will not even correspond to DWL under accurate perceptions, and will be essentially meaningless. In total, misperception increases annual DWL in our setting by $138 million, representing a 127% increase relative to the case in which perceptions were accurate. The increase in DWL from misperception is driven by heterogeneity in misperceptions across individuals. The fact that parents on average underestimate the tax rate from aid phaseout reduces DWL by 21%. However, the substantial variance of misperceptions across individuals increases DWL by 181%. The negative correlation between squared misperceptions and squared tax rate reduces DWL by 33%.

4.4.2 Robustness

Survey sample.—Our main estimates above pool survey data among respondents whose child is age 0 to 22 for precision, which is valid if perceptions are similar over this age range but can affect our estimates if perceptions are systematically different when a child is near college age. Table A.4 reports estimates of the relevant moments of misperceptions (comparable to Table 2) among respondents whose child is age 14 to 22 (high school and college age) or 17 to 22 (college age). While not identical, the estimates in the narrower samples are similar to those in our pooled main sample, and we cannot reject equality with the pooled sample for any of the moments. Using the misperception moments from the narrower age groups also produces a large net increase in DWL due to misperception, as we estimate that misperception increases DWL by 99% using the moments from the 14 to 22 sample and by 86% using the moments from the 17 to 22 sample.

Accurate measurement of moments.—A key assumption is that measurement error is not correlated between the focal tax rate perception measure and the instrument. Our survey includes design choices intended to reduce concerns of correlated measurement error driving our results – using a secondary tax rate that is slightly different from the focal rate, and asking for perceptions of aid levels at various points of income that differ between the two measures rather than asking for rates directly. Ad-

[27] We compute the relevant means in the tax data using middle-income households whose oldest child attends a four-year in-state public college or a four-year private college in 2016.

ditionally, the considerable reduction in variance between the raw and instrumented measures shown in Table 2 both demonstrates the utility of our instrument-based approach and points against respondents simply repeating the same (noisy) perceptions for both measures. Despite these factors, one may still think that less careful respondents might click through the survey quickly and repeat similar aid perceptions between the primary and secondary elicitations in order to exert less effort, which would violate error independence. Instead, we find that survey respondents who took above median time to complete the survey exhibit the same patterns as in our main estimates but to a greater extent, as shown in Column 4 of Table A.4. We find statistical evidence that this group underestimates their tax rate on average by a greater extent, has higher variance of misperceptions, and has a more negative covariance between squared misperceptions and squared tax rate. We also find a positive and statistically significant covariance between squared misperceptions and income in this group. Re-estimating DWL using the misperception moments from those who took above median time to complete the survey, we find that misperception increases DWL by 150%, while accounting for the positive covariance between squared misperceptions and income produces an increase of 172%. If anything, it appears that less careful respondents may attenuate our main results by providing less signal of their perceptions, rather than spuriously driving our results. In both the above-median-time sample and the main survey sample, overturning the qualitative finding that misperception increases DWL would require that more than 70% of the instrumented and bias-corrected variance of misperceptions reflects correlated measurement error across the two elicitations rather than genuine heterogeneity in perceptions.

Another potential concern is that our survey measures of tax rates are based on aggregate IPEDS data on aid across income bins rather than an individual's personal tax rate, and this discrepancy may drive our results. We use IPEDS measures because this is the only readily available, consistently defined measure of true aid phaseout rates (discussed in Section 3.1), which is needed to compute misperception. The drawback is that this schedule may differ from the respondent's own aid schedule on two margins. First, our ETI estimates come from families whose child attends a four-year in-state public college or a four-year private college. Our survey measures of misperceptions of aid phaseout at four-year in-state public colleges may not generalize to four-year private colleges. However, because aid schedules are generally more variable across private colleges than across public colleges, if anything misperceptions at four-year private colleges are likely to have higher variance than we measure among four-year in-state publics. Second, our measures are based on average aid across bins of family income, which do not align exactly with the tax rate a given family faces at their exact level of income and assets. Our approach is based on the intuition that a parent who overestimates (underestimates) average aid phaseout is also likely to overestimate (underestimate) their own specific aid phaseout rate, similar to evidence in Wiswall and Zafar (2015) and Bleemer and Zafar (2018) that individuals form perceptions of individual-specific parameters in part based on their perceptions of aggregate parameters. To contradict our main results, it would have to be the case that parents systematically have a much more precise perception of their own aid phaseout rate than of average rates. However, again the bias may well work in the opposite direction: because aid information does not transparently communicate individual-specific phaseout rates but aggregate statistics are more readily available,

misperceptions of individual rates are likely to be more variable than misperceptions of aggregate rates.

Welfare implications of the ETI.—Finally, an important consideration for our DWL calculations is the extent to which the ETI we measure in Section 3 reflects a reduction of income as opposed to simple re-timing. While we see evidence that capital gains may shift from when the child is age 17 to when they are 16, we do not see a clear spike in aggregate income at age 16, in contrast to the clear drop in aggregate income at age 17. Moreover, shifting income from when the child is age 17 to when they are age 18 would provide no benefit, as income at age 18 continues to affect college aid. Therefore, the increase of income we observe at age 18 is better explained by the introduction of college costs than by re-timing, which is also consistent with the fact that wage income increases at age 18 despite not decreasing at age 17. Thus, re-timing does not appear to be an important explanation of the effects we observe.

A related question is the extent to which the ETI we estimate reflects real economic distortions as opposed to costless misreporting or sheltering (Chetty 2009). Because the ETI we estimate is driven by non-wage factors, one might worry it reflects pure evasion. However, several facts suggest that the observed elasticity captures real efficiency costs. First, parents adjust real savings behavior, seen in reduced IRA distributions and increased pre-tax retirement contributions. Second, Chetty (2009) notes that shifting income between tax bases (e.g. shifting between corporate and personal income) tends to be a more important form of costless evasion than simple misreporting in practice. However, the taxable income measure in our setting is AGI, a relatively broad income measure that limits the ability to shift income between tax bases without real economic costs. Third, while some business income reductions could theoretically reflect costless misreporting, Kim (Forthcoming) uses small business checking account data and finds that entrepreneurs reduce revenues (measured as inflows to the business checking account) when a child is age 17 with little simultaneous change in expenses (measured as outflows on business checking, debit card, and credit card accounts) or even a slight reduction, which is very unlikely to be explained by misreporting. We cannot rule out some degree of costless misreporting, in which case our absolute DWL estimates should be scaled roughly by the share of income reduction that reflects real behavior change. Our core conclusion regarding the proportional impact of tax rate misperceptions on DWL is effectively unchanged by these considerations as long as some of the response reflects real distortions.

5 Discussion and Conclusion

This paper studies a single question: What is the welfare cost that results from the college aid implicit tax on parent income? Because the tax rates that result from college aid phaseout are difficult to learn, answering this single question requires information on two ingredients: the response of parent income to aid phaseout, and distributional features of tax rate misperceptions. In fact, to assess the welfare cost of any tax that is misperceived, a researcher must combine information on both of these components – simply estimating the elasticity of income is not sufficient to

identify the welfare cost of a misperceived tax. While recent theoretical work in behavioral public finance has shown that distributional features of misperceptions can be key determinants of the welfare cost of a tax (Farhi and Gabaix 2020), there is no previous method to implement this theory and measure the welfare cost of a misperceived tax without imposing structure on the form of misperception, which is a severely limiting requirement in many settings. To answer our motivating question, we provide the first estimates of the response of parent income to the college aid implicit tax and we develop and apply a new method to identify the welfare cost of this response given that parents often misperceive aid phaseout rates.

Our analysis of parent income response shows that, had one attempted to apply an "off-the-shelf" elasticity of income from an alternative setting to ours, they would have incorrectly overstated the elasticity to the college aid implicit tax. For example, the seminal work of Gruber and Saez (2002) estimates an elasticity of 0.11 for the \$50-100k income range (though with a standard error of 0.22), with higher elasticities at higher incomes. Saez et al. (2012) summarize from the literature that the best available elasticity estimates range from 0.12 to 0.40. Kleven and Schultz (2014) exploit panel variation in income tax schedules in Denmark, finding an overall ETI of 0.08, and an ETI of 0.18 for a large reform. Looney and Singhal (2006) use a similar style of identifying variation – they focus on anticipated changes in tax rates that occur over a \$10k-wide income range following a change in number of dependents – and estimate an elasticity of 0.75 (standard error 0.38). Possible explanations for our somewhat lower elasticity of 0.04 include that our population of parents of college-age children tend to be in the prime earning age range; that parents may face steady liquidity constraints around the time a child goes to college that limit their capability to reduce income; and that our estimates are specific to the \$50-\$150k income range.

At the same time, our welfare analysis finds substantial DWL from the college aid implicit tax, more than half of which is due to tax rate misperceptions and would be missed by a naive researcher who did not account for misperception. In fact, because the increase in DWL comes from heterogeneity in misperceptions across individuals, even a researcher who accounted for the average perceived tax rate would still miss the substantial increase in DWL due to misperception. The administrative cost of making information on financial aid schedules more easily accessible is likely much lower than the \$138 million in DWL that we estimate results annually from misperception among middle-income families.

Beyond allowing us to measure the welfare cost of the college aid implicit tax given parent misperception of aid phaseout rates, the tools we introduce can easily be adapted to study the welfare cost of other misperceived taxes and the impact of misperception on this cost. This includes progressive income tax schedules or other implicit taxes due to means-tested benefits such as the Social Security earnings test (Liebman and Luttmer 2015), the Child Tax Credit (Feldman et al. 2016), the EITC (Chetty et al. 2013), or others. This presents an exciting opportunity for further empirical implementation of the recent theoretical insights that heterogeneity in tax rate misperceptions can a crucial determinant of welfare and optimal tax design.

References

Allcott, Hunt, Daniel Cohen, William Morrison, and Dmitry Taubinsky (2025) "When Do "Nudges" Increase Welfare?" *American Economic Review*, 115 (5), 1555–96, 10.1257/aer.20231304.

Allcott, Hunt, Benjamin B Lockwood, and Dmitry Taubinsky (2019) "Regressive Sin Taxes, with an Application to the Optimal Soda Tax*," *The Quarterly Journal of Economics*, 134 (3), 1557–1626, 10.1093/qje/qjz017.

Bettinger, Eric P., Bridget Terry Long, Philip Oreopoulos, and Lisa Sanbonmatsu (2012) "The Role of Application Assistance and Information in College Decisions: Results from the H&R Block Fafsa Experiment," *The Quarterly Journal of Economics*, 127 (3), 1205–1242, 10.1093/qje/qjs017.

Bhargava, Palaash, Sandra E Black, Jeffrey T Denning, Robert W Fairlie, and Oded Gurantz (2025) "A Family Affair: The Effects of College on Parent and Student Finances," Working Paper 33497, National Bureau of Economic Research, 10.3386/w33497.

Bleemer, Zachary and Basit Zafar (2018) "Intended college attendance: Evidence from an experiment on college returns and costs," *Journal of Public Economics*, 157, 184–211, https://doi.org/10.1016/j.jpubeco.2017.11.002.

Brusco, Giacomo and Benjamin Glass (2020) "Attending to inattention: Identification of deadweight loss under nonsalient taxes," *Journal of Public Economic Theory*, 22 (1), 5–24, https://doi.org/10.1111/jpet.12401.

Bulman, George, Robert Fairlie, Sarena Goodman, and Adam Isen (2021) "Parental Resources and College Attendance: Evidence from Lottery Wins," *American Economic Review*, 111 (4), 1201–40, 10.1257/aer.20171272.

Case, Karl E. and Michael S. McPherson (1986) "Student Aid Incentives and Parental Effort: The Impact of Need-Based Aid on Savings and Labor Supply," Technical report prepared for the Washington office of the College Entrance Examination Board.

Chetty, Raj (2009) "Is the Taxable Income Elasticity Sufficient to Calculate Deadweight Loss? The Implications of Evasion and Avoidance," *American Economic Journal: Economic Policy*, 1 (2), 31–52, 10.1257/pol.1.2.31.

Chetty, Raj, John N. Friedman, and Emmanuel Saez (2013) "Using Differences in Knowledge across Neighborhoods to Uncover the Impacts of the EITC on Earnings," *American Economic Review*, 103 (7), 2683–2721, 10.1257/aer.103.7.2683.

Chetty, Raj, John N Friedman, Emmanuel Saez, Nicholas Turner, and Danny Yagan (2020) "Income Segregation and Intergenerational Mobility Across Colleges in the United States," *The Quarterly Journal of Economics*, 135 (3), 1567–1633, 10.1093/qje/qjaa005.

Chetty, Raj, Adam Looney, and Kory Kroft (2009) "Salience and Taxation: Theory and Evidence," *American Economic Review*, 99 (4), 1145–77, 10.1257/aer.99.4.1145.

Chorniy, Anna, Amy Finkelstein, and Matthew J Notowidigdo (2025) "Paternalistic Social Assistance: Evidence and Implications from Cash vs. In-Kind Transfers," Working Paper 34506, National Bureau of Economic Research, 10.3386/w34506.

Colas, Mark, Sebastian Findeisen, and Dominik Sachs (2021) "Optimal Need-Based Financial Aid," *Journal of Political Economy*, 129 (2), 492–533, 10.1086/711952.

College Board (2024) "Trends in Student Aid 2024: Data in Excel [Excel spreadsheet]," https://research.collegeboard.org/media/xlsx/Trends-in-Student-Aid-2024-excel-data.xlsx.

Cook, Emily E. and Sarah Turner (2022) "Progressivity of pricing at US public universities," *Economics of Education Review*, 88, 102239, https://doi.org/10.1016/j.econedurev.2022.102239.

Craig, Ashley C. and Joel Slemrod (2024) "Tax Knowledge and Tax Manipulation: A Unifying Model," *Journal of Political Economy Microeconomics*, 2 (2), 298–334, 10.1086/727557.

Cronin, Julie-Anne and Joseph Gray-Hancuch (2024) "Barriers to Claiming Education Tax Credits for Low-Income Students," Working Paper 125, Office of Tax Analysis.

Darolia, Rajeev (2017) "Assessing the College Financial Aid Work Penalty," *The Journal of Higher Education*, 88 (3), pp. 350–375, `https://www.jstor.org/stable/26772066`.

de Bartolome, Charles A.M. (1995) "Which tax rate do people use: Average or marginal?" *Journal of Public Economics*, 56 (1), 79–96, https://doi.org/10.1016/0047-2727(93)01409-4.

Dick, Andrew W. and Aaron S. Edlin (1997) "The implicit taxes from college financial aid," *Journal of Public Economics*, 65 (3), 295–322, https://doi.org/10.1016/S0047-2727(97)00018-2.

Dobkin, Carlos, Amy Finkelstein, Raymond Kluender, and Matthew J. Notowidigdo (2018) "The Economic Consequences of Hospital Admissions," *American Economic Review*, 108 (2), 308–52, 10.1257/aer.20161038.

Dynarski, Susan, CJ Libassi, Katherine Michelmore, and Stephanie Owen (2021) "Closing the Gap: The Effect of Reducing Complexity and Uncertainty in College Pricing on the Choices of Low-Income Students," *American Economic Review*, 111 (6), 1721–56, 10.1257/aer.20200451.

Dynarski, Susan M. and Judith E. Scott-Clayton (2006) "The Cost of Complexity in Federal Student Aid: Lessons from Optimal Tax Theory and Behavioral Economics," *National Tax Journal*, 59 (2), 319–356, 10.17310/ntj.2006.2.07.

Edlin, Aaron S. (1993) "Is College Financial Aid Equitable and Efficient?" *Journal of Economic Perspectives*, 7 (2), 143–158, 10.1257/jep.7.2.143.

Eissa, Nada and Hilary W. Hoynes (2006) "Behavioral Responses to Taxes: Lessons from the EITC and Labor Supply," *Tax Policy and the Economy*, 20, 73–110, 10.1086/tpe.20.20061905.

Fan, Kristy, Tyler J. Fisher, and Andrew A. Samwick (2025) "The Insurance Value of Financial Aid," *Education Finance and Policy*, 1–30, 10.1162/edfp_a_00442.

Farhi, Emmanuel and Xavier Gabaix (2020) "Optimal Taxation with Behavioral Agents," *American Economic Review*, 110 (1), 298–336, 10.1257/aer.20151079.

Feldman, Naomi E., Peter Katuščák, and Laura Kawano (2016) "Taxpayer Confusion: Evidence from the Child Tax Credit," *American Economic Review*, 106 (3), 807–35, 10.1257/aer.20131189.

Feldstein, Martin (1995) "College Scholarship Rules and Private Saving," *The American Economic Review*, 85 (3), 552–566, http://www.jstor.org/stable/2118187.

——— (1999) "Tax Avoidance and the Deadweight Loss of the Income Tax," *The Review of Economics and Statistics*, 81 (4), 674–680, http://www.jstor.org/stable/2646716.

Freyaldenhoven, Simon, Christian Hansen, Jorge Pérez Pérez, and Jesse M. Shapiro (2025) *Visualization, Identification, and Estimation in the Linear Panel Event-Study Design*, 225–268, Econometric Society Monographs: Cambridge University Press.

Fujii, Edwin T. and Clifford B. Hawley (1988) "On the Accuracy of Tax Perceptions," *The Review of Economics and Statistics*, 70 (2), 344–347, http://www.jstor.org/stable/1928321.

Fuster, Andreas, Greg Kaplan, and Basit Zafar (2021) "What Would You Do with $500? Spending Responses to Gains, Losses, News, and Loans," *The Review of Economic Studies*, 88 (4), 1760–1795, 10.1093/restud/rdaa076.

Fuster, Andreas and Basit Zafar (2023) "Chapter 4 - Survey experiments on economic expectations," in Bachmann, Rüdiger, Giorgio Topa, and Wilbert van der Klaauw eds. *Handbook of Economic Expectations*, 107–130: Academic Press, https://doi.org/10.1016/B978-0-12-822927-9.00010-0.

Gruber, Jon and Emmanuel Saez (2002) "The elasticity of taxable income: evidence and implications," *Journal of Public Economics*, 84 (1), 1–32, https://doi.org/10.1016/S0047-2727(01)00085-8.

Hastings, Justine and Jesse M. Shapiro (2018) "How Are SNAP Benefits Spent? Evidence from a Retail Panel," *American Economic Review*, 108 (12), 3493–3540, 10.1257/aer.20170866.

Hilger, Nathaniel G. (2016) "Parental Job Loss and Children's Long-Term Outcomes: Evidence from 7 Million Fathers' Layoffs," *American Economic Journal: Applied Economics*, 8 (3), 247–83, 10.1257/app.20150295.

Hoxby, Caroline M and Christopher Avery (2012) "The Missing "One-Offs": The Hidden Supply of High-Achieving, Low Income Students," Working Paper 18586, National Bureau of Economic Research, 10.3386/w18586.

Hoxby, Caroline and Sarah Turner (2013) "Expanding College Opportunities for High-Achieving, Low-Income Students," Discussion Paper 12-014, Stanford Institute for Economic Policy Research.

Internal Revenue Service (2023) "EITC reports and statistics," December, https://www.irs.gov/credits-deductions/individuals/earned-income-tax-credit/eitc-reports-and-statistics.

Kim, Olivia S. (Forthcoming) "Trading Off Business and Family Investments: Evidence from U.S. Entrepreneurial Households," *Management Science*, 10.2139/ssrn.3527860.

Kleven, Henrik Jacobsen and Esben Anton Schultz (2014) "Estimating Taxable Income Responses Using Danish Tax Reforms," *American Economic Journal: Economic Policy*, 6 (4), 271–301, 10.1257/pol.6.4.271.

Kleven, Henrik, Camille Landais, and Gabriel Leite-Mariante (2025) "The Child Penalty Atlas," *The Review of Economic Studies*, 92 (5), 3174–3207, 10.1093/restud/rdae104.

Kleven, Henrik, Camille Landais, and Jakob Egholt Søgaard (2019) "Children and Gender Inequality: Evidence from Denmark," *American Economic Journal: Applied Economics*, 11 (4), 181–209, 10.1257/app.20180010.

Kostøl, Andreas R. and Andreas S. Myhre (2021) "Labor Supply Responses to Learning the Tax and Benefit Schedule," *American Economic Review*, 111 (11), 3733–66, 10.1257/aer.20201877.

Levine, Phillip B. (2022) *A Problem of Fit: How the Complexity of College Pricing Hurts Students—and Universities*, Chicago: University of Chicago Press, 176.

Liebman, Jeffrey B. and Erzo F. P. Luttmer (2015) "Would People Behave Differently If They Better Understood Social Security? Evidence from a Field Experiment," *American Economic Journal: Economic Policy*, 7 (1), 275–99, 10.1257/pol.20120081.

Liebman, Jeffrey B and Richard J Zeckhauser (2004) ""Schmeduling"," Working paper.

Long, Mark (2004) "The impact of asset-tested college financial aid on household savings," *Journal of Public Economics*, 88 (1), 63–88, https://doi.org/10.1016/S0047-2727(02)00136-6.

Looney, Adam and Monica Singhal (2006) "The Effect of Anticipated Tax Changes on Intertemporal Labor Supply and the Realization of Taxable Income," Working Paper 12417, National Bureau of Economic Research, 10.3386/w12417.

Mas, Alexandre and Amanda Pallais (2017) "Valuing Alternative Work Arrangements," *American Economic Review*, 107 (12), 3722–59, 10.1257/aer.20161500.

Miller, Douglas L. (2023) "An Introductory Guide to Event Study Models," *Journal of Economic Perspectives*, 37 (2), 203–30, 10.1257/jep.37.2.203.

Monks, James (2004) "An Empirical Examination of the Impact of College Financial Aid on Family Savings," *National Tax Journal*, 57 (2.1), 189–207, 10.17310/ntj.2004.2.03.

Morrison, William and Dmitry Taubinsky (2023) "Rules of Thumb and Attention Elasticities: Evidence from Under- and Overreaction to Taxes," *The Review of Economics and Statistics*, 105 (5), 1110–1127, 10.1162/rest_a_01126.

National Center for Education Statistics (2024) "Digest of Education Statistics: 2023. Table 303.50. Total fall enrollment in degree-granting postsecondary institutions, by attendance status, age of student, and level and control of institution: 2021.," January, U.S. Department of Education. https://nces.ed.gov/programs/digest/d23/tables/dt23_303.50.asp.

Park, Rina Seung Eun and Judith Scott-Clayton (2018) "The Impact of Pell Grant Eligibility on Community College Students' Financial Aid Packages, Labor Supply, and Academic Outcomes," *Educational Evaluation and Policy Analysis*, 40 (4), 557–585, 10.3102/0162373718783868.

Rees-Jones, Alex and Dmitry Taubinsky (2018) "Taxing Humans: Pitfalls of the Mechanism Design Approach and Potential Resolutions," *Tax Policy and the Economy*, 32 (1), 107–133, 10.1086/697139.

——— (2020) "Measuring "Schmeduling"," *The Review of Economic Studies*, 87 (5), 2399–2438, 10.1093/restud/rdz045.

Saez, Emmanuel, Joel Slemrod, and Seth H. Giertz (2012) "The Elasticity of Taxable Income with Respect to Marginal Tax Rates: A Critical Review," *Journal of Economic Literature*, 50 (1), 3–50, 10.1257/jel.50.1.3.

Slemrod, Joel (2006) "The Role of Misconceptions in Support for Regressive Tax Reform," *National Tax Journal*, 59 (1), 57–75, 10.17310/ntj.2006.1.03.

Stantcheva, Stefanie (2021) "Understanding Tax Policy: How do People Reason?*," *The Quarterly Journal of Economics*, 136 (4), 2309–2369, 10.1093/qje/qjab033.

Taubinsky, Dmitry and Alex Rees-Jones (2018) "Attention Variation and Welfare: Theory and Evidence from a Tax Salience Experiment," *The Review of Economic Studies*, 85 (4), 2462–2496, 10.1093/restud/rdx069.

Wiswall, Matthew and Basit Zafar (2015) "Determinants of College Major Choice: Identification using an Information Experiment," *The Review of Economic Studies*, 82 (2), 791–824, 10.1093/restud/rdu044.

Supplemental Appendix for "Measuring the Welfare Cost of a Misperceived Tax: Evidence from College Financial Aid"

by Nick Gebbia, Joseph Gray-Hancuch, and Paul R. Organ

Appendix Table of Contents

Supplementary Appendix A **Additional Figures and Tables** **3**

Supplementary Appendix B **Proofs** **23**

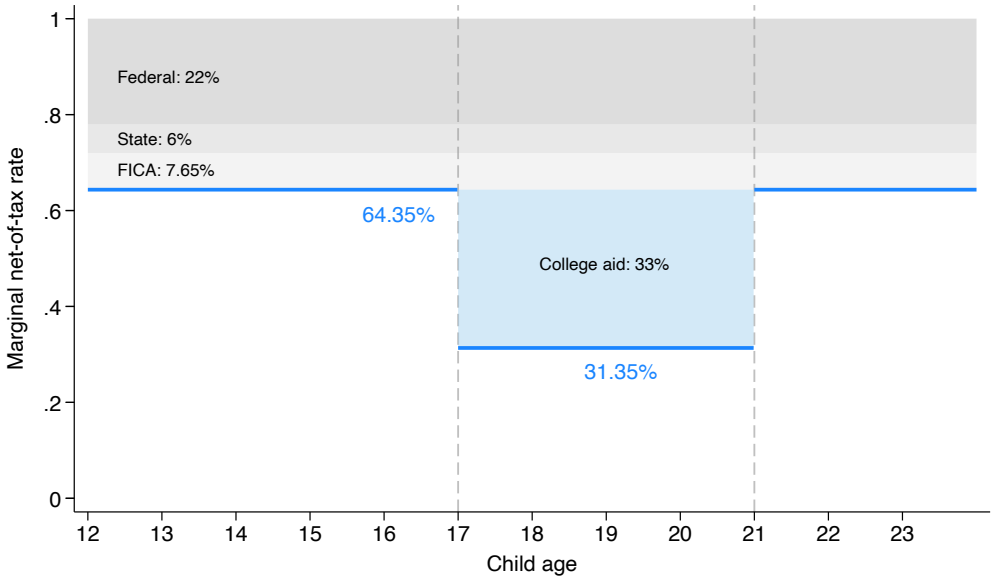
Supplementary Appendix C **Why DWL Increases with the Variance of Tax Rate Misperceptions** **37**

Supplementary Appendix D **Prolific Survey Description** **39**

Supplementary Appendix A Additional Figures and Tables

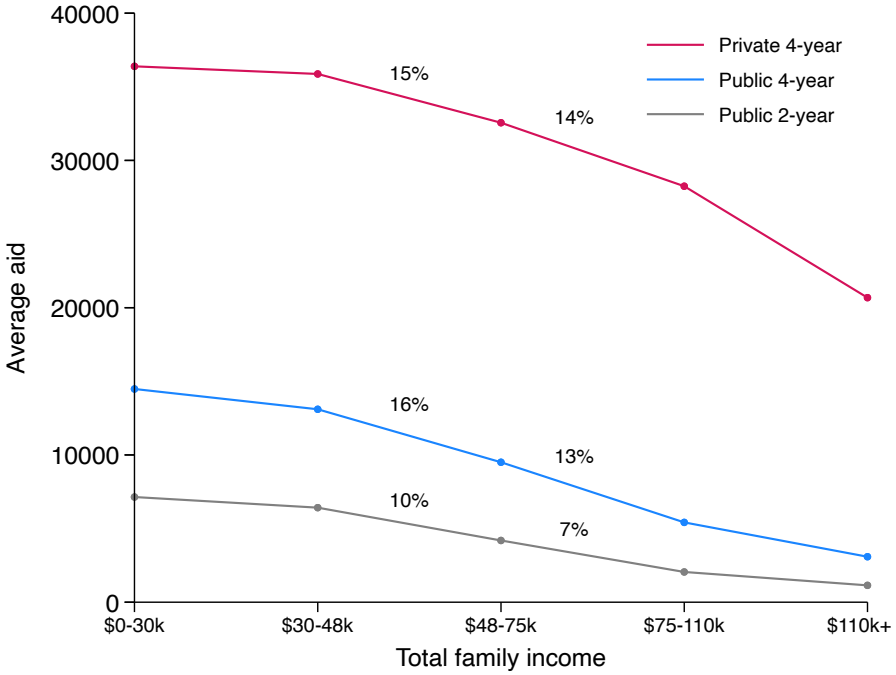
Figure A.1: Marginal Net-of-tax Rate for $125k Income Family around College

UC Berkeley, CA resident



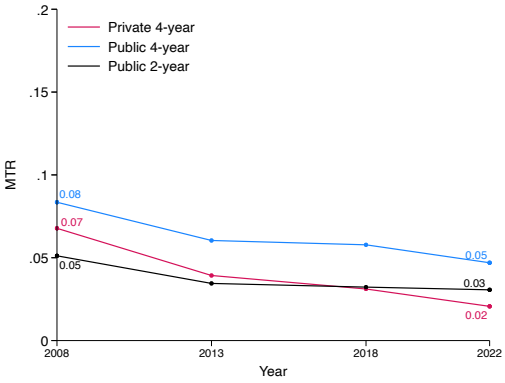
Notes: This figure plots the all-in marginal net-of-tax rate on income faced by a family with income $125k who live in California and enroll a child at UC Berkeley, around the years in which parent income is used to determine financial aid for the child. The marginal net-of-tax rate accounts for federal and state income taxes, FICA, and the implicit tax from college aid.

Figure A.2: Average Aid Schedules across the US

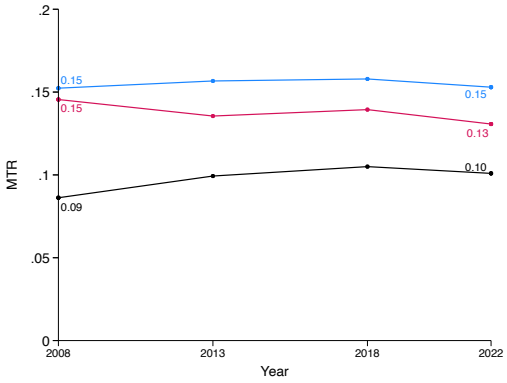


Notes: This figure uses 2019-20 IPEDS data to plot annual average grant and scholarship aid in each of five income bins, separately among private four-year colleges, public four-year colleges, and public two-year colleges. Averages are weighted by total fall enrollment of full-time, first-time, first-year, degree-seeking undergraduates. Estimated marginal implicit tax rates are listed in the figure, which are computed as the difference in average aid across adjacent income bins divided by the difference in the midpoints of each income bin.

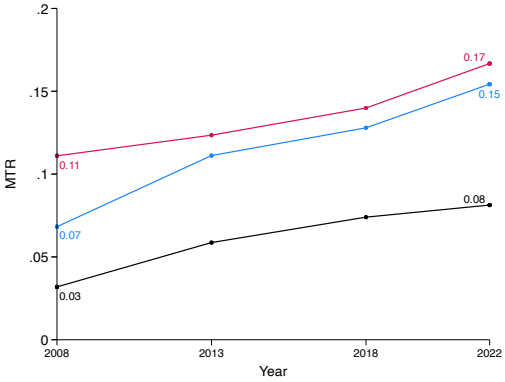
Figure A.3: Trends in the College Aid Implicit Tax



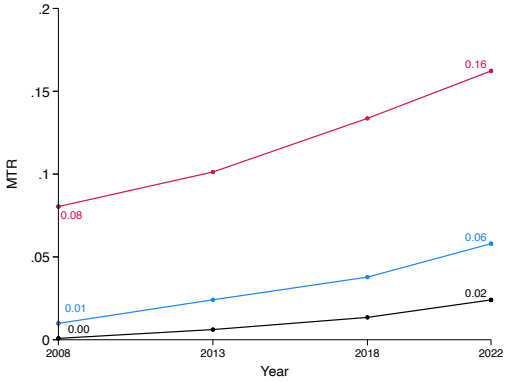
(A) \$0-30k to \$30-48k



(B) \$30-48k to \$48-75k



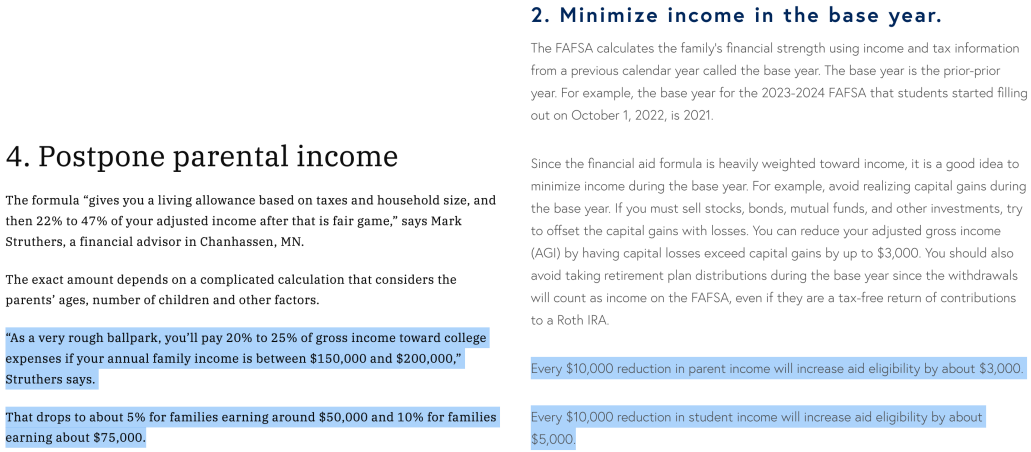
(C) \$48-75k to \$75-110k



(D) \$75-110k to \$110k+

Notes: This figure shows estimated marginal implicit tax rates from college aid phaseout based on IPEDS data. For colleges across the US, IPEDS annually records the average financial aid for students with total family income in each of five income bins: \$0-30k, \$30-48k, \$48-75k, \$75-110k, and above \$110k. Using this data, we construct estimates of marginal implicit tax rates from college aid for adjacent income bins by dividing the change in average aid by the difference in midpoints of each income range. For the top bin of \$110k and above, we use \$150k, which will influence the levels of the estimated tax rates in Panel D, but will not influence relative trends. We then compute averages of these tax rates by college type and year. We refer to each academic year based on the fall term, e.g. 2008 refers to the 2008-09 academic year. Averages are weighted by total fall enrollment of full-time, first-time, first-year, degree-seeking undergraduates.

Figure A.4: News Descriptions of the College Aid Implicit Tax



4. Postpone parental income

The formula "gives you a living allowance based on taxes and household size, and then 22% to 47% of your adjusted income after that is fair game," says Mark Struthers, a financial advisor in Chanhassen, MN.

The exact amount depends on a complicated calculation that considers the parents' ages, number of children and other factors.

"As a very rough ballpark, you'll pay 20% to 25% of gross income toward college expenses if your annual family income is between $150,000 and $200,000," Struthers says.

That drops to about 5% for families earning around $50,000 and 10% for families earning about $75,000.

2. Minimize income in the base year.

The FAFSA calculates the family's financial strength using income and tax information from a previous calendar year called the base year. The base year is the prior-prior year. For example, the base year for the 2023-2024 FAFSA that students started filling out on October 1, 2022, is 2021.

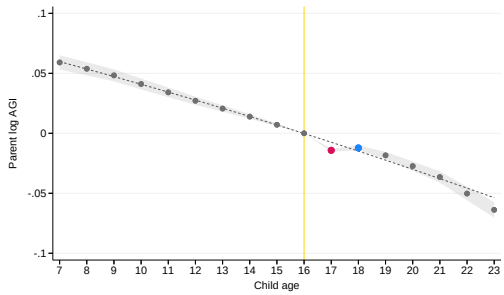
Since the financial aid formula is heavily weighted toward income, it is a good idea to minimize income during the base year. For example, avoid realizing capital gains during the base year. If you must sell stocks, bonds, mutual funds, and other investments, try to offset the capital gains with losses. You can reduce your adjusted gross income (AGI) by having capital losses exceed capital gains by up to $3,000. You should also avoid taking retirement plan distributions during the base year since the withdrawals will count as income on the FAFSA, even if they are a tax-free return of contributions to a Roth IRA.

Every $10,000 reduction in parent income will increase aid eligibility by about $3,000.

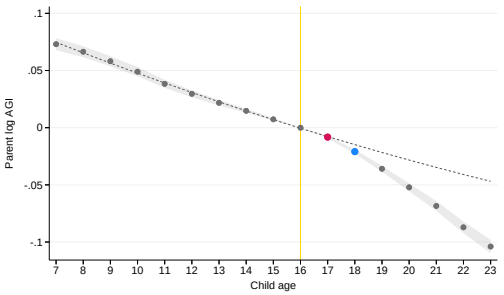
Every $10,000 reduction in student income will increase aid eligibility by about $5,000.

Notes: This figure shows excerpts from two popular news articles describing the college aid implicit tax and the incentive for parents to reduce income when it affects a child's college aid. Both excerpts emphasize that parents should reduce income when it affects aid. When describing the extent of the relationship between parent income and college aid, one article refers to average net college cost as a function of parent income, while the other refers to roughly the maximum marginal relationship of a $3,000 aid reduction for each additional $10,000 of income. Sources: Money, "FAFSA Tips: These 7 Moves Could Help You Score More Financial Aid" (https://money.com/more-financial-aid-money-fafsa/) and College Ave, "File the FAFSA to Get More Money in College" (https://www.collegeavestudentloans.com/blog/file-the-fafsa-to-get-more-money-in-college).

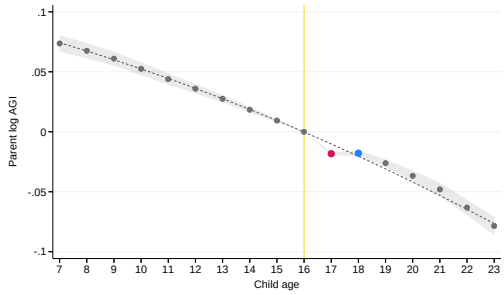
Figure A.5: Parent Income Responses to College Incentives – Levels



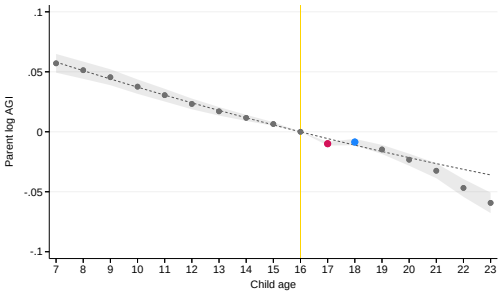
(A) 4-year college



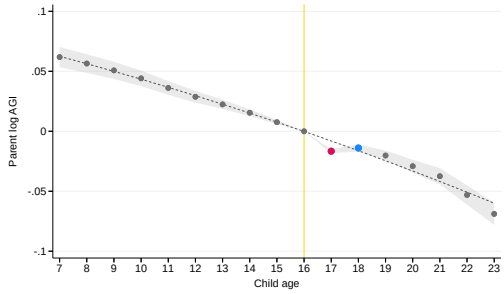
(B) No college



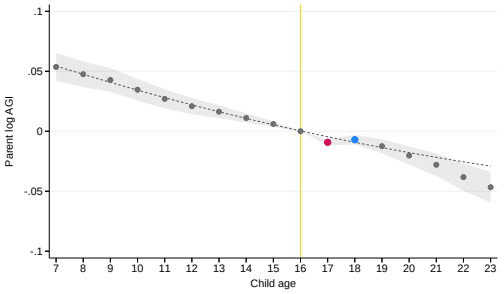
(C) High aid phaseout college



(D) Low aid phaseout college



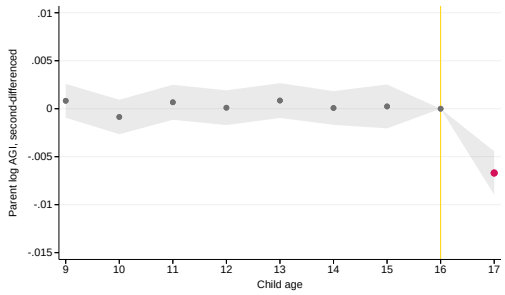
(E) High cost college



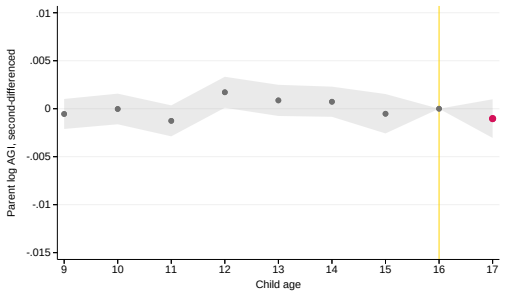
(F) Low cost college

Notes: This figure shows estimates of child age effects from the nonparametric event study regression in Equation (1), where the outcome is log AGI. Because our variation comes from parents who are different ages when they had their child, trends arise if families with different income trajectories have children at different ages. We overlay a quadratic trend fitted to the age 8 through 16 effects. Differences from this trend reflect our parametric event study estimates that control for the trend. All panels restrict to households that have baseline income averaged over 2005 and 2006 between \$50k to \$150k, filed taxes with positive AGI for the years 2001 through 2022, claimed the relevant child as a dependent when the child was age 7 through 17, and whose child turned 18 between 2012 and 2016. Each panel shows estimates for a separate subsample based on the child's college attendance at age 18: Panel (A) is a four-year in-state public college or four-year private college that has IPEDS data on average aid by family income bin; Panel (B) is no college; Panel (C) is for children who receive at least one 1098-T at age 18 from a college with an aid phaseout rate above 10%, defined by the change in average aid between families in income bins \$30-48k and \$75-110k divided by the difference in the midpoints of each range; Panel (D) is for children who receive a 1098-T at age 18, but any 1098-T they receive is from a college with an aid phaseout rate below 10%; Panel (E) is for children who receive at least one 1098-T at age 18 from a college with a total Cost of Attendance above \$30k; Panel (F) is for children who receive a 1098-T at age 18, but any 1098-T they receive is from a college with total Cost of Attendance below \$20k. Standard errors are clustered at the tax unit level.

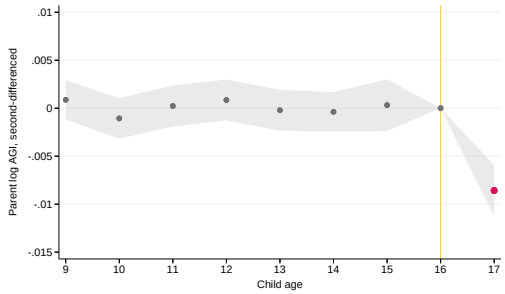
Figure A.6: Parent Income Responses to College Incentives – Second-differences



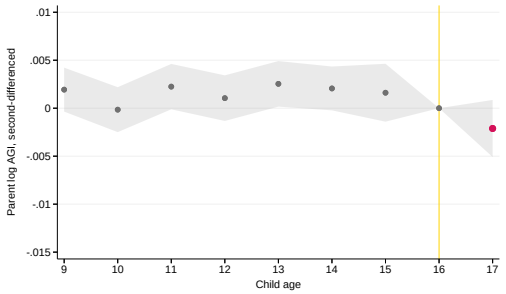
(A) 4-year college



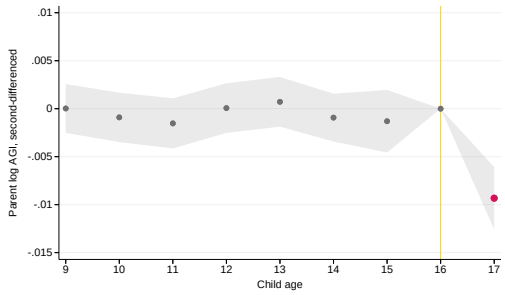
(B) No college



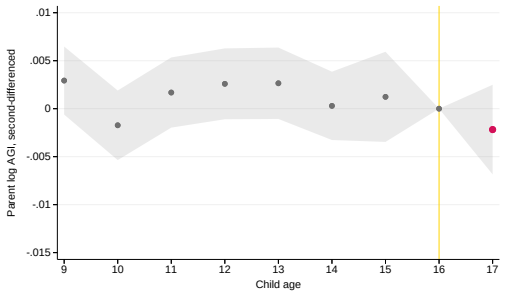
(C) High aid phaseout college



(D) Low aid phaseout college



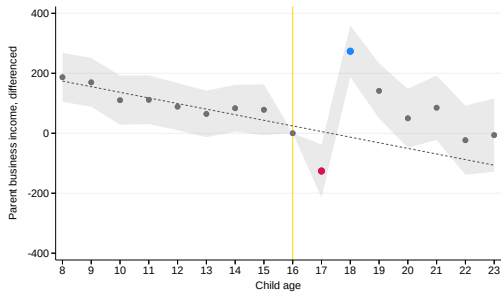
(E) High cost college



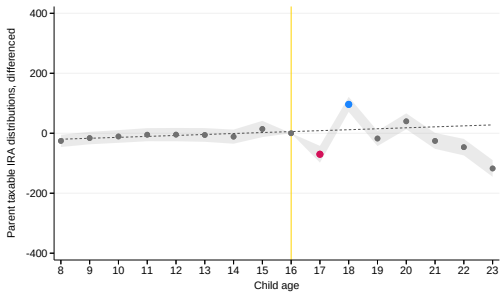
(F) Low cost college

Notes: This figure shows estimates of child age effects from the nonparametric event study regression in Equation (1), where the outcome is second-differenced log AGI. Because our variation comes from parents who are different ages when they had their child, trends arise if families with different income trajectories have children at different ages. All panels restrict to households that have baseline income averaged over 2005 and 2006 between \$50k to \$150k, filed taxes with positive AGI for the years 2001 through 2022, claimed the relevant child as a dependent when the child was age 7 through 17, and whose child turned 18 between 2012 and 2016. Each panel shows estimates for a separate subsample based on the child's college attendance at age 18: Panel (A) is a four-year in-state public college or four-year private college that has IPEDS data on average aid by family income bin; Panel (B) is no college; Panel (C) is for children who receive at least one 1098-T at age 18 from a college with an aid phaseout rate above 10%, defined by the change in average aid between families in income bins \$30-48k and \$75-110k divided by the difference in the midpoints of each range; Panel (D) is for children who receive a 1098-T at age 18, but any 1098-T they receive is from a college with an aid phaseout rate below 10%; Panel (E) is for children who receive at least one 1098-T at age 18 from a college with a total Cost of Attendance above \$30k; Panel (F) is for children who receive a 1098-T at age 18, but any 1098-T they receive is from a college with total Cost of Attendance below \$20k. Standard errors are clustered at the tax unit level.

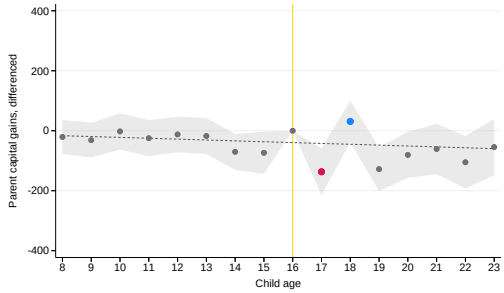
Figure A.7: Which Types of Income Parents Adjust Non-wage Components



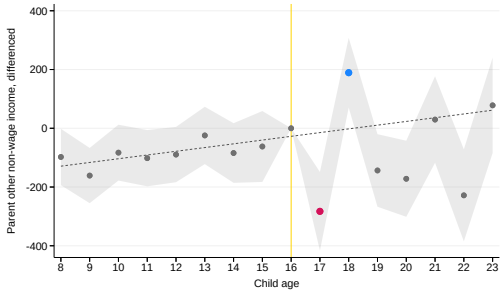
(A) Business income



(B) Taxable IRA distributions



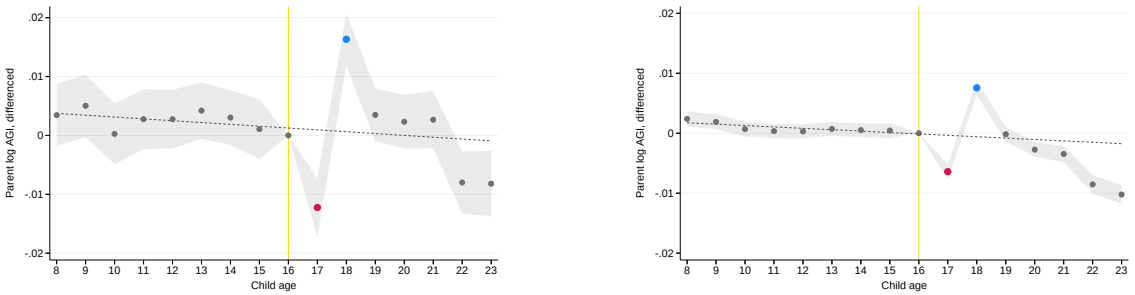
(C) Capital gains



(D) Other non-wage

Notes: This figure shows estimates of child age effects from the nonparametric event study regression in Equation (1). The outcome in Panel (A) is first-differenced business income, defined as the sum of profit or loss from an S-corporation, partnership, or sole proprietorship, in Panel (B) is first-differenced taxable IRA distributions, in Panel (C) is first-differenced capital gains (Schedule D), and in Panel (D) is the first difference of AGI minus W-2 Box 1 minus the sum of business income, taxable IRA distributions, and capital gains. Because our variation comes from parents who are different ages when they had their child, trends arise if families with different income trajectories have children at different ages. We overlay a linear trend fitted to the age 8 through 16 effects. Differences from this trend reflect our parametric event study estimates that control for the trend. All panels restrict to households that have baseline income averaged over 2005 and 2006 between \$50k to \$150k, filed taxes with positive AGI for the years 2001 through 2022, claimed the relevant child as a dependent when the child was age 7 through 17, whose child turned 18 between 2012 and 2016, and whose child attended a four-year in-state public college or four-year private college that has IPEDS data on average aid by family income bin at age 18. Standard errors are clustered at the tax unit level.

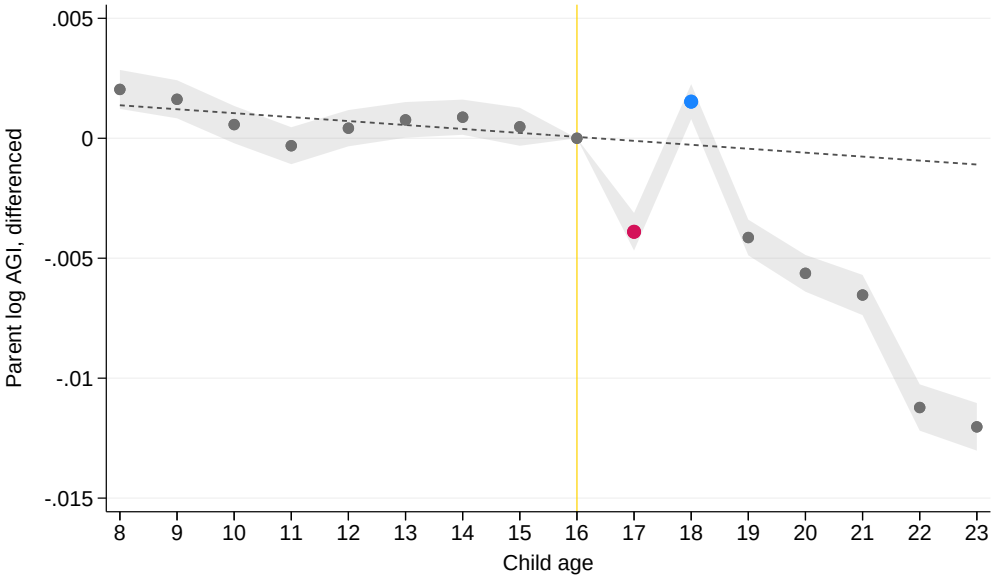
Figure A.8: Parent Income Responses to College Incentives
By Non-wage Income Share



(A) High non-wage (B) Low non-wage

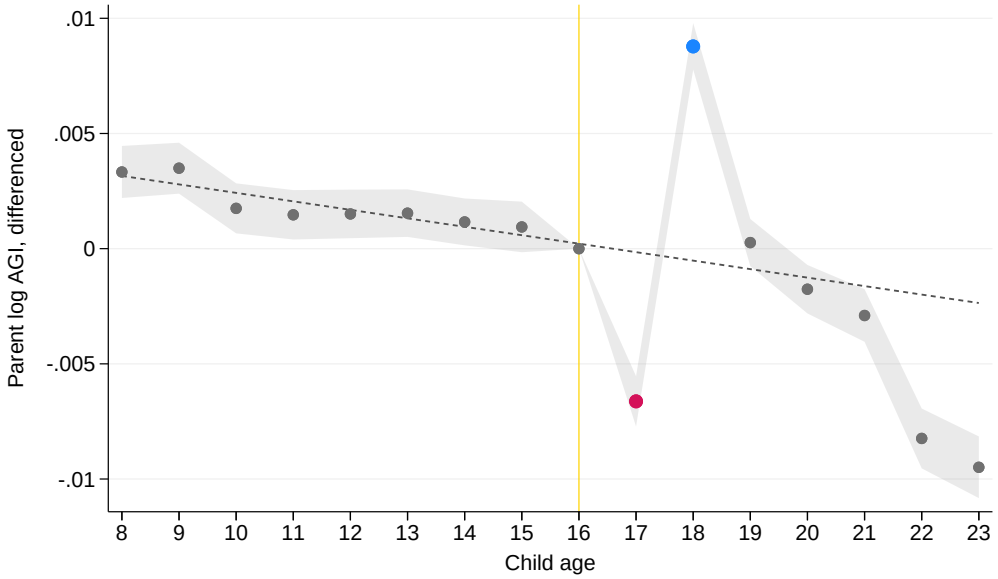
Notes: This figure shows estimates of child age effects from the nonparametric event study regression in Equation (1), where the outcome is first-differenced log AGI. Because our variation comes from parents who are different ages when they had their child, trends arise if families with different income trajectories have children at different ages. We overlay a linear trend fitted to the age 8 through 16 effects. Differences from this trend reflect our parametric event study estimates that control for the trend. All panels restrict to households that have baseline income averaged over 2005 and 2006 between \$50k to \$150k, filed taxes with positive AGI for the years 2001 through 2022, claimed the relevant child as a dependent when the child was age 7 through 17, whose child turned 18 between 2012 and 2016, and whose child attended a four-year in-state public college or four-year private college that has IPEDS data on average aid by family income bin at age 18. Panel (A) restricts to households whose W-2 income accounts for less than 80% of their AGI on average over 2005 and 2006, and Panel (B) restricts to households whose W-2 income accounts for more than 80% of their AGI on average over 2005 and 2006. Standard errors are clustered at the tax unit level.

Figure A.9: Effects without Conditioning on College Status



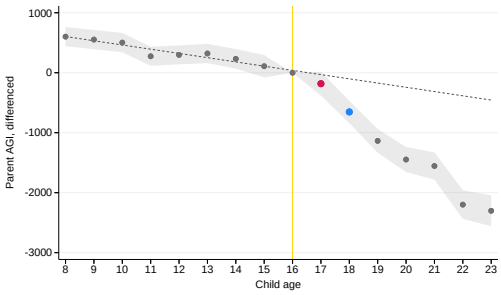
Notes: This figure shows estimates of child age effects from the nonparametric event study regression in Equation (1), where the outcome is first-differenced log AGI. Because our variation comes from parents who are different ages when they had their child, trends arise if families with different income trajectories have children at different ages. We overlay a linear trend fitted to the age 8 through 16 effects. Differences from this trend reflect our parametric event study estimates that control for the trend. The sample restricts to households that have baseline income averaged over 2005 and 2006 between \$50k to \$150k, filed taxes with positive AGI for the years 2001 through 2022, claimed the relevant child as a dependent when the child was age 7 through 17, and whose child turned 18 between 2012 and 2016. Standard errors are clustered at the tax unit level.

Figure A.10: Effects among Any-College Families

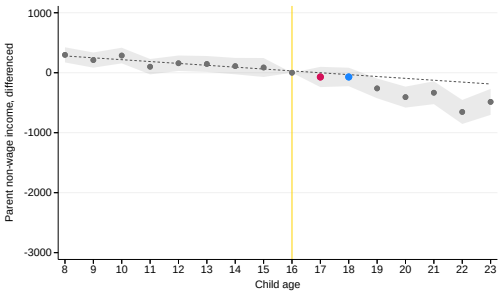


Notes: This figure shows estimates of child age effects from the nonparametric event study regression in Equation (1), where the outcome is first-differenced log AGI. Because our variation comes from parents who are different ages when they had their child, trends arise if families with different income trajectories have children at different ages. We overlay a linear trend fitted to the age 8 through 16 effects. Differences from this trend reflect our parametric event study estimates that control for the trend. The sample restricts to households that have baseline income averaged over 2005 and 2006 between $50k to $150k, filed taxes with positive AGI for the years 2001 through 2022, claimed the relevant child as a dependent when the child was age 7 through 17, whose child turned 18 between 2012 and 2016, and whose child received a 1098-T from any college at age 18. Standard errors are clustered at the tax unit level.

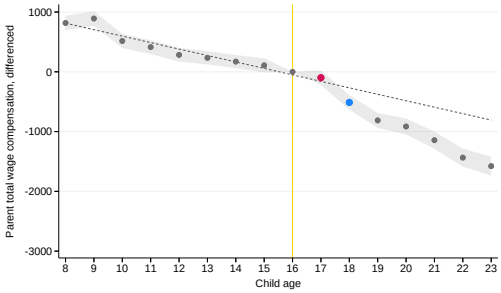
Figure A.11: Responses by Income Type – Non-college Families



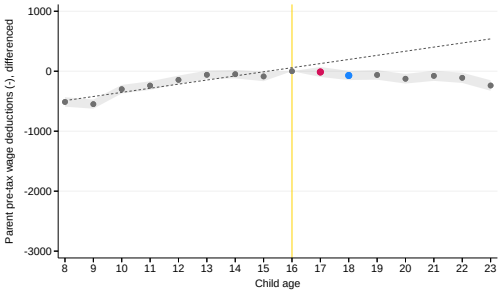
(A) Total income



(B) Non-wage income



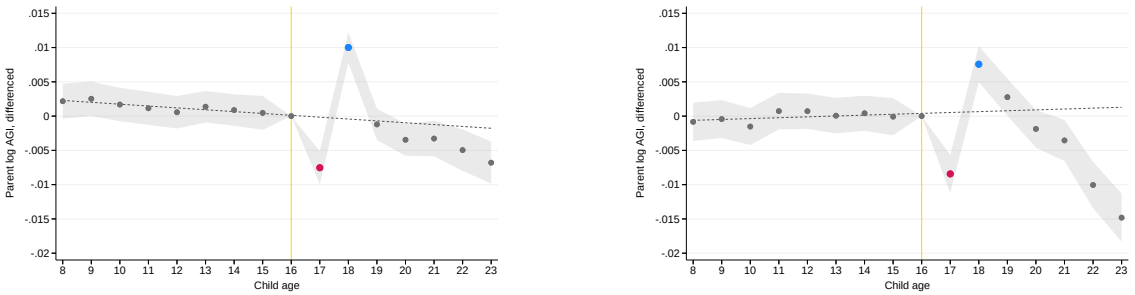
(C) Total wage compensation



(D) Pre-tax wage deductions (-)

Notes: This figure shows estimates of child age effects from the nonparametric event study regression in Equation (1). The outcome in Panel (A) is first-differenced AGI, in Panel (B) is the first difference of AGI minus W-2 Box 1, in Panel (C) is the first difference of W-2 Box 5, and in Panel (D) is the first difference of W-2 Box 1 minus Box 5. Because our variation comes from parents who are different ages when they had their child, trends arise if families with different income trajectories have children at different ages. We overlay a linear trend fitted to the age 8 through 16 effects. Differences from this trend reflect our parametric event study estimates that control for the trend. All panels restrict to households that have baseline income averaged over 2005 and 2006 between \$50k to \$150k, filed taxes with positive AGI for the years 2001 through 2022, claimed the relevant child as a dependent when the child was age 7 through 17, whose child turned 18 between 2012 and 2016, and whose child received no 1098-T at age 18. Standard errors are clustered at the tax unit level.

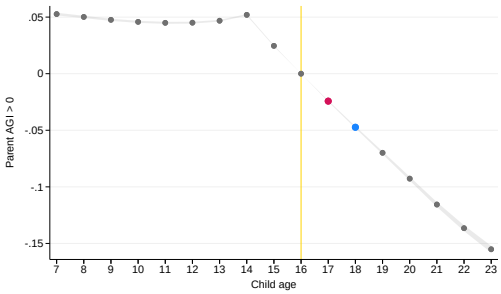
Figure A.12: Parent Income Responses to College Incentives
By Younger Sibling Presence



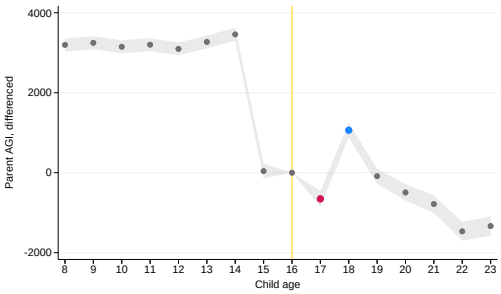
(A) Has younger sibling (B) No younger sibling

Notes: This figure shows estimates of child age effects from the nonparametric event study regression in Equation (1), where the outcome is first-differenced log AGI. Because our variation comes from parents who are different ages when they had their child, trends arise if families with different income trajectories have children at different ages. We overlay a linear trend fitted to the age 8 through 16 effects. Differences from this trend reflect our parametric event study estimates that control for the trend. All panels restrict to households that have baseline income averaged over 2005 and 2006 between \$50k to \$150k, filed taxes with positive AGI for the years 2001 through 2022, claimed the relevant child as a dependent when the child was age 7 through 17, whose child turned 18 between 2012 and 2016, and whose child attended a four-year in-state public college or four-year private college that has IPEDS data on average aid by family income bin at age 18. Panel (A) restricts to households that claim an age 1-14 dependent when the focal child is age 18, and Panel (B) restricts to households that claim no younger dependent when the focal child is age 18. Standard errors are clustered at the tax unit level.

Figure A.13: Effects Remain Apparent without Filing or Claiming Restrictions



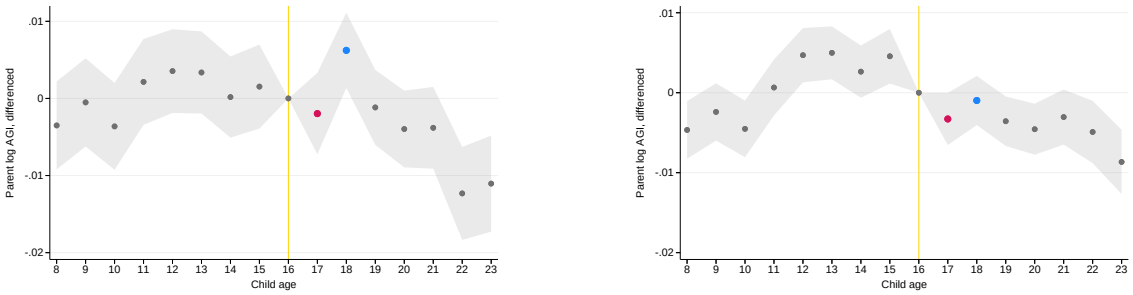
(A) Files with positive AGI



(B) AGI, first differenced

Notes: This figure shows estimates of child age effects from the nonparametric event study regression in Equation (1). The outcome in Panel (A) is an indicator for filing with positive AGI, and in Panel (B) is first-differenced AGI. All panels restrict to households that have baseline income averaged over 2005 and 2006 between $50k to $150k, whose child turned 18 between 2012 and 2016, and whose child attended a four-year in-state public college or four-year private college that has IPEDS data on average aid by family income bin at age 18. Standard errors are clustered at the tax unit level.

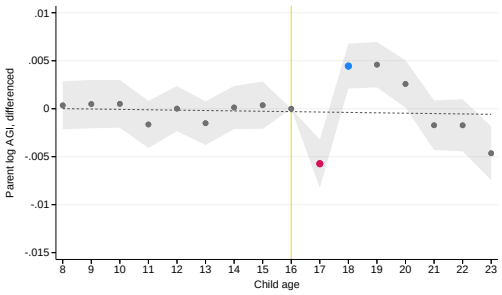
Figure A.14: Effects among Low-income Families



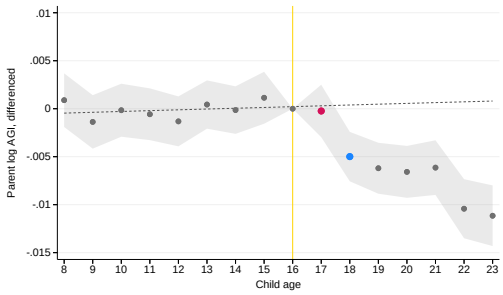
(A) Any college (B) No college

Notes: This figure shows estimates of child age effects from the nonparametric event study regression in Equation (1), where the outcome is first-differenced log AGI. Because our variation comes from parents who are different ages when they had their child, trends arise if families with different income trajectories have children at different ages. We do not overlay a linear trend in this figure due to the apparent nonlinear pretrends. All panels restrict to households that have baseline income averaged over 2009 and 2010 below $25k, no taxable interest income in 2009 and 2010, filed taxes with positive AGI for the years 2001 through 2022, claimed the relevant child as a dependent when the child was age 7 through 17, and whose child turned 18 between 2012 and 2016. Each panel shows estimates for a separate subsample based on the child's college attendance at age 18: Panel (A) is any college; Panel (B) is no college. Standard errors are clustered at the tax unit level.

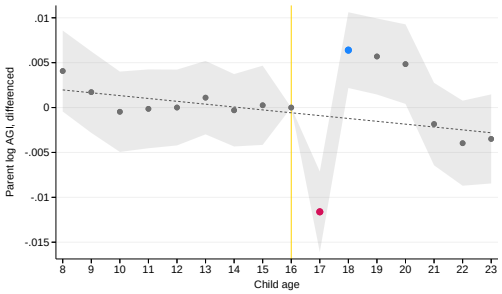
Figure A.15: Effects among High-income Families



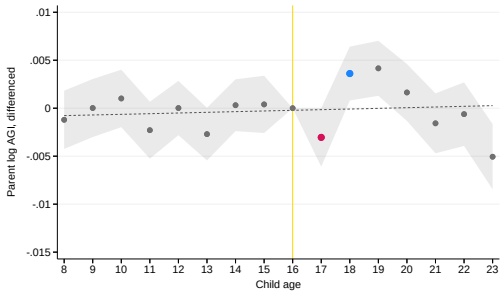
(A) Any college



(B) No college



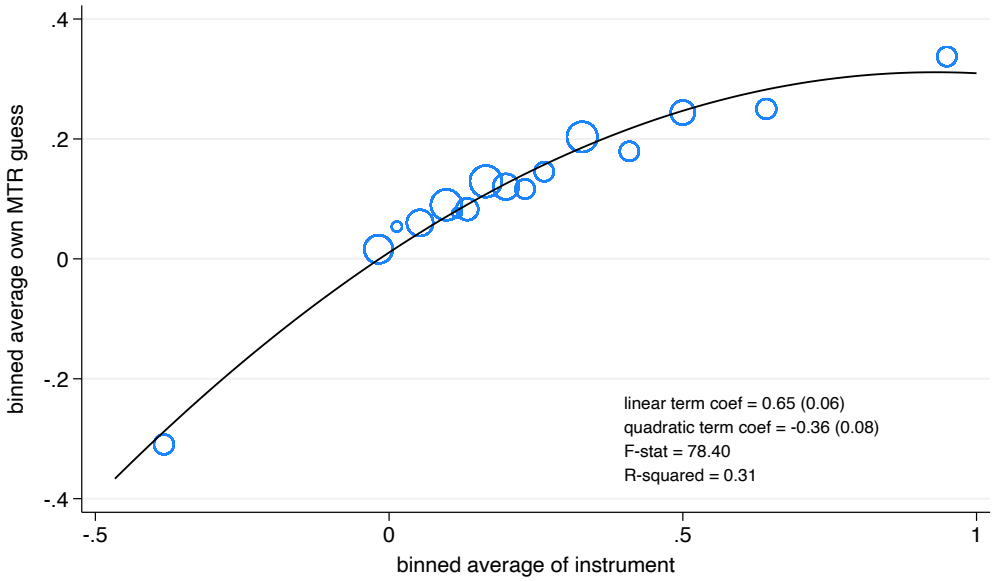
(C) Cost $>$ \$50k



(D) Cost $<$ \$50k

Notes: This figure shows estimates of child age effects from the nonparametric event study regression in Equation (1), where the outcome is first-differenced log AGI. Because our variation comes from parents who are different ages when they had their child, trends arise if families with different income trajectories have children at different ages. We overlay a linear trend fitted to the age 8 through 16 effects. Differences from this trend reflect our parametric event study estimates that control for the trend. All panels restrict to households that have baseline income averaged over 2009 and 2010 between \$150k and \$300k, filed taxes with positive AGI for the years 2001 through 2022, claimed the relevant child as a dependent when the child was age 7 through 17, and whose child turned 18 between 2012 and 2016. Each panel shows estimates for a separate subsample based on the child's college attendance at age 18: Panel (A) is any college; Panel (B) is no college; Panel (C) is for children who receive at least one 1098-T at age 18 from a college with a total Cost of Attendance above \$50k; and Panel (D) is for children who receive a 1098-T at age 18, but any 1098-T they receive is from a college with total Cost of Attendance below \$50k. Standard errors are clustered at the tax unit level.

Figure A.16: First Stage Relationship between Survey Instrument and Focal Tax Perception



Notes: This figure shows the variation in survey respondent perceptions of their own MTR from college aid phaseout (the focal rate) that is explained by variation in reported perceptions of the secondary rate (the instrument). For respondents with previous year income between \$50-75k, the focal rate is the difference in the respondent's reported perception of average aid for an in-state student at a four-year public college in the respondent's state between family income bins \$30-48k and \$48-75k, divided by the difference in midpoints of these bins. For respondents with previous year income between \$75-150k, the focal rate is the difference in the respondent's reported perception of average aid for an in-state student at a four-year public college in the respondent's state between family income bins \$48-75k and \$75-110k, divided by the difference in midpoints of these bins. For all respondent's, the instrument is the difference in their reported perception of average aid for an in-state student at a four-year public college in the respondent's state between students with family income \$45k or \$75k, divided by the \$30k difference between these points. For 20 bins of the survey instrument, the figure shows the average value of the respondent's own MTR guess against the average value of the instrument, along with a quadratic fit. The figure notes results from regressing the respondent's own MTR guess on the survey instrument at the individual level, with robust standard errors in parentheses. All values are weighted to match the US college population based on state and family income range using IPEDS data.

Table A.1: Tax Data Descriptive Statistics

	Mean	Std. dev.	Median	# of households
Household characteristics				
AGI (2005-06 average)	87,490	25,698	83,100	1,284,879
W-2 income (2005-06 average)	80,057	29,087	77,200	1,284,879
Primary filer age when child is 17	48	6	47	1,284,838
% any college at age 18	51%			1,284,879
% Broad aid college at age 18	43%			1,284,879
% 4-year aid college at age 18	37%			1,284,879
College characteristics				
Broad aid sample				
MTR from college aid	12%	16%	12%	557,849
Cost of attendance	34,046	16,401	27,400	557,849
Net price	20,104	9,001	19,600	557,849
4-year aid sample				
MTR from college aid	13%	16%	13%	475,384
Cost of attendance	36,853	16,117	29,400	475,384
Net price	21,967	8,365	20,700	475,384

Notes: This table shows descriptive statistics for the tax data analysis sample. The sample restricts to households that have baseline income averaged over 2005 and 2006 between \$50k to \$150k, filed taxes with positive AGI for the years 2001 through 2022, claimed the relevant child as a dependent when the child was age 7 through 17, and whose child turned 18 between 2012 and 2016. Median values for AGI, W-2 wages, cost of attendance, and net price are rounded to the nearest \$100 for disclosure purposes.

Table A.2: ETI due to the College Aid Implicit Tax – Quadratic Trend

	By college type		By non-wage income, among 4-year aid sample	
	Broad aid sample	4-year aid sample	Low non-wage	High non-wage
Reduced form	-0.0068	-0.0073	-0.0069	-0.0118
	(0.0006)	(0.0007)	(0.0007)	(0.0027)
First stage	-0.1478	-0.1570	-0.1566	-0.1598
	(0.0003)	(0.0003)	(0.0003)	(0.0008)
ETI (2SLS)	0.0461	0.0467	0.0442	0.0741
	(0.0043)	(0.0044)	(0.0044)	(0.0170)
N	11,119,582	9,477,378	8,165,556	1,311,822

Notes: This table shows estimates of age 17 effects from a version of the parametric event study regression in Equation (2), which also includes a quadratic trend in child age. The first row shows the reduced form effect on the outcome of first-differenced log AGI. The second row shows the first stage effect on the first difference of the marginal net of tax rate, which is one minus the marginal tax rate from college aid phaseout. The third row shows the ETI, which is the ratio of the reduced form to the first stage. All panels restrict to households that have baseline income averaged over 2005 and 2006 between \$50k to \$150k, filed taxes with positive AGI for the years 2001 through 2022, claimed the relevant child as a dependent when the child was age 7 through 17, and whose child turned 18 between 2012 and 2016. The first column also restricts to households whose child attends a "4-year aid" college at age 18, and the second column restricts to households whose child attends a "Broad aid" college at age 18. Among households whose child attends a "4-year aid" college at age 18, the third column restricts to households whose W-2 income accounts for more than 80% of their AGI on average over 2005 and 2006, and the fourth column restricts to households whose W-2 income accounts for less than 80% of their AGI on average over 2005 and 2006. Standard errors are clustered at the tax unit level.

Table A.3: Survey Descriptive Statistics

	Mean	SD	Min	Max
Total income	$98,380	$31,727	$50,000	$150,000
Nonlabor income share >= 20%	0.05	0.22	0	1
Respondent BA	0.63	0.5	0	1
Respondent age	40.3	8.8	20	74
Own aid MTR, true	0.14	0.04	0.05	0.24
Own aid MTR, guess	0.11	0.23	-0.81	0.65
Ratio of MTR guess to true	0.89	1.64	-4.87	5.23
N	1,279			

Notes: This table shows descriptive statistics for the main analysis sample from the Prolific survey measuring parent perceptions of the college aid implicit tax. All values are weighted to match the US college population based on state and family income range using IPEDS data.

Table A.4: Misperception Moments – by Subgroups

	Main sample	Ages 14-22 High school and college age		Ages 17-22 College age		Above median survey time	
	Estimates	Estimates	p-value: diff vs. (1)	Estimates	p-value: diff vs. (1)	Estimates	p-value: diff vs. (1)
	(1)	(2)		(3)		(4)	
Elasticity conversion ratio	1.01	1.04	0.72	1.03	0.84	1.12	0.13
	(0.05)	(0.08)		(0.13)		(0.10)	
Bias	-0.21	-0.23	0.87	-0.29	0.65	-0.40	0.02
	(0.09)	(0.14)		(0.18)		(0.11)	
Variance	1.81	1.50	0.17	1.30	0.11	2.35	0.01
	(0.20)	(0.26)		(0.32)		(0.33)	
Cov with income (thousands)	10.26	1.27	0.27	2.36	0.43	26.56	0.03
	(6.86)	(9.73)		(9.31)		(11.89)	
Cov with elasticity (Non-labor income share >= 20%)	0.11	0.10	0.90	-0.01	0.23	0.26	0.10
	(0.08)	(0.13)		(0.08)		(0.18)	
Cov with tax rate	-0.013	-0.011	0.50	-0.006	0.11	-0.018	0.08
	(0.002)	(0.004)		(0.005)		(0.004)	
N	1,279	467		243		640	

Notes: For the main survey analysis sample and several subgroups, this table shows estimates of the components of Equation (3) involving tax rate perceptions, which includes the moments involving individual misperception θ_i and the conversion ratio that translates between the reduced form ETI and the structural elasticity as shown in Equation (4). We define Bias as $E[\theta_i]^2 - 1$. All estimates incorporate the survey instrument and bias correction. Standard errors are bootstrapped with 1,000 repetitions. All values are weighted to match the US college population based on state and family income range using IPEDS data.

Supplementary Appendix B Proofs

B.1 A Model of Deadweight Loss under Tax Rate Misperception and Proof of Equation (3)

In the model, agents value consumption and they dislike producing income. There are two periods, with period 0 representing years in which agents face only a baseline tax and period 1 representing years in which agents face the baseline tax plus an additional temporary tax. Each agent chooses how much income to earn in period 0 and in period 1, and their consumption is then determined by their budget constraint. Agents can misperceive the period 1 tax rate in a heterogeneous manner.

Specifically, each agent solves the following two-period dynamic income problem:

$$\max_{C_i,\, Z_{i,0},\, Z_{i,1}} U \;=\; C_i \;-\; \sum_{t=0}^{1} \frac{a_i}{1+\frac{1}{e(X_i)}} \left(\frac{Z_{i,t}}{a_i}\right)^{1+\frac{1}{e(X_i)}} \tag{B.5}$$

$$\text{s.t.} \quad C_i \;=\; (1-\tau_i^b)Z_{i,0} + [(1-(\tau_i^b+\tilde{\tau}_i)]Z_{i,1}$$

where C_i is agent i's total consumption; $Z_{i,0}$ is income earned in period 0; $Z_{i,1}$ is income earned in period 1; a_i is ability; e is an elasticity parameter that is allowed to vary based on observables in the vector X_i; τ_i^b is the baseline tax rate the agent faces in both periods; and $\tilde{\tau}_i$ is the additional tax rate the agent *perceives* to apply to their period 1 income. The true additional tax rate the agent faces in period 1 is τ_i. For simplicity, we assume utility is quasilinear in consumption, and as a result agents are indifferent to the allocation of consumption across the two periods, which emphasizes that the problem is one of dynamic income choice rather than of consumption allocation.

We assume that each agent makes their income decisions *as if* $\tilde{\tau}_i$ is the true additional tax rate they face. If an agent misperceives this tax rate, so that $\tilde{\tau}_i \neq \tau_i$, we assume this is accounted for when the true budget constraint – which replaces $\tilde{\tau}_i$ with τ_i – binds, so that the agent's actual consumption might differ from their chosen consumption.

Taking the first order conditions and solving yields the solution for chosen income in each period as:

$$\begin{aligned} Z_{i,0} \;&=\; a_i(1-\tau_i^b)^{e(X_i)} \\ Z_{i,1} \;&=\; a_i[1-(\tau_i^b+\tilde{\tau}_i)]^{e(X_i)} \end{aligned} \tag{B.6}$$

In period 0, when individuals face only the baseline tax, each agent produces income that is lower than their ability a_i and depends on the baseline net-of-tax rate $(1-\tau_i^b)$. In period 1, when the temporary tax is added, each agent produces income that is lower than their ability and depends on the total perceived net-of-tax rate $[1-(\tau_i^b+\tilde{\tau}_i)]$. $e(X_i)$ is the elasticity of earned income with respect to the total perceived net-of-tax rate for an individual with observables X_i. Below, we refer to e as the structural elasticity as it is the parameter that determines preferences over consumption and earning income.

To assess DWL, assume all tax revenue collected from each agent is returned to the agent as a lump sum transfer. To simplify exposition, define the function for the

disutility of producing income $v_i(z) \coloneqq \frac{a_i}{1+\frac{1}{e(X_i)}}\left(\frac{z}{a_i}\right)^{1+\frac{1}{e(X_i)}}$. Then the welfare (utility) of agent i is:

$$W_i \;=\; \left\{(1-\tau_i^b)Z_{i,0} + [1-(\tau_i^b+\tau_i)]Z_{i,1} - v_i(Z_{i,0}) - v_i(Z_{i,1})\right\} + \underbrace{\tau_i^b Z_{i,0} + (\tau_i^b+\tau_i)Z_{i,1}}_{\text{Lump sum transfer}} \tag{B.7}$$

We are interested in the change in welfare that results from introducing a tax τ_i relative to a baseline of no such tax, or our object of interest is $\Delta W_i \equiv W_i(\tau_i) - W_i(0)$ where we write $W_i(\cdot)$ to represent agent i's utility as a function of the period 1 added tax τ_i. As is standard in the public economics literature, we derive an expression for ΔW_i by approximating $W_i(\tau_i)$ using a second order Taylor expansion around a tax of zero:

$$W_i(\tau_i) \;\approx\; W_i(0) \;+\; \frac{dW_i}{d\tau_i}\Big|_{\tau_i=0}\tau_i \;+\; \frac{1}{2}\frac{d^2W_i}{d\tau_i^2}\Big|_{\tau_i=0}\tau_i^2$$

which implies:

$$\begin{aligned}\Delta W_i \;&=\; W_i(\tau_i) \;-\; W_i(0) \\ &\approx\; \frac{dW_i}{d\tau_i}\Big|_{\tau_i=0}\tau_i \;+\; \frac{1}{2}\frac{d^2W_i}{d\tau_i^2}\Big|_{\tau_i=0}\tau_i^2\end{aligned}$$

To simplify further, we must find expressions for the first and second derivatives of W_i with respect to τ_i. Taking the first derivative of Equation (B.7) with respect to τ_i yields:

$$\frac{dW_i}{d\tau_i} \;=\; [1-(\tau_i^b+\tau_i)]\frac{dZ_{i,1}}{d\tau_i} \;-\; Z_{i,1} \;-\; v_i'(Z_{i,1})\frac{dZ_{i,1}}{d\tau_i} \;+\; (\tau_i^b+\tau_i)\frac{dZ_{i,1}}{d\tau_i} \;+\; Z_{i,1}$$

which simplifies to:

$$\frac{dW_i}{d\tau_i} \;=\; \frac{dZ_{i,1}}{d\tau_i}(\tau_i^b+\tilde{\tau}_i)$$

where we make use of the first order conditions for Equation (B.5): $v_i'(Z_{i,0}) = (1-\tau_i^b)$ and $v_i'(Z_{i,1}) = [1-(\tau_i^b+\tilde{\tau}_i)]$.

The second derivative of welfare with respect to the tax is:

$$\frac{d^2W_i}{d\tau_i^2} \;=\; \frac{dZ_{i,1}}{d\tau_i}\frac{d\tilde{\tau}_i}{d\tau_i} \;+\; \frac{d^2Z_{i,1}}{d\tau_i^2}(\tau_i^b+\tilde{\tau}_i)$$

Plugging in the expressions for the first and second derivatives of W_i with respect to τ_i yields:

$$\begin{aligned}\Delta W_i \;&\approx\; \frac{dW_i}{d\tau_i}\Big|_{\tau_i=0}\tau_i \;+\; \frac{1}{2}\frac{d^2W_i}{d\tau_i^2}\Big|_{\tau_i=0}\tau_i^2 \\ &=\; \left[\frac{dZ_{i,1}}{d\tau_i}(\tau_i^b+\tilde{\tau}_i)\right]\Big|_{\tau_i=0}\tau_i \;+\; \frac{1}{2}\left[\frac{dZ_{i,1}}{d\tau_i}\frac{d\tilde{\tau}_i}{d\tau_i} \;+\; \frac{d^2Z_{i,1}}{d\tau_i^2}(\tau_i^b+\tilde{\tau}_i)\right]\Big|_{\tau_i=0}\tau_i^2 \\ &=\; \frac{1}{2}\left[\frac{dZ_{i,1}}{d\tau_i}\frac{d\tilde{\tau}_i}{d\tau_i}\right]\Big|_{\tau_i=0}\tau_i^2\end{aligned}$$

where the terms multiplied by $(\tau_i^b + \tilde{\tau}_i)$ drop out because (1) we assume the baseline tax τ_i^b is zero and (2) we evaluate the Taylor expansion beginning from $\tau_i = 0$ and we assume there is no misperception when the true tax is zero (i.e. $\tilde{\tau}_i\big|_{\tau_i=0} = 0$). In fact, individuals face a nonzero marginal tax on income at baseline due to federal and state income taxes along with incentives created by means testing in other social programs. Assuming a baseline tax of zero here focuses the analysis on DWL that is produced directly from the implicit tax in college aid regardless of the structure of other tax programs, while omitting the portion of DWL that depends on the interaction between college aid incentives and existing programs. In this regard, our empirical DWL estimates omit the portion of DWL that is due to the interaction with pre-existing distortions, and as such our estimates provide a lower bound on the full DWL from the college aid implicit tax that includes interactions with pre-existing distortions.

Now we define the individual misperception parameter θ_i such that $\tilde{\tau}_i = \theta_i \tau_i$, which captures the degree to which an individual over- or under-estimates their own marginal tax rate. Then we can rewrite the preceding expression for the individual change in welfare as[28]:

$$\Delta W_i \;\approx\; \frac{1}{2}\frac{dZ_{i,1}}{d\tilde{\tau}_i}\Big|_{\tau_i=0} \tau_i^2 \theta_i^2$$

We can then incorporate the elasticity $e(X_i)$ as:

$$\begin{aligned}
\Delta W_i \;&\approx\; \frac{1}{2}\frac{dZ_{i,1}}{d\tilde{\tau}_i}\Big|_{\tau_i=0} \tau_i^2 \theta_i^2 \\
&= -\frac{1}{2}\left[Z_{i,1} \frac{1}{1-(\tau_i^b + \tilde{\tau}_i)} \right] \Big|_{\tau_i=0} e(X_i)\tau_i^2 \theta_i^2 \\
&= -\frac{1}{2} Z_{i,1}\Big|_{\tau_i=0} e(X_i)\tau_i^2 \theta_i^2
\end{aligned}$$

where, as above, the last step uses the fact that $(\tau_i^b + \tilde{\tau}_i)\big|_{\tau_i=0} = 0$ under our assumptions. $Z_{i,1}$ appears in the final expression above as a weight on the elasticity $e(X_i)$, which accounts for the fact that, conditional on the elasticity $e(X_i)$, an individual with a higher baseline income $Z_{i,1}$ will reduce income by a greater amount. From here, we will suppress the notation indicating that $Z_{i,1}$ is to be evaluated at a tax of $\tau_i = 0$, but this remains implicit.

The total deadweight loss in a population with N individuals is then:

$$DWL \;=\; N \cdot E[-\Delta W_i] \;\approx\; N \cdot \frac{1}{2} \cdot E\left[Z_{i,1} e(X_i) \tau_i^2 \theta_i^2 \right] \tag{B.8}$$

Equation (B.8) demonstrates that welfare can be computed using information on income, the structural elasticity of income with respect to a change in the perceived tax, the true tax, and misperception. In order to make sense of Equation (B.8) and to facilitate identification using survey data measuring perceptions, we rearrange terms

[28] Note θ_i enters this equation as $\frac{d\tilde{\tau}_i}{d\tau_i}\big|_{\tau_i=0} = \theta_i$.

to produce the following:

$$DWL \approx N \cdot \underbrace{\frac{1}{2} \cdot \Big\{ E\left[Z_{i,1} e(X_i) \tau_i^2\right]}_{\substack{\text{Correct Perceptions} \\ \text{(standard Harberger triangle)}}} \cdot \Big(E[\theta_i]^2 + Var(\theta_i) \Big) \; + \; E\left[e(X_i)\right] E\left[\tau_i^2\right] \cdot Cov\left(\theta_i^2, Z_{i,1}\right)$$

$$+ \; E\left[Z_{i,1}\right] E\left[\tau_i^2\right] \cdot Cov\left(\theta_i^2, e(X_i)\right) \; + \; E\left[Z_{i,1}\right] E\left[e(X_i)\right] \cdot Cov\left(\theta_i^2, \tau_i^2\right) \Big\}$$

which is the key result of Equation (3). This last step relies on the assumption that third and fourth central moments are zero, which rules out a higher order dependence structure. Specifically, the assumptions are that $E\left[\bar{Z}_{i,1} \overline{e(X_i)} \bar{\theta}_i^2\right] = 0$, $E\left[\bar{Z}_{i,1} \bar{\tau}_i^2 \bar{\theta}_i^2\right] = 0$, $E\left[\overline{e(X_i)} \bar{\tau}_i^2 \bar{\theta}_i^2\right] = 0$, and $E\left[\bar{Z}_{i,1} \overline{e(X_i)} \bar{\tau}_i^2 \bar{\theta}_i^2\right] = 0$, where here bars are used to denote deviations from the expectation, or generally $\bar{X} \equiv X - E(X)$.

B.2 ETI vs Structural Elasticity and Proof of Equation (4)

We will distinguish between the reduced form elasticity and the structural elasticity. The reduced form elasticity corresponds to the standard approach for estimating the ETI in the public economics literature, which measures the change in income with respect to the true change in the marginal tax rate. The structural elasticity instead measures the change in income with respect to the perceived change in the marginal tax rate. The structural elasticity is solely a function of the parameters of the utility function, but the reduced form elasticity depends additionally on the extent to which individuals misperceive the tax rate they face. Under the standard assumption of perfect perceptions, the reduced form elasticity will equal the structural elasticity; however, in the presence of misperception, the two are not necessarily equal. This distinction is essential for identifying DWL in the presence of misperception, as Equation (3) expresses DWL as a function of the structural elasticity $e(X_i)$ which is not identified by standard ETI estimation techniques when there is misperception. Here, we clarify the difference between the reduced form elasticity and the structural elasticity, and we demonstrate how to identify the structural elasticity – a necessary input to identify DWL – by combining the reduced form elasticity with a measure of perceptions.

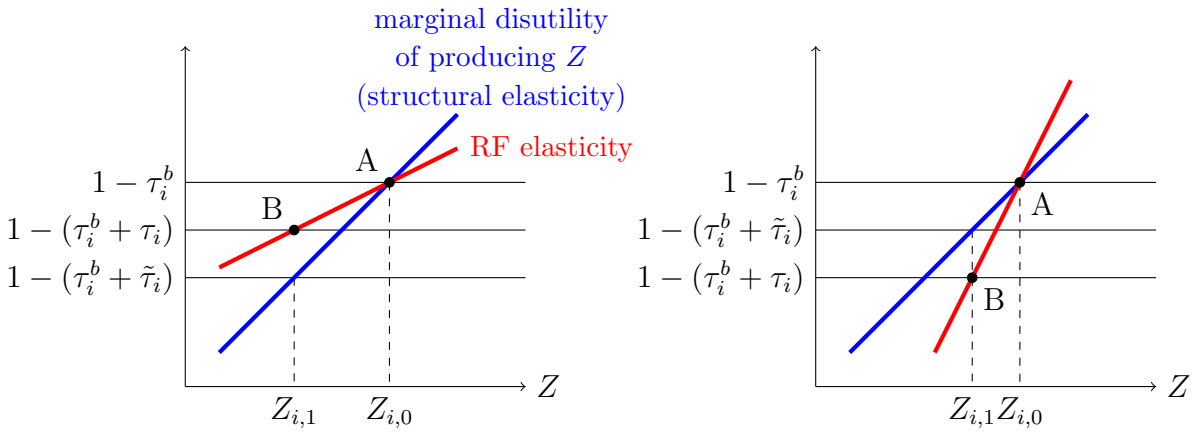
Figure B.1 illustrates this distinction. The figure considers how income changes between period 0 and period 1 in the preceding model. The individual dislikes producing income at an increasing marginal disutility, but likes consumption.[29] $Z_{i,0}$ is income in period 0, which occurs where the marginal disutility of earning income equals the marginal utility of consumption, normalized to the net-of-tax rate in period 0, $(1 - \tau_i^b)$. In period 1, if the individual perceives her additional tax rate to be $\tilde{\tau} > 0$, she will choose a lower income level that equates her marginal disutility of earnings with her total perceived net-of-tax rate $[1 - (\tau_i^b + \tilde{\tau}_i)]$. This lower income level is labeled as $Z_{i,1}$ in the figure, and the blue curve traces out how income will vary with the perceived net-of-tax rate. The slope of this curve depends on the individual's utility function, and in particular on the rate at which the marginal disutility of earning declines as income falls. This is the structural elasticity.

The econometrician typically doesn't observe the perceived tax rate, but only the actual rate, τ_i^b and τ_i, along with income in each period. In other words, the econometrician observes the points labeled "A" and "B" in the data. The ETI that the econometrician will estimate (the "reduced form ETI") is the change in log earnings divided by the change in the log net-of-tax rate between the two periods, which extrapolates the line through points A and B. The left panel of Figure 7 shows a case where $\tilde{\tau} > \tau$. Here, the econometrician will estimate a larger ETI than is implied by the utility function parameters alone. The right panel shows a different case, where the structural elasticity is the same but now $\tilde{\tau} < \tau$. Here, the econometrician's reduced form elasticity will be smaller than the structural elasticity.

The two panels in Figure B.1 provide intuition to support the following result: the relationship by which the structural elasticity and the perceived tax rate determine the reduced form elasticity can be inverted, so that knowledge of the reduced form elasticity and the perceived tax rate is sufficient to identify the structural elasticity. We now demonstrate this result formally.

[29] Figure B.1 assumes utility is quasilinear in consumption, so that the marginal utility of consumption curves are flat.

Figure B.1: Reduced Form ETI vs. Structural Elasticity



(A) $\tilde{\tau}_i > \tau_i$ (B) $\tilde{\tau}_i < \tau_i$

Notes: This figure details the relationship between the reduced form ETI and the structural elasticity under tax rate misperception. In the left panel, the perceived tax rate is higher than the true tax rate. The blue curve is the structural relationship between income earned and the marginal disutility of earning income, and this curve along with the perceived tax determines income decisions. The econometrician observes data points A and B, and uses these to fit the reduced form relationship between income and the actual tax rate and thereby estimate the reduced form elasticity. In this case, the reduced form elasticity is larger than the underlying structural elasticity. The right panel shows the opposite case, where the perceive tax is lower than the true tax rate. In this case, the reduced form elasticity is smaller than the underlying structural elasticity.

In most empirical studies, the reduced form ETI estimand is an instrumental variables ratio of an average policy effect on log income scaled by the first stage average effect of the policy on the log net-of-tax rate. In such cases, the reduced form elasticity can be expressed following the notation in Figure B.1 as:

$$\beta^{ETI}(X_i) \coloneqq \frac{E[\ln(Z_{i,1}) - \ln(Z_{i,0}) \,|\, X_i]}{E[\ln(1-(\tau_i^b+\tau_i)) - \ln(1-\tau_i^b) \,|\, X_i]}$$

where we allow for heterogeneity by observables in vector X_i.

We can use the solution of the model in Appendix B.1 as written in Equation (B.6) to plug in for $Z_{i,1}$ and $Z_{i,0}$ to express the reduced form ETI β^{ETI} as a function of the structural elasticity e, conditional on observables X_i:

$$\begin{aligned}
\beta^{ETI}(X_i) &\coloneqq \frac{E[\ln(Z_{i,1}) - \ln(Z_{i,0}) \,|\, X_i]}{E[\ln(1-(\tau_i^b+\tau_i)) - \ln(1-\tau_i^b) \,|\, X_i]} \\
&= \frac{E[\ln(a_i[1-(\tau_i^b+\tilde{\tau}_i)]^{e(X_i)}) - \ln(a_i(1-\tau_i^b)^{e(X_i)}) \,|\, X_i]}{E[\ln(1-(\tau_i^b+\tau_i)) - \ln(1-\tau_i^b) \,|\, X_i]} \\
&= e(X_i) \cdot \frac{E[\ln(1-(\tau_i^b+\tilde{\tau}_i)) - \ln(1-\tau_i^b) \,|\, X_i]}{E[\ln(1-(\tau_i^b+\tau_i)) - \ln(1-\tau_i^b) \,|\, X_i]}
\end{aligned}$$

As a result, we can also solve for the structural elasticity as a function of the reduced form elasticity and a measure of both the perceived and actual tax:

$$e(X_i) = \beta^{ETI}(X_i) \cdot \frac{E[\ln(1-(\tau_i^b+\tau_i)) - \ln(1-\tau_i^b) \,|\, X_i]}{E[\ln(1-(\tau_i^b+\tilde{\tau}_i)) - \ln(1-\tau_i^b) \,|\, X_i]}$$

Under the assumption that the baseline tax is zero, which we will maintain in subsequent sections, this simplifies to:

$$e(X_i) \;=\; \beta^{ETI}(X_i) \cdot \frac{E[\ln(1-\tau_i) \,|\, X_i]}{E[\ln(1-\tilde{\tau}_i) \,|\, X_i]}$$

which is the result shown in Equation (4).

B.3 Proof of Identification Results with Survey-Based Measures of Misperception

1. $E[\ln(1-\tilde{\tau}_i)|X_i] \approx E[\ln(1-\hat{\tilde{\tau}}_{1i})|X_i]$

Using a first order Taylor approximation of $\ln(1-\hat{\tilde{\tau}}_i)$ starting from $\epsilon_{1i}=0$, we have:

$$\begin{aligned}
E[\ln(1-\hat{\tilde{\tau}}_i)|X_i] &= E[\ln(1-(\tilde{\tau}_i+\epsilon_{1i}))|X_i] \\
&\approx E\left[\ln(1-\tilde{\tau}_i) + \epsilon_{1i}\frac{d}{d\epsilon_{1i}}[\ln(1-(\tilde{\tau}_i+\epsilon_{1i}))]|_{\epsilon_{1i}=0}|X_i\right] \\
&= E\left[\ln(1-\tilde{\tau}_i) - \frac{\epsilon_{1i}}{1-\tilde{\tau}_i}|X_i\right] \\
&= E[\ln(1-\tilde{\tau}_i)|X_i] - E\left[\frac{1}{1-\tilde{\tau}_i}\underbrace{E\left[\epsilon_{1i}|X_i,\tilde{\tau}_i\right]}_{=0,\text{ by assumption}}\Big|X_i\right] \\
&= E[\ln(1-\tilde{\tau}_i)|X_i]
\end{aligned}$$

2. $E[\theta_i] = E[\hat{\theta}_{1i}]$

$$\begin{aligned}
E[\hat{\theta}_{1i}] &= E\left[\frac{\hat{\tilde{\tau}}_{1i}}{\tau_i}\right] \\
&= E\left[\frac{\tilde{\tau}_i+\epsilon_{1i}}{\tau_i}\right] \\
&= E\left[\frac{\tilde{\tau}_i}{\tau_i}\right] + \underbrace{E\left[\frac{\epsilon_{1i}}{\tau_i}\right]}_{=0,\text{ by } E[\epsilon_{1i}|\tau_i]=0} \\
&= E[\theta_i]
\end{aligned}$$

3. $Var(\theta_i) = Cov(\hat{\theta}_{1i},\hat{\theta}_{2i}) - Cov\left(\frac{E[\hat{\tilde{\tau}}_{1i}|\tau_i]}{\tau_i}, \frac{E[\hat{\tilde{\tau}}_{2i}-\hat{\tilde{\tau}}_{1i}|\tau_i]}{\tau_i}\right)$

First we derive an expression for $Cov(\hat{\theta}_{1i},\hat{\theta}_{2i})$:

$$\begin{aligned}
Cov(\hat{\theta}_{1i},\hat{\theta}_{2i}) &= Cov\left(\frac{\hat{\tilde{\tau}}_{1i}}{\tau_i}, \frac{\hat{\tilde{\tau}}_{2i}}{\tau_i}\right) \\
&= Cov\left(\frac{\tilde{\tau}_i+\epsilon_{1i}}{\tau_i}, \frac{\tilde{\tau}_i+\epsilon_{2i}}{\tau_i}\right) \\
&= Cov\left(\frac{\tilde{\tau}_i}{\tau_i}, \frac{\tilde{\tau}_i}{\tau_i}\right) + Cov\left(\frac{\tilde{\tau}_i}{\tau_i}, \frac{\epsilon_{1i}}{\tau_i}\right) + Cov\left(\frac{\tilde{\tau}_i}{\tau_i}, \frac{\epsilon_{2i}}{\tau_i}\right) + Cov\left(\frac{\epsilon_{1i}}{\tau_i}, \frac{\epsilon_{2i}}{\tau_i}\right) \\
&= Var(\theta_i) + Cov\left(\frac{\tilde{\tau}_i}{\tau_i}, \frac{\epsilon_{1i}}{\tau_i}\right) + Cov\left(\frac{\tilde{\tau}_i}{\tau_i}, \frac{\epsilon_{2i}}{\tau_i}\right) + Cov\left(\frac{\epsilon_{1i}}{\tau_i}, \frac{\epsilon_{2i}}{\tau_i}\right)
\end{aligned}$$

Then we simplify the three covariance terms immediately above:

$$\begin{aligned} Cov\left(\frac{\tilde{\tau}_i}{\tau_i}, \frac{\epsilon_{1i}}{\tau_i}\right) &= E\underbrace{\left[Cov\left(\frac{\tilde{\tau}_i}{\tau_i}, \frac{\epsilon_{1i}}{\tau_i}\middle|\tau_i\right)\right]}_{=0, \text{ by } Cov(\tilde{\tau}_i, \epsilon_{1i}|\tau_i)=0} + Cov\left(E\left[\frac{\tilde{\tau}_i}{\tau_i}\middle|\tau_i\right], \underbrace{E\left[\frac{\epsilon_{1i}}{\tau_i}\middle|\tau_i\right]}_{=0, \text{ by } E[\epsilon_{1i}|\tau_i]=0}\right) \\ &= 0 \end{aligned}$$

And,

$$\begin{aligned} Cov\left(\frac{\tilde{\tau}_i}{\tau_i}, \frac{\epsilon_{2i}}{\tau_i}\right) &= E\underbrace{\left[Cov\left(\frac{\tilde{\tau}_i}{\tau_i}, \frac{\epsilon_{2i}}{\tau_i}\middle|\tau_i\right)\right]}_{=0, \text{ by } Cov(\tilde{\tau}_i, \epsilon_{2i}|\tau_i)=0} + Cov\left(E\left[\frac{\tilde{\tau}_i}{\tau_i}\middle|\tau_i\right], E\left[\frac{\epsilon_{2i}}{\tau_i}\middle|\tau_i\right]\right) \\ &= Cov\left(\frac{E[\tilde{\tau}_i|\tau_i]}{\tau_i}, \frac{E[\epsilon_{2i}|\tau_i]}{\tau_i}\right) \end{aligned}$$

And,

$$\begin{aligned} Cov\left(\frac{\epsilon_{1i}}{\tau_i}, \frac{\epsilon_{2i}}{\tau_i}\right) &= E\underbrace{\left[Cov\left(\frac{\epsilon_{1i}}{\tau_i}, \frac{\epsilon_{2i}}{\tau_i}\middle|\tau_i\right)\right]}_{=0, \text{ by } Cov(\epsilon_{1i}, \epsilon_{2i}|\tau_i)=0} + Cov\left(\underbrace{E\left[\frac{\epsilon_{1i}}{\tau_i}\middle|\tau_i\right]}_{=0, \text{ by } E[\epsilon_{1i}|\tau_i]=0}, E\left[\frac{\epsilon_{2i}}{\tau_i}\middle|\tau_i\right]\right) \\ &= 0 \end{aligned}$$

This leaves us with

$$Cov(\hat{\theta}_{1i}, \hat{\theta}_{2i}) = Var(\theta_i) + Cov\left(\frac{E[\tilde{\tau}_i|\tau_i]}{\tau_i}, \frac{E[\epsilon_{2i}|\tau_i]}{\tau_i}\right)$$

We now show that we can identify the nuisance term as $Cov\left(\frac{E[\hat{\tilde{\tau}}_{1i}|\tau_i]}{\tau_i}, \frac{E[\hat{\tilde{\tau}}_{2i} - \hat{\tilde{\tau}}_{1i}|\tau_i]}{\tau_i}\right) = Cov\left(\frac{E[\tilde{\tau}_i|\tau_i]}{\tau_i}, \frac{E[\epsilon_{2i}|\tau_i]}{\tau_i}\right)$:

$$\begin{aligned} Cov\left(\frac{E[\hat{\tilde{\tau}}_{1i}|\tau_i]}{\tau_i}, \frac{E[\hat{\tilde{\tau}}_{2i} - \hat{\tilde{\tau}}_{1i}|\tau_i]}{\tau_i}\right) &= Cov\left(\frac{\overbrace{E[\tilde{\tau}_i + \epsilon_{1i}|\tau_i]}^{=E[\tilde{\tau}_i|\tau_i], \text{ by } E[\epsilon_{1i}|\tau_i]=0}}{\tau_i}, \frac{E[(\tilde{\tau}_i + \epsilon_{2i}) - (\tilde{\tau}_i + \epsilon_{1i})|\tau_i]}{\tau_i}\right) \\ &= Cov\left(\frac{E[\tilde{\tau}_i|\tau_i]}{\tau_i}, \frac{E[\epsilon_{2i}|\tau_i] - \overbrace{E[\epsilon_{1i}|\tau_i]}^{=0, \text{ by } E[\epsilon_{1i}|\tau_i=0]}}{\tau_i}\right) \\ &= Cov\left(\frac{E[\tilde{\tau}_i|\tau_i]}{\tau_i}, \frac{E[\epsilon_{2i}|\tau_i]}{\tau_i}\right) \end{aligned}$$

Therefore, combining the above we have:

$$\begin{aligned} Cov(\hat{\theta}_{1i}, \hat{\theta}_{2i}) - Cov\left(\frac{E[\hat{\tilde{\tau}}_{1i}|\tau_i]}{\tau_i}, \frac{E[\hat{\tilde{\tau}}_{2i} - \hat{\tilde{\tau}}_{1i}|\tau_i]}{\tau_i}\right) &= Var(\theta_i) + Cov\left(\frac{E[\tilde{\tau}_i|\tau_i]}{\tau_i}, \frac{E[\epsilon_{2i}|\tau_i]}{\tau_i}\right) \\ &\quad - Cov\left(\frac{E[\tilde{\tau}_i|\tau_i]}{\tau_i}, \frac{E[\epsilon_{2i}|\tau_i]}{\tau_i}\right) \\ &= Var(\theta_i) \end{aligned}$$

4. $Cov(\theta_i^2, Z_{i,1}) = Cov(\hat{\theta}_{1i}\hat{\theta}_{2i}, Z_{i,1}) - Cov\left(\frac{E[\hat{\tilde{\tau}}_{2i} - \hat{\tilde{\tau}}_{1i}|\tau_i]}{\tau_i^2}\hat{\tilde{\tau}}_{1i}, Z_{i,1}\right)$

First we derive an expression for $Cov(\hat{\theta}_{1i}\hat{\theta}_{2i}, Z_{i,1})$:

$$\begin{aligned}
Cov(\hat{\theta}_{1i}\hat{\theta}_{2i}, Z_{i,1}) &= Cov\left(\frac{\tilde{\tau}_i + \epsilon_{1i}}{\tau_i}\frac{\tilde{\tau}_i + \epsilon_{2i}}{\tau_i}, Z_{i,1}\right) \\
&= Cov\left(\left(\frac{\tilde{\tau}_i}{\tau_i}\right)^2, Z_{i,1}\right) + Cov\left(\frac{\tilde{\tau}_i\epsilon_{1i}}{\tau_i^2}, Z_{i,1}\right) + Cov\left(\frac{\tilde{\tau}_i\epsilon_{2i}}{\tau_i^2}, Z_{i,1}\right) \\
&\quad + Cov\left(\frac{\epsilon_{1i}\epsilon_{2i}}{\tau_i^2}, Z_{i,1}\right) \\
&= Cov\left(\theta_i^2, Z_{i,1}\right) + Cov\left(\frac{\tilde{\tau}_i\epsilon_{1i}}{\tau_i^2}, Z_{i,1}\right) + Cov\left(\frac{\tilde{\tau}_i\epsilon_{2i}}{\tau_i^2}, Z_{i,1}\right) + Cov\left(\frac{\epsilon_{1i}\epsilon_{2i}}{\tau_i^2}, Z_{i,1}\right)
\end{aligned}$$

Then we simplify the three covariance terms immediately above that follow $Cov(\theta_i^2, Z_{i,1})$.

For the first term:

$$Cov\left(\frac{\tilde{\tau}_i\epsilon_{1i}}{\tau_i^2}, Z_{i,1}\right) = E\left[Cov\left(\frac{\tilde{\tau}_i\epsilon_{1i}}{\tau_i^2}, Z_{i,1}\middle|\tau_i\right)\right] + Cov\left(E\left[\frac{\tilde{\tau}_i\epsilon_{1i}}{\tau_i^2}\middle|\tau_i\right], E\left[Z_{i,1}|\tau_i\right]\right)$$

Of the two terms involved, we have:

$$\begin{aligned}
E\left[Cov\left(\frac{\tilde{\tau}_i\epsilon_{1i}}{\tau_i^2}, Z_{i,1}\middle|\tau_i\right)\right] &= E\left[\frac{1}{\tau_i^2}Cov\left(\tilde{\tau}_i\epsilon_{1i}, Z_{i,1}|\tau_i\right)\right] \\
&= E\left[\frac{1}{\tau_i^2}\left(E\left[\tilde{\tau}_i\epsilon_{1i}Z_{i,1}|\tau_i\right] - \underbrace{E[\tilde{\tau}_i\epsilon_{1i}|\tau_i]}_{=0, \text{ by } Cov(\tilde{\tau}_i,\epsilon_{1i}|\tau_i)=0, E[\epsilon_{1i}|\tau_i]=0} E[Z_{i,1}|\tau_i]\right)\right] \\
&= E\left[\frac{1}{\tau_i^2}E\left[Z_{i,1}\underbrace{E[\tilde{\tau}_i\epsilon_{1i}|Z_{i,1},\tau_i]}_{=0, \text{ by } Cov(\tilde{\tau}_i,\epsilon_{1i}|Z_{i,1},\tau_i)=0, E[\epsilon_{1i}|Z_{i,1},\tau_i]=0}|\tau_i\right]\right] \\
&= 0
\end{aligned}$$

And for the second term, observe that $E\left[\frac{\tilde{\tau}_i\epsilon_{1i}}{\tau_i^2}|\tau_i\right] = \frac{1}{\tau_i^2}\left(Cov(\tilde{\tau}_i, \epsilon_{1i}|\tau_i) + E[\tilde{\tau}_i|\tau_i]E[\epsilon_{1i}|\tau_i]\right)$, which equals zero by the assumptions $Cov(\tilde{\tau}_i, \epsilon_{1i}|\tau_i) = 0$ and $E[\epsilon_{1i}|\tau_i] = 0$. Therefore, $Cov\left(E\left[\frac{\tilde{\tau}_i\epsilon_{1i}}{\tau_i^2}\middle|\tau_i\right], E\left[Z_{i,1}|\tau_i\right]\right) = 0$. Then we have for the first of the three covariance terms:

$$\begin{aligned}
Cov\left(\frac{\tilde{\tau}_i\epsilon_{1i}}{\tau_i^2}, Z_{i,1}\right) &= E\left[Cov\left(\frac{\tilde{\tau}_i\epsilon_{1i}}{\tau_i^2}, Z_{i,1}\middle|\tau_i\right)\right] + Cov\left(E\left[\frac{\tilde{\tau}_i\epsilon_{1i}}{\tau_i^2}\middle|\tau_i\right], E\left[Z_{i,1}|\tau_i\right]\right) \\
&= 0
\end{aligned}$$

For the second term:

$$\begin{aligned}
Cov\left(\frac{\tilde{\tau}_i\epsilon_{2i}}{\tau_i^2}, Z_{i,1}\right) &= E\left[\frac{\tilde{\tau}_i\epsilon_{2i}}{\tau_i^2}Z_{i,1}\right] - E\left[\frac{\tilde{\tau}_i\epsilon_{2i}}{\tau_i^2}\right]E\left[Z_{i,1}\right] \\
&= E\left[\frac{Z_{i,1}}{\tau_i^2}E[\tilde{\tau}_i\epsilon_{2i}|Z_{i,1},\tau_i]\right] - E\left[\frac{1}{\tau_i^2}E[\tilde{\tau}_i\epsilon_{2i}|\tau_i]\right]E[Z_{i,1}]
\end{aligned}$$

Then we can simplify $E[\tilde{\tau}_i\epsilon_{2i}|Z_{i,1},\tau_i] = E[\tilde{\tau}_i|Z_{i,1},\tau_i]E[\epsilon_{2i}|\tau_i]$ by the assumptions $Cov(\tilde{\tau}_i,\epsilon_{2i}|Z_{i,1},\tau_i) = 0$ and $E[\epsilon_{2i}|Z_{i,1},\tau_i] = E[\epsilon_{2i}|\tau_i]$. We can also simplify $E[\tilde{\tau}_i\epsilon_{2i}|\tau_i] = E[\tilde{\tau}_i|\tau_i]E[\epsilon_{2i}|\tau_i]$ by the assumption $Cov(\tilde{\tau}_i,\epsilon_{2i}|\tau_i) = 0$. Therefore, continuing:

$$\begin{aligned}
Cov\left(\frac{\tilde{\tau}_i\epsilon_{2i}}{\tau_i^2}, Z_{i,1}\right) &= E\left[\frac{Z_{i,1}}{\tau_i^2}E[\tilde{\tau}_i\epsilon_{2i}|Z_{i,1},\tau_i]\right] - E\left[\frac{1}{\tau_i^2}E[\tilde{\tau}_i\epsilon_{2i}|\tau_i]\right]E[Z_{i,1}] \\
&= E\left[\frac{Z_{i,1}}{\tau_i^2}E[\tilde{\tau}_i|Z_{i,1},\tau_i]E[\epsilon_{2i}|\tau_i]\right] - E\left[\frac{1}{\tau_i^2}E[\tilde{\tau}_i|\tau_i]E[\epsilon_{2i}|\tau_i]\right]E[Z_{i,1}] \\
&= E\left[\frac{E[\epsilon_{2i}|\tau_i]}{\tau_i^2}\tilde{\tau}_i Z_{i,1}\right] - E\left[\frac{E[\epsilon_{2i}|\tau_i]}{\tau_i^2}\tilde{\tau}_i\right]E[Z_{i,1}] \\
&= Cov\left(\frac{E[\epsilon_{2i}|\tau_i]}{\tau_i^2}\tilde{\tau}_i, Z_{i,1}\right)
\end{aligned}$$

Finally, for the third term:

$$Cov\left(\frac{\epsilon_{1i}\epsilon_{2i}}{\tau_i^2}, Z_{i,1}\right) = E\left[Cov\left(\frac{\epsilon_{1i}\epsilon_{2i}}{\tau_i^2}, Z_{i,1}\middle|\tau_i\right)\right] + Cov\left(E\left[\frac{\epsilon_{1i}\epsilon_{2i}}{\tau_i^2}\middle|\tau_i\right], E[Z_{i,1}|\tau_i]\right)$$

Of these, the first term can be simplified as:

$$\begin{aligned}
E\left[Cov\left(\frac{\epsilon_{1i}\epsilon_{2i}}{\tau_i^2}, Z_{i,1}\middle|\tau_i\right)\right] &= E\left[\frac{1}{\tau_i^2}(E[\epsilon_{1i}\epsilon_{2i}Z_{i,1}|\tau_i] - \underbrace{E[\epsilon_{1i}\epsilon_{2i}|\tau_i]}_{=0,\text{ by } Cov(\epsilon_{1i},\epsilon_{2i}|\tau_i)=0, E[\epsilon_{1i}|\tau_i]=0} E[Z_{i,1}|\tau_i])\right] \\
&= E\left[\frac{1}{\tau_i^2}E[Z_{i,1}\underbrace{E[\epsilon_{1i}\epsilon_{2i}|Z_{i,1},\tau_i]}_{=0,\text{ by } Cov(\epsilon_{1i},\epsilon_{2i}|Z_{i,1},\tau_i)=0, E[\epsilon_{1i}|Z_{i,1},\tau_i]=0}|\tau_i]\right] \\
&= 0
\end{aligned}$$

The second term can be simplified as:

$$\begin{aligned}
Cov\left(E\left[\frac{\epsilon_{1i}\epsilon_{2i}}{\tau_i^2}\middle|\tau_i\right], E[Z_{i,1}|\tau_i]\right) &= Cov\left(\frac{1}{\tau_i^2}\underbrace{E[\epsilon_{1i}\epsilon_{2i}|\tau_i]}_{=0,\text{ by } Cov(\epsilon_{1i},\epsilon_{2i}|\tau_i)=0, E[\epsilon_{1i}|\tau_i]=0}, E[Z_{i,1}|\tau_i]\right) \\
&= 0
\end{aligned}$$

So we conclude the third term $Cov\left(\frac{\epsilon_{1i}\epsilon_{2i}}{\tau_i^2}, Z_{i,1}\right) = 0$.

This leaves us with:

$$Cov\left(\hat{\theta}_{1i}\hat{\theta}_{2i}, Z_{i,1}\right) = Cov\left(\theta_i^2, Z_{i,1}\right) + Cov\left(\frac{E[\epsilon_{2i}|\tau_i]}{\tau_i^2}\tilde{\tau}_i, Z_{i,1}\right)$$

Now we show that we can identify the nuisance term as $Cov\left(\frac{E[\hat{\tilde{\tau}}_{2i}-\hat{\tilde{\tau}}_{1i}|\tau_i]}{\tau_i^2}\hat{\tilde{\tau}}_{1i}, Z_{i,1}\right) =$

$Cov\left(\frac{E[\epsilon_{2i}|\tau_i]}{\tau_i^2}\tilde{\tau}_i, Z_{i,1}\right)$:

$$
\begin{aligned}
Cov\left(\frac{E[\hat{\tilde{\tau}}_{2i}-\hat{\tilde{\tau}}_{1i}|\tau_i]}{\tau_i^2}\hat{\tilde{\tau}}_{1i}, Z_{i,1}\right) &= Cov\left(\frac{E[(\tilde{\tau}_i+\epsilon_{2i})-(\tilde{\tau}_i+\epsilon_{1i})|\tau_i]}{\tau_i^2}(\tilde{\tau}_i+\epsilon_{1i}), Z_{i,1}\right) \\
&= Cov\left(\frac{E[\epsilon_{2i}|\tau_i] - \overbrace{E[\epsilon_{1i}|\tau_i]}^{=0,\text{ by assumption}}}{\tau_i^2}(\tilde{\tau}_i+\epsilon_{1i}), Z_{i,1}\right) \\
&= Cov\left(\frac{E[\epsilon_{2i}|\tau_i]}{\tau_i^2}\tilde{\tau}_i, Z_{i,1}\right) + Cov\left(\frac{E[\epsilon_{2i}|\tau_i]}{\tau_i^2}\epsilon_{1i}, Z_{i,1}\right) \\
&= Cov\left(\frac{E[\epsilon_{2i}|\tau_i]}{\tau_i^2}\tilde{\tau}_i, Z_{i,1}\right) + E\left[\frac{E[\epsilon_{2i}|\tau_i]}{\tau_i^2}\underbrace{Cov(\epsilon_{1i}, Z_{i,1}|\tau_i)}_{=0,\text{ by assumption}}\right] + Cov\left(\frac{E[\epsilon_{2i}|\tau_i]}{\tau_i^2}\underbrace{E[\epsilon_{1i}|\tau_i]}_{=0,\text{ by assumption}}, E[Z_{i,1}|\tau_i]\right) \\
&= Cov\left(\frac{E[\epsilon_{2i}|\tau_i]}{\tau_i^2}\tilde{\tau}_i, Z_{i,1}\right)
\end{aligned}
$$

Therefore, combining we have:

$$
\begin{aligned}
Cov(\hat{\theta}_{1i}\hat{\theta}_{2i}, Z_{i,1}) - Cov\left(\frac{E[\hat{\tilde{\tau}}_{2i}-\hat{\tilde{\tau}}_{1i}|\tau_i]}{\tau_i^2}\hat{\tilde{\tau}}_{1i}, Z_{i,1}\right) &= Cov(\theta_i^2, Z_{i,1}) \\
&\quad + Cov\left(\frac{E[\epsilon_{2i}|\tau_i]}{\tau_i^2}\tilde{\tau}_i, Z_{i,1}\right) - Cov\left(\frac{E[\epsilon_{2i}|\tau_i]}{\tau_i^2}\tilde{\tau}_i, Z_{i,1}\right) \\
&= Cov(\theta_i^2, Z_{i,1})
\end{aligned}
$$

5. $Cov(\theta_i^2, e(X_i)) = Cov(\hat{\theta}_{1i}\hat{\theta}_{2i}, e(X_i)) - Cov\left(\frac{E[\hat{\tilde{\tau}}_{2i}-\hat{\tilde{\tau}}_{1i}|\tau_i]}{\tau_i^2}\hat{\tilde{\tau}}_{1i}, e(X_i)\right)$

The steps proceed identically to the preceding proof demonstrating that $Cov(\theta_i^2, Z_{i,1}) = Cov(\hat{\theta}_{1i}\hat{\theta}_{2i}, Z_{i,1}) - Cov\left(\frac{E[\hat{\tilde{\tau}}_{2i}-\hat{\tilde{\tau}}_{1i}|\tau_i]}{\tau_i^2}\hat{\tilde{\tau}}_{1i}, Z_{i,1}\right)$, except with $e(X_i)$ in place of $Z_{i,1}$ and using the following assumptions: $Cov(\tilde{\tau}_i, \epsilon_{1i}|e(X_i), \tau_i) = 0$, $E[\epsilon_{1i}|e(X_i), \tau_i] = 0$, $Cov(\epsilon_{1i}, \epsilon_{2i}|e(X_i), \tau_i) = 0$, $Cov(\tilde{\tau}_i, \epsilon_{2i}|e(X_i), \tau_i) = 0$, $E[\epsilon_{2i}|e(X_i), \tau_i] = E[\epsilon_{2i}|\tau_i]$, and $Cov(\epsilon_{1i}, e(X_i)|\tau_i) = 0$.

6. $Cov(\theta_i^2, \tau_i^2) = Cov(\hat{\theta}_{1i}\hat{\theta}_{2i}, \tau_i^2) - Cov\left(\frac{E[\hat{\tilde{\tau}}_{2i}-\hat{\tilde{\tau}}_{1i}|\tau_i]}{\tau_i^2}\hat{\tilde{\tau}}_{1i}, \tau_i^2\right)$

$$
\begin{aligned}
Cov(\hat{\theta}_{1i}\hat{\theta}_{2i}, \tau_i^2) &= Cov\left(\frac{\tilde{\tau}_i+\epsilon_{1i}}{\tau_i}\frac{\tilde{\tau}_i+\epsilon_{2i}}{\tau_i}, \tau_i^2\right) \\
&= Cov\left(\left(\frac{\tilde{\tau}_i}{\tau_i}\right)^2, \tau_i^2\right) + Cov\left(\frac{\tilde{\tau}_i\epsilon_{1i}}{\tau_i^2}, \tau_i^2\right) + Cov\left(\frac{\tilde{\tau}_i\epsilon_{2i}}{\tau_i^2}, \tau_i^2\right) + Cov\left(\frac{\epsilon_{1i}\epsilon_{2i}}{\tau_i^2}, \tau_i^2\right) \\
&= Cov(\theta_i^2, \tau_i^2) + Cov\left(\frac{\tilde{\tau}_i\epsilon_{1i}}{\tau_i^2}, \tau_i^2\right) + Cov\left(\frac{\tilde{\tau}_i\epsilon_{2i}}{\tau_i^2}, \tau_i^2\right) + Cov\left(\frac{\epsilon_{1i}\epsilon_{2i}}{\tau_i^2}, \tau_i^2\right)
\end{aligned}
$$

Then we simplify the three covariance terms immediately above that follow $Cov(\theta_i^2, \tau_i^2)$.

For the first term:

$$
Cov\left(\frac{\tilde{\tau}_i\epsilon_{1i}}{\tau_i^2}, \tau_i^2\right) = E\left[Cov\left(\frac{\tilde{\tau}_i\epsilon_{1i}}{\tau_i^2}, \tau_i^2\middle|\tau_i\right)\right] + Cov\left(E\left[\frac{\tilde{\tau}_i\epsilon_{1i}}{\tau_i^2}\middle|\tau_i\right], E[\tau_i^2|\tau_i]\right)
$$

Of these, for the first term we have:

$$
\begin{aligned}
E\left[Cov\left(\frac{\tilde{\tau}_i\epsilon_{1i}}{\tau_i^2}, \tau_i^2\middle|\tau_i\right)\right] &= E[Cov(\tilde{\tau}_i\epsilon_{1i}, 1|\tau_i)] \\
&= 0
\end{aligned}
$$

And for the second we have:

$$\begin{aligned} Cov\left(E\left[\frac{\tilde{\tau}_i\epsilon_{1i}}{\tau_i^2}\Big|\tau_i\right], E[\tau_i^2|\tau_i]\right) &= Cov\left(\frac{1}{\tau_i^2}\underbrace{E[\tilde{\tau}_i\epsilon_{1i}|\tau_i]}_{=0, \text{ by } Cov(\tilde{\tau}_i,\epsilon_{1i}|\tau_i)=0, E[\epsilon_{1i}|\tau_i]=0}, \tau_i^2\right) \\ &= 0 \end{aligned}$$

Therefore, we have $Cov\left(\frac{\tilde{\tau}_i\epsilon_{1i}}{\tau_i^2}, \tau_i^2\right) = 0$.

For the second term:

$$Cov\left(\frac{\tilde{\tau}_i\epsilon_{2i}}{\tau_i^2}, \tau_i^2\right) = Cov\left(\frac{E[\tilde{\tau}_i\epsilon_{2i}|\tau_i]}{\tau_i^2}, \tau_i^2\right) + E\left[\underbrace{Cov\left(\frac{\tilde{\tau}_i\epsilon_{2i}}{\tau_i^2}, \tau_i^2|\tau_i\right)}_{=0}\right]$$

Then using the assumption $Cov(\tilde{\tau}_i, \epsilon_{2i}|\tau_i) = 0$, we get:

$$Cov\left(\frac{\tilde{\tau}_i\epsilon_{2i}}{\tau_i^2}, \tau_i^2\right) = Cov\left(\frac{E[\epsilon_{2i}|\tau_i]}{\tau_i^2}E[\tilde{\tau}_i|\tau_i], \tau_i^2\right)$$

For the third term:

$$Cov\left(\frac{\epsilon_{1i}\epsilon_{2i}}{\tau_i^2}, \tau_i^2\right) = E\left[Cov\left(\frac{\epsilon_{1i}\epsilon_{2i}}{\tau_i^2}, \tau_i^2\Big|\tau_i\right)\right] + Cov\left(E\left[\frac{\epsilon_{1i}\epsilon_{2i}}{\tau_i^2}\Big|\tau_i\right], E[\tau_i^2|\tau_i]\right)$$

For the first term we have:

$$\begin{aligned} E\left[Cov\left(\frac{\epsilon_{1i}\epsilon_{2i}}{\tau_i^2}, \tau_i^2\Big|\tau_i\right)\right] &= E[Cov(\epsilon_{1i}\epsilon_{2i}, 1|\tau_i)] \\ &= 0 \end{aligned}$$

For the second term we have:

$$\begin{aligned} Cov\left(E\left[\frac{\epsilon_{1i}\epsilon_{2i}}{\tau_i^2}\Big|\tau_i\right], E[\tau_i^2|\tau_i]\right) &= Cov\left(\frac{1}{\tau_i^2}\underbrace{E[\epsilon_{1i}\epsilon_{2i}|\tau_i]}_{=0, \text{ by } Cov(\epsilon_{1i},\epsilon_{2i}|\tau_i)=0, E[\epsilon_{1i}|\tau_i]=0}, \tau_i^2\right) \\ &= 0 \end{aligned}$$

Therefore, $Cov\left(\frac{\epsilon_{1i}\epsilon_{2i}}{\tau_i^2}, \tau_i^2\right) = 0$.

This leaves us with:

$$Cov(\hat{\theta}_{1i}\hat{\theta}_{2i}, \tau_i^2) = Cov(\theta_i^2, \tau_i^2) + Cov\left(\frac{E[\epsilon_{2i}|\tau_i]}{\tau_i^2}E[\tilde{\tau}_i|\tau_i], \tau_i^2\right)$$

Now we show that we can identify the nuisance term as $Cov\left(\frac{E[\hat{\tilde{\tau}}_{2i}-\hat{\tilde{\tau}}_{1i}|\tau_i]}{\tau_i^2}\hat{\tilde{\tau}}_{1i}, \tau_i^2\right) =$

$Cov\left(\frac{E[\epsilon_{2i}|\tau_i]}{\tau_i^2}E[\tilde{\tau}_i|\tau_i], \tau_i^2\right)$:

$$
\begin{aligned}
Cov\left(\frac{E[\hat{\tilde{\tau}}_{2i} - \hat{\tilde{\tau}}_{1i}|\tau_i]}{\tau_i^2}\hat{\tilde{\tau}}_{1i}, \tau_i^2\right) &= Cov\left(\frac{E[(\tilde{\tau}_i + \epsilon_{2i}) - (\tilde{\tau}_i + \epsilon_{1i})|\tau_i]}{\tau_i^2}(\tilde{\tau}_i + \epsilon_{1i}), \tau_i^2\right) \\
&= Cov\left(\frac{E[\epsilon_{2i}|\tau_i] - \overbrace{E[\epsilon_{1i}|\tau_i]}^{=0, \text{ by assumption}}}{\tau_i^2}(\tilde{\tau}_i + \epsilon_{1i}), \tau_i^2\right) \\
&= Cov\left(\frac{E[\epsilon_{2i}|\tau_i]}{\tau_i^2}\tilde{\tau}_i, \tau_i^2\right) + Cov\left(\frac{E[\epsilon_{2i}|\tau_i]}{\tau_i^2}\epsilon_{1i}, \tau_i^2\right) \\
&= Cov\left(\frac{E[\epsilon_{2i}|\tau_i]}{\tau_i^2}E[\tilde{\tau}_i|\tau_i], \tau_i^2\right) + E\left[\underbrace{Cov\left(\frac{E[\epsilon_{2i}|\tau_i]}{\tau_i^2}\tilde{\tau}_i, \tau_i^2|\tau_i\right)}_{=0}\right] \\
&\quad + Cov\left(\frac{E[\epsilon_{2i}|\tau_i]}{\tau_i^2}\underbrace{E[\epsilon_{1i}|\tau_i]}_{=0, \text{ by assumption}}, \tau_i^2\right) + E\left[\underbrace{Cov\left(\frac{E[\epsilon_{2i}|\tau_i]}{\tau_i^2}\epsilon_{1i}, \tau_i^2|\tau_i\right)}_{=0}\right] \\
&= Cov\left(\frac{E[\epsilon_{2i}|\tau_i]}{\tau_i^2}E[\tilde{\tau}_i|\tau_i], \tau_i^2\right)
\end{aligned}
$$

Therefore, combining we have:

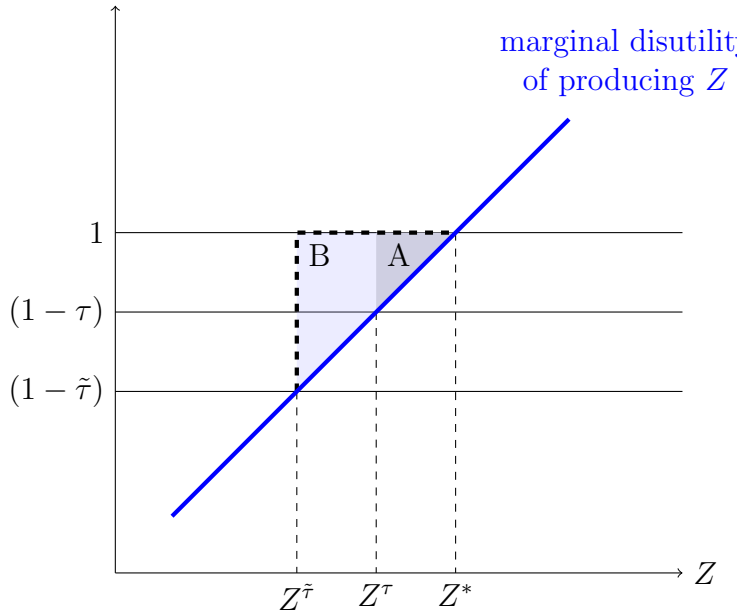
$$
\begin{aligned}
Cov(\hat{\theta}_{1i}\hat{\theta}_{2i}, \tau_i^2) - Cov\left(\frac{E[\hat{\tilde{\tau}}_{2i} - \hat{\tilde{\tau}}_{1i}|\tau_i]}{\tau_i^2}\hat{\tilde{\tau}}_{1i}, \tau_i^2\right) &= Cov(\theta_i^2, \tau_i^2) \\
&\quad + Cov\left(\frac{E[\epsilon_{2i}|\tau_i]}{\tau_i^2}E[\tilde{\tau}_i|\tau_i], \tau_i^2\right) - Cov\left(\frac{E[\epsilon_{2i}|\tau_i]}{\tau_i^2}E[\tilde{\tau}_i|\tau_i], \tau_i^2\right) \\
&= Cov(\theta_i^2, \tau_i^2)
\end{aligned}
$$

Supplementary Appendix C Why DWL Increases with the Variance of Tax Rate Misperceptions

To see why DWL increases with the variance of misperceptions, consider Figure C.1. Assume there are two agents with the same structural curve of marginal disutility of producing income Z who both face a true tax τ. In the first scenario, assume both agents correctly perceive the tax and choose income Z^{τ}. This produces total DWL of $2A$.[30] In the second scenario, assume one agent misperceives the tax as $\tilde{\tau} = 2\tau$, and the other agent misperceives the tax as zero. The first agent chooses income $Z^{\tilde{\tau}}$, contributing DWL equal to B, the area of the larger triangle. The second agent chooses income Z^{*} and contributes zero DWL, making the total DWL in this scenario equal to B. In both scenarios, the average perceived tax equals the true tax τ. However, DWL in the second scenario with higher variance of misperceptions is *twice* that of the first scenario. The intuition for this result can also be seen in Figure C.1. Increasing one agent's perceived tax above τ produces additional DWL from two factors: first, the triangle between $(1-\tau)$ and $(1-\tilde{\tau})$ that reflects increasing misperception from zero to $\tilde{\tau}-\tau$; and second, the rectangle between 1 and $(1-\tau)$ on the vertical axis and between $Z^{\tilde{\tau}}$ and Z^{τ} on the horizontal axis, which is the DWL that results from further reducing income from Z^{τ} to $Z^{\tilde{\tau}}$ when there is a pre-existing distortion at Z^{τ}. On the other hand, reducing the other agent's perceived tax from τ to 0 does not produce a symmetric reduction of the pre-existing distortion, as it simply removes the baseline triangle. This relates to the common result that DWL increases with the square of the tax rate, where misperception can be viewed equivalently to an additional tax levied on top of the pre-existing true tax – it is the average of this square, not the square of the average, that determines DWL.

[30] Recall that under the maintained assumption of quasilinear utility, the marginal utility of consumption is constant, which we set equal to one.

Figure C.1: The Impact of Variance of Tax Rate Misperceptions on DWL



Notes: This figure illustrates why variance of tax rate misperceptions increases DWL. The figure depicts a hypothetical marginal disutility of earning income curve alongside choices of income and deadweight loss triangles that result when the agent perceives a tax rate of zero, the true tax τ, or a higher tax $\tilde{\tau}$. In a scenario with two individuals who both correctly perceive a tax rate τ, DWL is $2A$. In a scenario where one individual overestimates their tax rate as $\tilde{\tau} = 2\tau$ and the second individual underestimates their tax rate as zero, DWL is B, which is twice as large as the DWL in the first scenario with correct perceptions. This is despite the fact that the average perception equals τ in both cases.

Supplementary Appendix D Prolific Survey Description

In this section, we describe the online survey administered through Prolific to measure parents' perceptions of college aid. Unless otherwise noted, each image below from the survey represents a complete and separate page of the survey. Figures are placed at the end of this section for readability.

D.1 Introduction

The survey began with introductory content. The first screen shown in Figure D.1 collected the respondent's Prolific ID, which was used as a quality check and to assist with payments. It also asked which device the respondent was using. The next screen in Figure D.2 asked for the year in which the respondent's oldest and youngest child were born, which were used as prescreeners for the survey. The survey was only available to individuals whose Prolific account listed an oldest or youngest child in a certain age range, so this question was used as an additional check. Respondents without a child in the required age range were sent to the end of the survey and were generally paid for their participation nonetheless. The next screen in Figure D.3 obtained consent. Respondents who did not select "Accept" were sent to the end of the survey and were generally paid for their participation nonetheless. The compensation amount varied from $3 to $4. The next screen in Figure D.4 presented a brief preamble describing the survey.

D.2 Background Questions

The next several screens collected background information, especially information that would influence later portions of the survey. The survey collected the respondent's state of residence (Figure D.5), which was used later to provide references to four-year public colleges in the respondent's state. The next screen (Figure D.6) measured the respondent's understanding of when income is used to determine a child's college aid.

The next few screens identified the relevant child, whose college aid would be a focus for the remainder of the survey. Respondents were first asked if they have any children age 0-22 who have not yet attended a four-year college as a full-time student, but who might in the future (Figure D.7). If the respondent selected "Yes", they would then see the subsequent page (Figure D.8) instructing the respondent to answer the remaining questions of the survey with this child in mind (and to select the oldest such child in the case of multiple children meeting the criteria). If the respondent instead selected "No" regarding having a child age 0-22 who had not yet attended a four-year college, they would instead be asked in order whether they have any dependent children age 17-25 who are currently enrolled in a four-year college full-time and then whether they have any dependent children age 17-25 who were previously enrolled in a four-year college full-time. If the respondent answered "Yes" to either of these, they would be shown a page similar to Figure D.8, with relevant language modifications. If the respondent answered "No" to all three prompts, or if they answered "No" when asked to confirm that they have identified a child meeting the relevant conditions, they were sent to the end of the survey and were

generally paid nonetheless. The figures shown in this document correspond to what a respondent would see if they selected "Yes" in response to having any children age 0-22 who have not yet attended a four-year college as a full-time student, but who might in the future. Some survey questions that related to the child's own college attendance were phrased differently or omitted for other respondents as relevant; for example, instead of being asked the likelihood that the respondent or a guardian would submit a FAFSA if the child eventually attends a four-year college, a respondent with a child who had previously attended college was asked whether the child or a guardian has ever submitted a FAFSA for the child.

The next couple screens regarded whether the respondent is a primary provider for the relevant child. First, the respondent was asked whether they are married to or living with another parent of the child (Figure D.9). If the respondent answered "Yes. I am married to or living with this child's other parent.", then they were shown a screen noting that the remainder of the survey will refer to the child's other parent as the respondent's "spouse or partner", and asking the respondent to answer questions with this person in mind even if they are not a "spouse or partner". If the respondent instead answered "No. This child does not have another living parent, or I am not married to or living with that person.", then they were asked whether the child lives with them for more than half of the year (Figure D.10).

D.3 Baseline Questions about College

The next screen asked questions related to the relevant child's anticipated college enrollment and aid (Figure D.11). These included questions about the child's current age, whether and when the child would attend college, who would pay for most of the child's college education, whether the child or a guardian would submit a FAFSA, the respondent and their spouse or partner's anticipated labor supply and total household income in the years when income would affect the child's college aid, and the respondent's anticipated savings toward the child's college education. Respondents were randomized into being asked labor supply and household income questions either with language reflecting expected labor supply and income or the labor supply and corresponding household income they would like if they could work the number of hours they would like at their usual hourly wage.

D.4 Questions about Financial Aid at In-state Public Four-year Colleges

The next section asked the respondent various questions about their perceptions of financial aid at public four-year colleges in their home state. This section included the questions that were used to construct the primary measure and the secondary instrument measure of college aid perceptions used in "Behavioral Responses to Large, Complex Tax and Transfer Incentives: Evidence from College Financial Aid".

This section began with an introduction to such colleges. The first screen (Figure D.12) in this section noted that the section would ask questions about financial aid at four-year public colleges in the respondent's state (the state was indicated), and listed several such colleges that were specific to that state as reference points. It then asked how likely the child is to attend such a college, if they attend any four-year college.

The next screen (Figure D.13) asked the respondent for their best guess of the average sticker cost at public four-year colleges in their state, their best guess of the total annual financial aid their child would receive at such a college, and their best guess of the total aid their child would receive if their household income were $10,000 higher than it otherwise would be in the years affecting their child's aid. Respondents were randomized to see the earlier questions regarding the respondent and their spouse or partner's anticipated labor supply and total household income in the years when income would affect the child's college aid and the respondent's anticipated savings toward the child's college education either on the screen noted earlier or on a screen immediately following the screen in Figure D.13.

The next screen (Figure D.14) asked the respondent for their best guess of grants and scholarships for an in-state student attending a public four-year college in their state with the hypothetical student's family income varying over $45k, $55k, $65k, or $75k. We use the $45k and $75k responses to form our instrument for college aid implicit tax perceptions. Respondents were randomized into providing guesses for a student whose total parent non-retirement savings was either $5,000 or $500,000.

The next several screens asked the respondent for their best guess of average grant and scholarship aid at public four-year college in their state for in-state students with family income in one of several ranges. The first screen (Figure D.15) introduced these questions. The next five screens (Figure D.16) asked for the respondent's best guess of average grant and scholarship aid for a student with family income in each of the following ranges: $0-30k, $30-48k, $48-75k, $75-110k, and above $110k. We use responses to these five questions to form our main measure of perceptions of the college aid implicit tax. The next screen (Figure D.17) introduced the following questions, which would ask for average grant and scholarship aid for a family with two students enrolled at a public four-year college in their state, and asked the respondent to provide their best guess of total aid awarded to two students. The next five screens (Figure D.18) asked for the respondent's best guess of average grant and scholarship aid for two students with family income in each of the same preceding ranges.

D.5 Treatments

The survey then randomly exposed respondents to two types of treatments. The first (Figure D.19) treatment explained the concept of college aid phaseout operating implicitly like a temporary tax on income. Individuals were randomized to either see this screen or to not see this screen. The next treatment was an information treatment, in which each respondent was randomized to either see the treated condition screen (Figure D.20) or the control condition screen (Figure D.21). Both screens reminded the respondent of their guesses of the average annual sticker cost at public four-year colleges in their state as well as average financial aid for a single student from each of the preceding five family income ranges. In the treated condition, the respondent was also shown the true sticker cost and the true average aid for a single student from each of the five family income ranges alongside their guesses. In the control condition, the respondent was instead shown true statistics of total enrollment at public four-year colleges in their state by each of the five family income ranges. In each condition, the respondent was then asked how reliable they find the U.S. Department of Education - which is the underlying source for the actual aid and enrollment statistics shown in the treatment and control conditions, respectively - as

an information source.

D.6 Post-treatment Questions about College and Own Aid

The next screen (Figure D.22) then re-asked several of the preceding questions including the respondent's best guess of the total annual financial aid their child would receive at such a college, their best guess of the total aid their child would receive if their household income were $10,000 higher than it otherwise would be in the years affecting their child's aid, the respondent and their spouse or partner's anticipated labor supply and total household income in the years when income would affect the child's college aid, the respondent's anticipated savings toward the child's college education, whether and where the child would attend college, and whether the child or a guardian would submit a FAFSA.

D.7 Demographics and Attention Check

The next screen (Figure D.23) presented a number of demographic questions along with an attention check. Demographic questions included information about household income (which was used to map respondents to their true implicit tax rate from college aid), non-wage income (which was used to proxy for higher versus lower elasticity when estimating the impact of misperception on welfare), savings, respondent age, and others.

D.8 Feedback and Survey End

The next screen (Figure D.24) asked for feedback. After this screen, the respondent was thanked for their participation and re-directed back to Prolific to register their submission.

D.9 Survey Figures

Figure D.1: Prolific ID

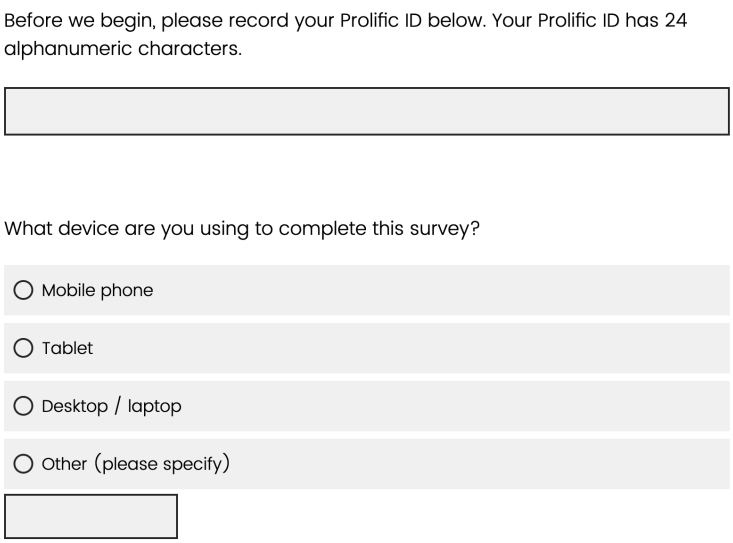


Figure D.2: Child Age

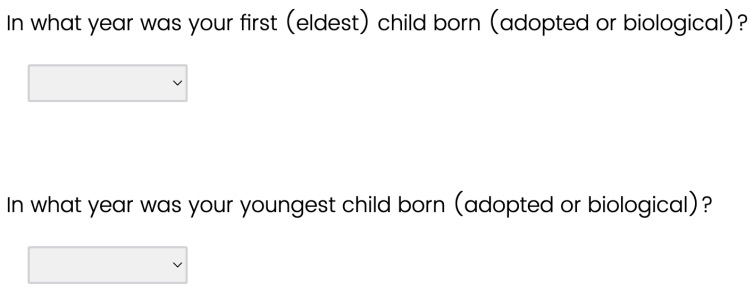


Figure D.3: Consent

University of California at Berkeley
Consent to Participate in Research

Introduction
We are Nick Gebbia and Patrick Kline, researchers at UC Berkeley. We would like to invite you to take part in our research study about college financial aid.

Procedures
If you agree to participate in our research, we will ask you to complete the following online survey. The survey will involve questions about college financial aid and some questions about your family, and it should take about 15 minutes to complete.

Confidentiality
Your study data will be handled as confidentially as possible. All responses to this survey are anonymous.

Compensation
To thank you for participating in this study, you will receive $4 in Prolific credit upon completion of the survey.

Participation is voluntary
This is a research study, and your participation is completely voluntary.

Questions
If you have any questions about this research, please feel free to contact us. You can reach me, Nick Gebbia, at nickgebbia@berkeley.edu. If you have any questions about your rights or treatment as a research participant in this study, please contact the University of California at Berkeley's Committee for Protection of Human Subjects at 510-642-7461, or e-mail subjects@berkeley.edu, where this project ID is CPHS #2023-05-16405.

If you agree to take part in the research, please print a copy of this page for future reference, then click on the "Accept" button below.

◯ Accept

◯ Do not accept

Figure D.4: Preamble

In this survey we will ask you questions regarding college financial aid, along with some questions about your family. Throughout the survey, there are questions to which you might not know an exact answer. If you are uncertain of an answer, simply provide your best guess.

Note also that we will use the word "college" to refer to either a college or a university.

Figure D.5: State

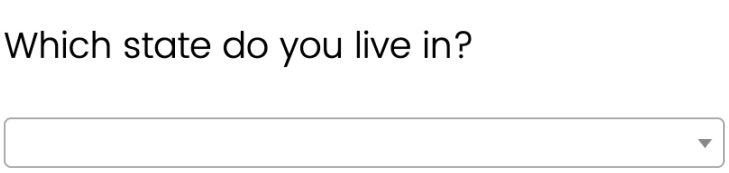


Figure D.6: Aid Timing

In order to be eligible for certain types of college financial aid, a student or their family must fill out a Free Application for Federal Student Aid (or, "FAFSA"). The FAFSA asks questions about family income and savings, which affects how much financial aid a student will receive.

Imagine you are filing a FAFSA to apply for financial aid for the 2023-24 school year (i.e. fall 2023 through spring 2024). The FAFSA will ask about your income. Which year of income will the FAFSA ask about?

◯ 2016	◯ 2021
◯ 2017	◯ 2022
◯ 2018	◯ 2023
◯ 2019	◯ 2024
◯ 2020	◯ 2025

Figure D.7: Child 1

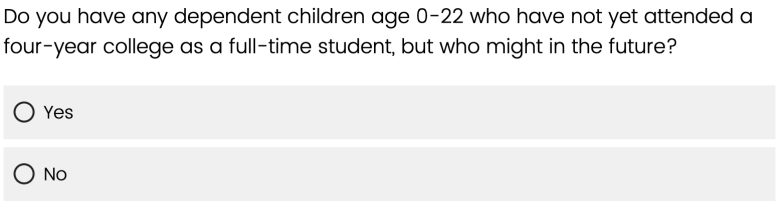


Figure D.8: Child 2

The remainder of the survey will be about your child age 0-22 who might attend a four-year college in the future. If you have more than one child who meets these criteria, select the oldest among them.

Please answer the remaining questions with this child in mind.

Have you identified one child who meets the above conditions?

- ◯ Yes
- ◯ No

Figure D.9: Provider 1

If this child has another living parent, are you married to or living with that person?

- ◯ Yes. I am married to or living with this child's other parent.
- ◯ No. This child does not have another living parent, or I am not married to or living with that person.

Figure D.10: Provider 2

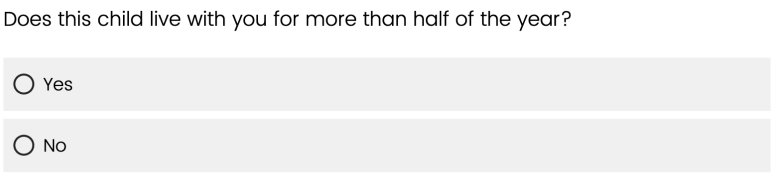


Figure D.11: Own Child's College - Baseline

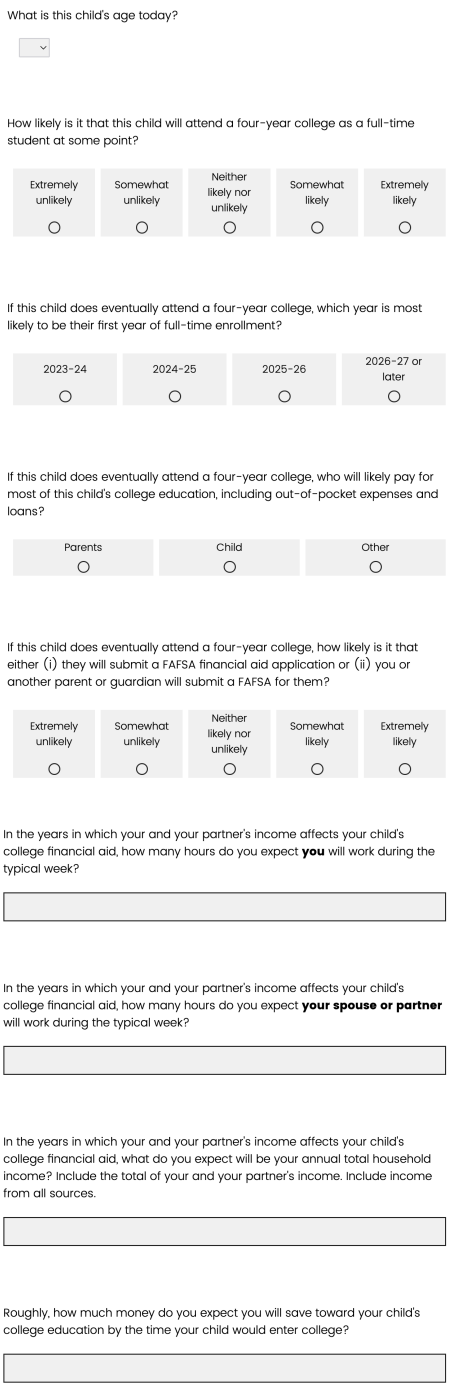


Figure D.12: Public Four-year In-state Colleges: Introduction

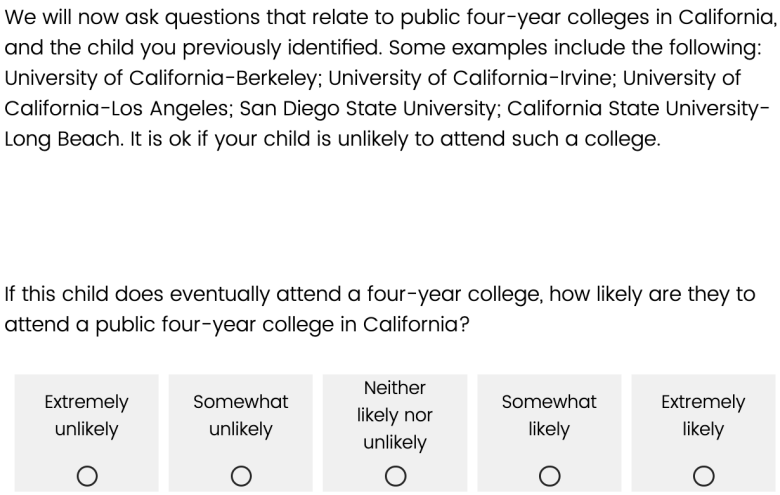


Figure D.13: Public Four-year In-state Colleges: Sticker Cost and Own Aid

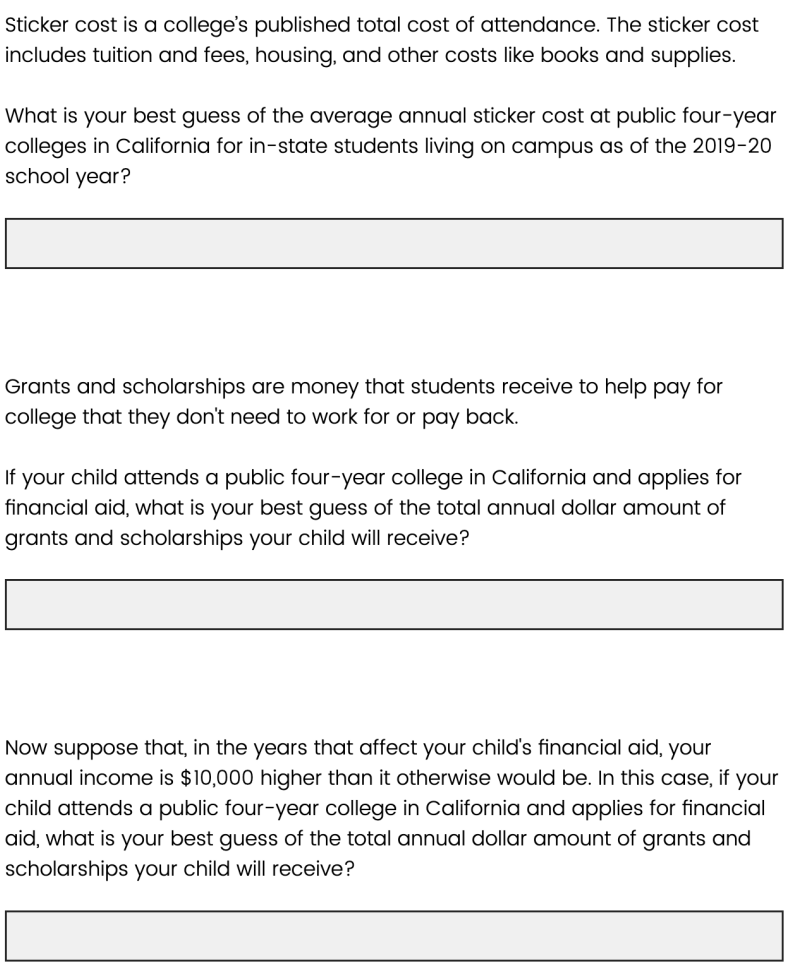


Figure D.14: Public Four-year In-state Colleges: Aid at $45-75k (Instrument Measure)

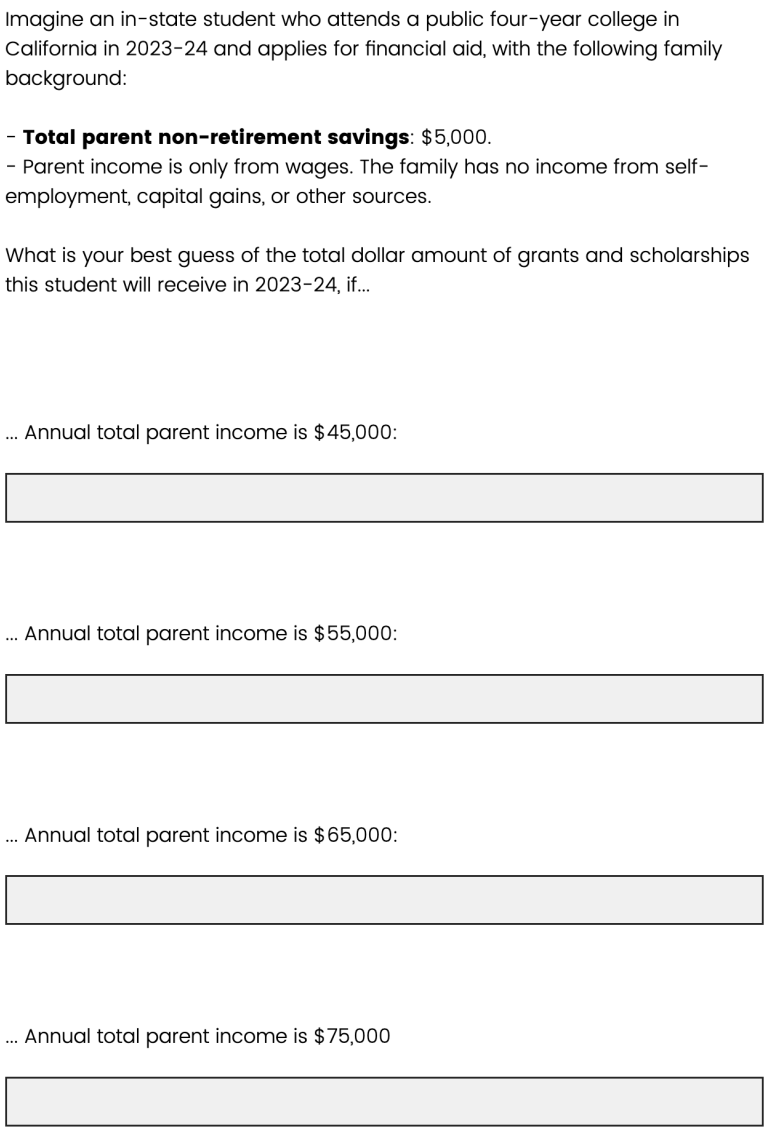


Figure D.15: Public Four-year In-state Colleges: Aid by Family Income Ranges, Introduction

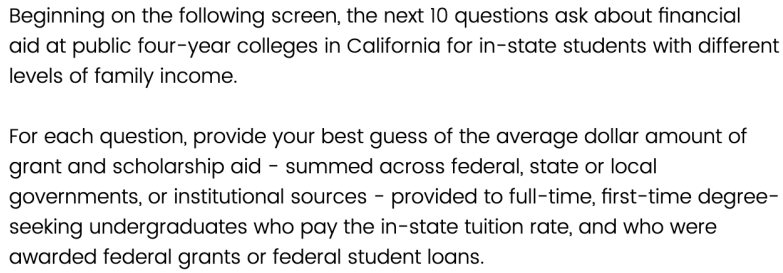
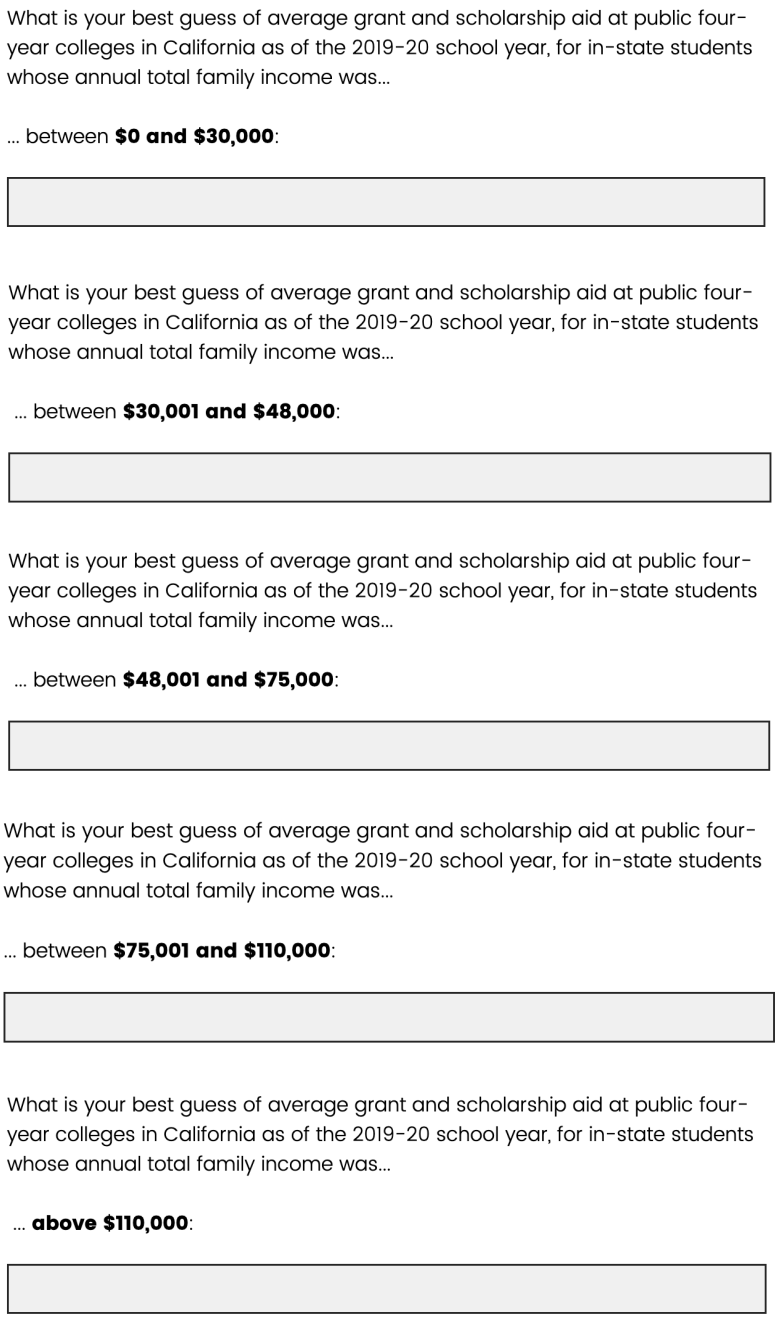


Figure D.16: Public Four-year In-state Colleges: Aid by Family Income Ranges (Main Perceptions Measure)



Notes: Each income range question was presented on a separate screen.

Figure D.17: Public Four-year In-state Colleges: Aid by Family Income Ranges, Two Students Introduction

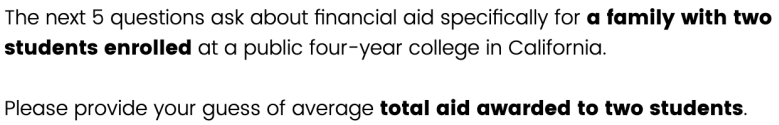
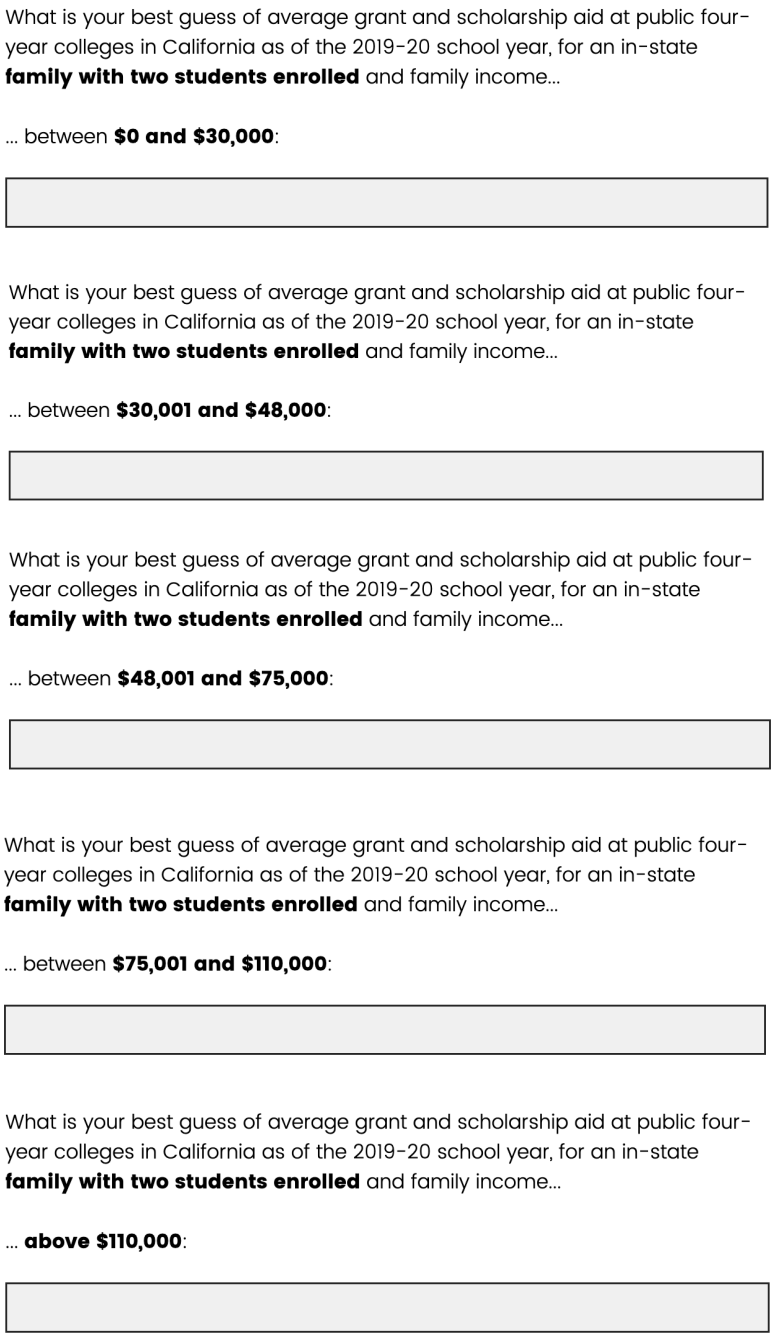


Figure D.18: Public Four-year In-state Colleges: Aid by Family Income Ranges, Two Students



Notes: Each income range question was presented on a separate screen.

Figure D.19: Explainer Treatment

Did you know?

College financial aid is similar to a temporary increase in your income tax rate. For each additional dollar of income, you will generally receive less aid, just like a tax.

As a result, each additional dollar of income produces less additional take home pay in the years when income affects college aid compared to the years before or after.

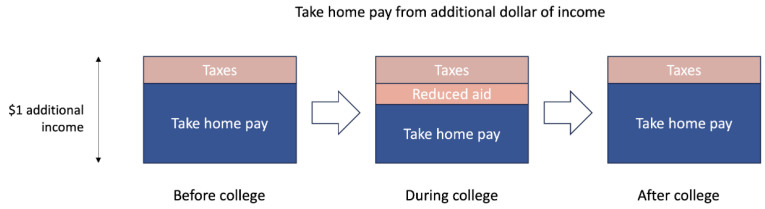


Figure D.20: Information Treatment: Treated Condition

We would like to share with you data on sticker price and average grant and scholarship aid by family income at public four-year colleges in California in 2019-20. We also list your guesses of sticker price and average grant and scholarship aid by income at public four-year colleges in California for reference. Data on sticker price and average grant and scholarship aid comes from the U.S. Department of Education.

For in-state on campus sticker price, your guess was $, and the actual price was $30,676.

For each income group, the table below shows your guess of average grants and scholarships and the actual amount for students who are full-time, first-time, degree-seeking, in-state, and received either federal grants or federal student loans:

Income	Aid: Your Guess	Aid: Actual
$0 to $30,000	$	$18,897
$30,001 to $48,000	$	$17,654
$48,001 to $75,000	$	$14,127
$75,001 to $110,000	$	$8,872
Above $110,000	$	$2,688

Note: The above averages are weighted by total enrollment at each college.

How reliable do you find the U.S. Department of Education as an information source?

Very unreliable	Somewhat unreliable	Neither reliable nor unreliable	Somewhat reliable	Very reliable
○	○	○	○	○

Figure D.21: Information Treatment: Control Condition

We would like to share with you data on the number of students who were enrolled at a public four-year college in California in 2019-20 by family income. We also list your guesses of sticker price and average grant and scholarship aid by income at public four-year colleges in California for reference. Data on the number of students enrolled and receiving federal grants or federal student loans comes from the U.S. Department of Education.

For in-state on campus sticker price, your guess was $.

For each income group, the table below shows your guess of average grants and scholarships and the actual number of students enrolled who are full-time, first-time, degree-seeking, in-state, and received either federal grants or federal student loans:

Income	Aid: Your Guess	Number enrolled: Actual
$0 to $30,000	$	23,520
$30,001 to $48,000	$	12,755
$48,001 to $75,000	$	10,922
$75,001 to $110,000	$	5,718
Above $110,000	$	8,098

How reliable do you find the U.S. Department of Education as an information source?

Very unreliable	Somewhat unreliable	Neither reliable nor unreliable	Somewhat reliable	Very reliable
○	○	○	○	○

Figure D.22: Post-treatment Questions

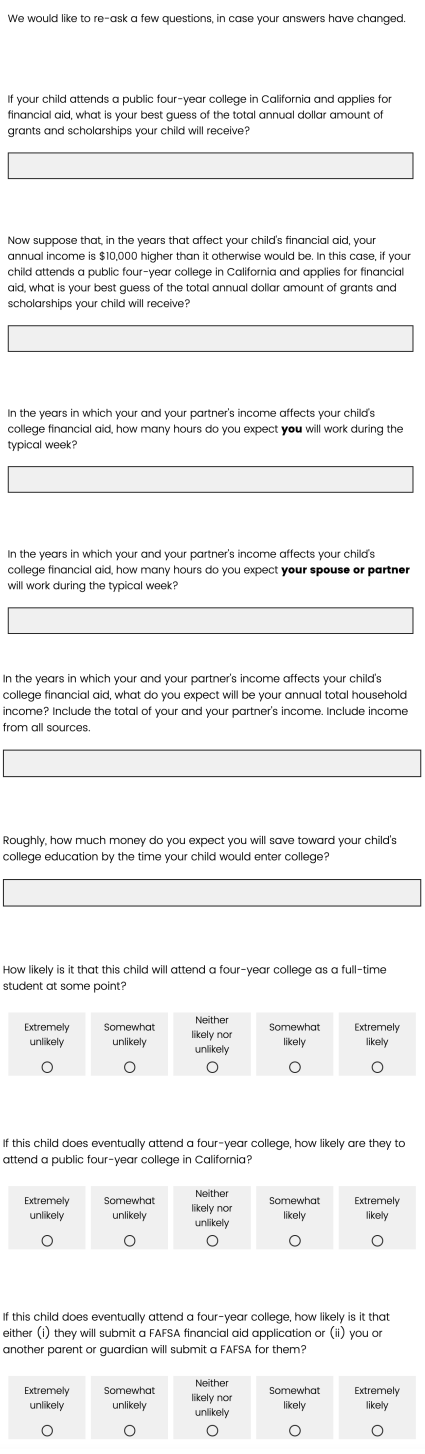
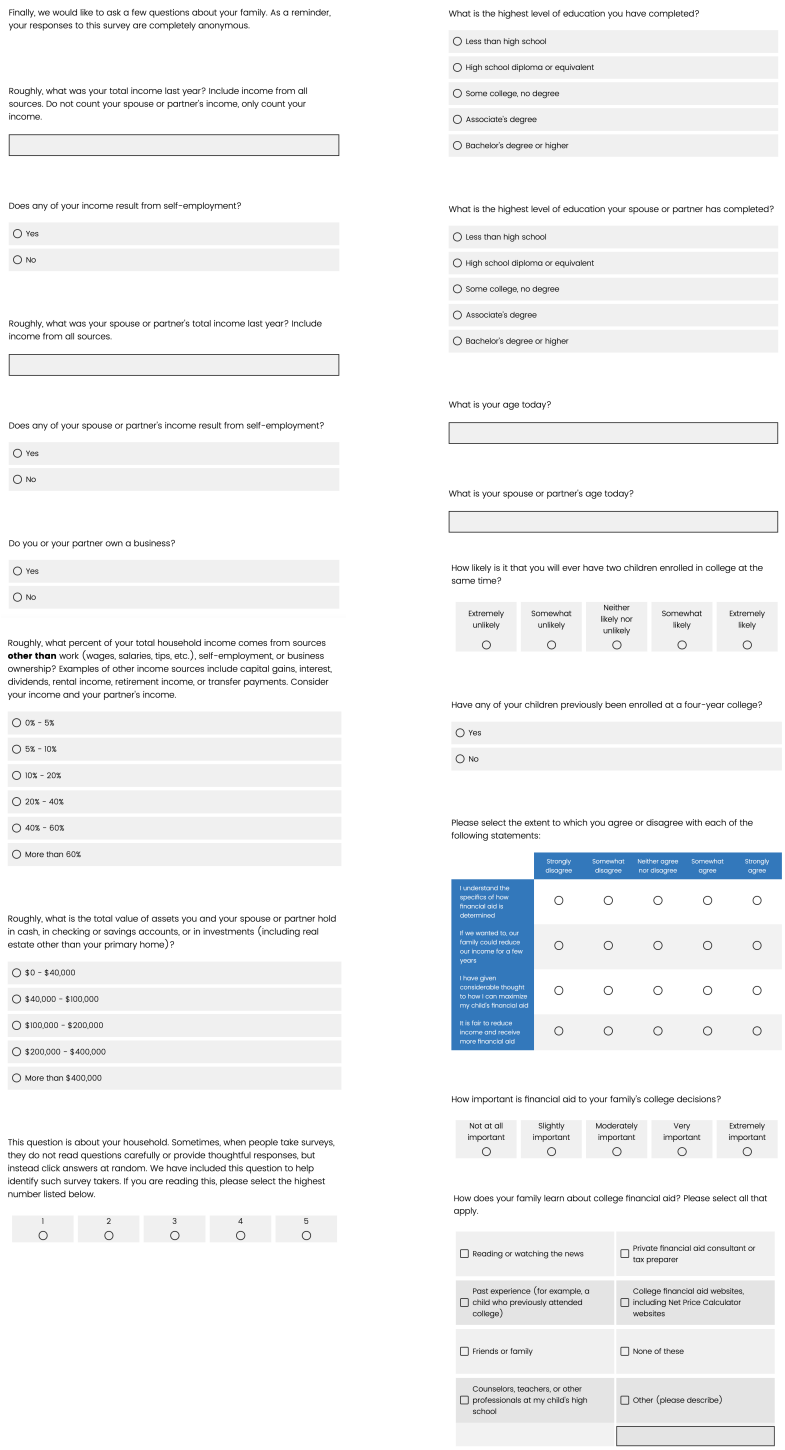


Figure D.23: Demographics and Attention Check



Finally, we would like to ask a few questions about your family. As a reminder, your responses to this survey are completely anonymous.

Roughly, what was your total income last year? Include income from all sources. Do not count your spouse or partner's income, only count your income.

Does any of your income result from self-employment?

- Yes
- No

Roughly, what was your spouse or partner's total income last year? Include income from all sources.

Does any of your spouse or partner's income result from self-employment?

- Yes
- No

Do you or your partner own a business?

- Yes
- No

Roughly, what percent of your total household income comes from sources **other than** work (wages, salaries, tips, etc.), self-employment, or business ownership? Examples of other income sources include capital gains, interest, dividends, rental income, retirement income, or transfer payments. Consider your income and your partner's income.

- 0% - 5%
- 5% - 10%
- 10% - 20%
- 20% - 40%
- 40% - 60%
- More than 60%

Roughly, what is the total value of assets you and your spouse or partner hold in cash, in checking or savings accounts, or in investments (including real estate other than your primary home)?

- $0 - $40,000
- $40,000 - $100,000
- $100,000 - $200,000
- $200,000 - $400,000
- More than $400,000

This question is about your household. Sometimes, when people take surveys, they do not read questions carefully or provide thoughtful responses, but instead click answers at random. We have included this question to help identify such survey takers. If you are reading this, please select the highest number listed below.

1	2	3	4	5

What is the highest level of education you have completed?

- Less than high school
- High school diploma or equivalent
- Some college, no degree
- Associate's degree
- Bachelor's degree or higher

What is the highest level of education your spouse or partner has completed?

- Less than high school
- High school diploma or equivalent
- Some college, no degree
- Associate's degree
- Bachelor's degree or higher

What is your age today?

What is your spouse or partner's age today?

How likely is it that you will ever have two children enrolled in college at the same time?

Extremely unlikely	Somewhat unlikely	Neither likely nor unlikely	Somewhat likely	Extremely likely

Have any of your children previously been enrolled at a four-year college?

- Yes
- No

Please select the extent to which you agree or disagree with each of the following statements:

	Strongly disagree	Somewhat disagree	Neither agree nor disagree	Somewhat agree	Strongly agree
I understand the specifics of how financial aid is determined					
If we wanted to, our family could reduce our income for a few years					
I have given considerable thought to how I can maximize my child's financial aid					
It is fair to reduce income and receive more financial aid					

How important is financial aid to your family's college decisions?

Not at all important	Slightly important	Moderately important	Very important	Extremely important

How does your family learn about college financial aid? Please select all that apply.

- Reading or watching the news
- Past experience (for example, a child who previously attended college)
- Friends or family
- Counselors, teachers, or other professionals at my child's high school
- Private financial aid consultant or tax preparer
- College financial aid websites, including Net Price Calculator websites
- None of these
- Other (please describe)

Notes: The questions were presented on a single screen all vertically stacked.

Figure D.24: Feedback

