

SELF-EMPLOYMENT DYNAMICS AND THE RETURNS TO ENTREPRENEURSHIP*

Eleanor W. Dillon Christopher T. Stanton

Arizona State University Harvard Business School,

NBER, and CEPR

January 30, 2017

First Version: February 12, 2016

Abstract

Small business owners and others in self-employment have the option to transition to paid work. If there is initial uncertainty about entrepreneurial earnings, this option increases the expected lifetime value of self-employment relative to pay in a single year. This paper first documents that moves between paid work and self-employment are common and consistent with experimentation to learn about earnings. This pattern motivates estimating the expected returns to entrepreneurship within a dynamic lifecycle model that allows for non-random selection and gradual learning about the entrepreneurial earnings process. The model accurately fits entry patterns into self-employment by age. The option value of returning to paid work is found to constitute a substantial portion of the monetary value of entrepreneurship. The model is then used to evaluate policies that change incentives for entry into self-employment.

*ewdillon@asu.edu, christopher.t.stanton@gmail.com. We thank Kelly Bishop, Luis Garicano, Erik Hurst, Bill Kerr, Edward Lazear, David de Meza, Robert Miller, Ramana Nanda, Paul Oyer, Luis Rayo, Yona Rubinstein, Nathan Seegert, Jeff Smith, Chris Taber, Galina Vereschagina, Matthew Wiswall, and seminar and conference participants at Arizona State University, IZA, the London School of Economics, the NBER Summer Institute, the Utah Winter Business Economics Conference, Wisconsin, and Wharton for helpful comments. All errors are our own.

1 Introduction

Nearly half of all workers who enter self-employment return to paid work within five years. In the model considered here, individuals cycle in and out of self-employment in part to resolve initial uncertainty about their potential earnings. Workers who enter self-employment observe their realized earnings and update their beliefs about the distribution of their future performance. In time, unsuccessful entrepreneurs who expect low future earnings shift back to paid work.¹ We document that movements between paid work and self-employment in the Panel Study of Income Dynamics (PSID) are consistent with this kind of learning.²

If workers view self-employment as an experimental trial that can be reversed, then static comparisons of earnings between entrepreneurs and paid workers provide a biased measure of the lifetime expected returns from entering self-employment.³ The first goal of the paper is to estimate the lifetime returns to entrepreneurship while accounting for the option value of returning to paid employment. This uncertainty about potential entrepreneurial earnings also has implications for policy evaluation. Policies that provide incentives or disincentives for workers to enter self-employment change whether entrants learn about their entrepreneurial ability. Encouraging entry has the potential to improve sorting based on comparative advantage. The second goal of the paper is to explore the long-run sorting and distributional implications of policies that promote entry rates into self-employment.

To accomplish these goals, we build and estimate a semi-structural model of lifecycle choices to work as an entrepreneur that incorporates features of the Bayesian learning model presented in seminal papers by Jovanovic (1979) and Miller (1984). To model mobility in and out of entrepreneurship, we specify a reasonably flexible forward-looking discrete choice model that encompasses the key monetary and non-monetary determinants of labor sector choice: sector-specific

¹We define anyone who reports working for themselves in their main job as an entrepreneur, therefore using entrepreneurship and self-employment as synonyms. We discuss this choice in the data section.

²This may be due to uncertainty about sector-specific ability or uncertainty about the value of an idea. For exposition purposes we focus on uncertainty about sector-specific ability, but both are consistent with the empirical exercises documented later. Kerr et al. (2014) discuss the role of venture capital in experimenting with different businesses or products, suggesting there is some uncertainty that must be resolved about the idea itself.

³Several earlier studies, reviewed in Section 2, have documented the puzzling finding that the median entrepreneur earns less in a given year than the median paid worker.

expected earnings, shocks to earnings, tastes for self-employment, which we allow to vary across workers, and costs to entering self-employment. We estimate the model on data from the PSID. The model is able to match movements between paid work and self-employment over the lifecycle quite accurately despite its relative sparsity, even along dimensions that are not explicitly targeted in estimation.

Using each individual’s sequence of predicted choices between paid work and entrepreneurship, we then project expected lifetime earnings conditional on the choice made while allowing workers to move between sectors in future years to maximize utility. The mean difference across workers between the expected discounted lifetime earnings conditional on choosing entrepreneurship and choosing paid work is an unbiased estimate of the expected monetary returns to entrepreneurship. We also calculate these projections under the counterfactual assumption that workers must remain in one sector in all future periods. The difference between the estimated value of entrepreneurship with and without the opportunity for later sector changes defines the value of the option to return to paid employment. Finally, we use the model to estimate the effects of two stylized counterfactual policies: a subsidy to encourage entry into self-employment and flat taxes, rather than progressive taxes, which may encourage more individuals to try self-employment because of the increased after-tax value of right-tail earnings.

Our analysis yields five key contributions.

1. Because of non-random selection, moments of the cross-sectional distribution of entrepreneurial earnings are biased estimates of the population moments that would prevail if everyone were observed in self-employment.⁴ Most existing work on this topic relies on observed worker characteristics to correct for selection. Our focus on mobility between paid work and self-employment, and the strong observed correlation in individual earnings between sectors, enables a different approach. We estimate population moments of the conditional distribution of

⁴This is not the first paper to highlight the bias that results from this kind of experimentation. Vereshchagina and Hopenhayn (2009) develop a model of entrepreneurial experimentation and Manso (2016) and Daly (2015) compare realized lifetime earnings for workers with and without any experience in self-employment, highlighting the role of experimentation in dynamic returns. Our model-based approach enables two key contributions. It allows us to separately estimate selection bias in cross-sectional earnings comparisons and the option value of experimentation. It also allows us to consider counterfactual policies.

entrepreneurial earnings given paid earnings, which embodies both observed and unobserved worker traits. We find that the raw median of cross-sectional entrepreneurial earnings for those observed in self-employment is biased downward while the mean is biased upward.⁵

2. Expected lifetime earnings from entering entrepreneurship look substantially more attractive than estimates from either the raw or corrected cross-sectional distribution of entrepreneurial earnings. For example, the mean 30 year old would expect to earn 6.6% less after taxes if he worked in self-employment in all futures years relative to his expected lifetime earnings in paid work. In contrast, if he can choose to change sectors in future years, a 30 year old who chooses self-employment next year can expect to earn 0.8% more in after-tax income over his life than if he chooses paid work.⁶
3. Tastes for self-employment vary considerably across workers.⁷ 15% of men experience a positive non-pecuniary benefit from self-employment, while the majority of the sample would require substantial compensation to overcome their dis-utility from working for themselves. Modeling this heterogeneity helps to match observed patterns of entrepreneurial choice, particularly at the entry margin.
4. Workers face a tradeoff when deciding when over their lives to enter self-employment. Workers who learn their entrepreneurial potential earlier have more working years in which to use that knowledge, but experience in paid work may increase entrepreneurial earnings. We find that the payoff to early experimentation dominates: the expected lifetime earnings gains to entering entrepreneurship fall with age, with most of the decline coming after age 50, as retirement approaches.
5. Policies that encourage entry into self-employment are likely to be most effective for raising average earnings when targeted to those at the top of the paid earnings distribution. Earnings

⁵This mixed result stems from a combination of factors: some adverse selection in who enters self-employment, along with longer run positive selection because successful entrepreneurs are more likely to remain in self-employment.

⁶ This comparison does not include costs to enter self-employment, in keeping with prior studies that compare self-employment and paid earnings. In the data, average monetary investments at entry are small.

⁷Within our model, these preferences are best interpreted as capturing a mix of true tastes for self-management and un-modeled heterogeneity in opportunities, access to non-job-linked health insurance, etc.

in paid work and entrepreneurship are highly positively correlated and the entrepreneurial earnings distribution has a relatively thick right tail. As a result, while lower-earning workers are the most responsive to fixed-dollar subsidies, the biggest gains come from encouraging experimentation among the most able workers. Flat taxes are effective for inducing more high-ability paid workers to experiment with self-employment, but substantially reduce overall tax revenue.

The next section discusses how these findings fit into the literature. Section 3 describes the data and section 4 presents stylized facts to motivate considering returns to entrepreneurship in a dynamic context. Section 5 presents the full model while section 6 describes the estimation algorithm. Section 7 presents our estimates and discusses fit, section 8 characterizes the value of entering entrepreneurship, and section 9 evaluates counterfactual policies. Finally, section 10 concludes.

2 Literature and Motivation

Evans and Leighton (1989) and Hamilton (2000) use survey data to document that the median self-employed worker earns less than the median paid worker after controlling for worker characteristics that are included in typical earnings regressions.⁸ These and other subsequent papers (see, for example, Hurst et al. (2014)) propose a variety of plausible explanations for this puzzling finding, including non-pecuniary benefits from being one's own boss, adverse selection (on difficult-to-observe traits) into self-employment, and more opportunities to manipulate or under-report income. Daly (2015) and Manso (2016) both emphasize that cross-sectional estimates of entrepreneurial earnings are biased because of dynamic concerns. These papers use a matching on observables approach to show that realized pre-tax lifetime earnings in the PSID and NLSY79, respectively, are higher for workers who spend at least some time as entrepreneurs than for those who never work for themselves.

⁸A related literature (Moskowitz and Vissing-Jorgensen (2002) and Hall and Woodward (2010)) finds even larger gaps between the returns to private businesses and investments in publicly traded equities, despite the higher risk of these largely undiversified business holdings.

Jovanovic (1979) and Miller (1984) develop models of how workers move between jobs and occupations to learn about their relative skills. Miller (1984) is the first paper to estimate a model where individuals choose jobs according to the optimal order from a multi-armed bandit process. In addition to estimating the model, he lays out implications for the duration of occupational spells, which are borne out for self-employment spells in the PSID. Subsequent papers expand on this learning model. Gibbons et al. (2005) estimate wage processes across industries or sectors taking seriously that workers may learn strategically. Moscarini (2005) integrates search costs into these models of experimentation and Papageorgiou (2014) finds, as we do, that skills are correlated across sectors. Vereshchagina and Hopenhayn (2009) derive the potential value of learning about entrepreneurial ability in a dynamic model of business development. We adapt a tractable experimental framework to the unique context of self-employment, including modeling entry costs.

Relative to Daly (2015) and Manso (2016), our more structural approach allows us to isolate the option value of experimenting in entrepreneurship from other selection concerns and to characterize how the value of the option changes over the lifecycle. Few other papers examine returns to self-employment by age, although a related literature looks at rates of entrepreneurship by age (Liang et al., 2014) and other papers look at creativity or innovation by age (Jones, 2010). Our model also allows us to consider the likely effects of some counterfactual policies. The findings in Daly (2015) and Manso (2016) on the realized lifetime earnings of workers reinforce our estimates that rely on projecting lifetime earnings from an unbalanced panel. Like them, we find that, on balance, cross-sectional earnings comparisons understate the expected value of entering self-employment by failing to account for this option value. In contrast to other studies, we also account for taxes, including progressive income taxes and the differential payroll tax treatment of earnings in self-employment.

The policy shifts we consider are highly stylized, but they correspond in spirit to several realized policy changes. Within the context of our model, a subsidy to the utility cost of entering self-employment might represent a loosening of borrowing constraints, such as the mortgage reform studied in Jensen et al. (2014), or a guaranteed option to return to one's old job, such as the maternity leave extension studied in Gottlieb et al. (2016). Both papers find that lowering the cost of entering self-employment, or re-entering the paid workforce from self-employment, generates

moderate increases in the share of workers who become self-employed.

Bruce (2000) surveys the literature on the effects of tax reform on self-employment rates, starting with Blau (1987). Most of these papers find that decreases in marginal tax rates decrease self-employment. These empirical studies face a confounding factor: in addition to any formal differences in the treatment of income from paid work and self-employment, self-employment offers more opportunities to reduce tax liability by claiming deductions or simply under-reporting earnings. Our model has no avenues for earnings manipulation, and we find that in this context reducing the progressivity of income taxes increases self-employment because it increases the potential upside of experimentation.

A related strand of the literature considers the characteristics and circumstances that drive entry into self-employment, while a smaller set of papers considers the determinants of exit. Caliendo and Uhlenhorff (2008), Caliendo et al. (2010), Levine and Rubinstein (2013), and Hurst and Pugsley (2011) discuss the heterogeneity of self-employed workers and the reasons they select into self-employment. We deal with the determinants of entry in an abstract way, but show that learning about entrepreneurial ability appears to be important for a wide range of workers, and that this process is not strongly affected by the path through which workers become self-employed. In line with several of these studies, we find some adverse selection into entrepreneurship by earnings ability. Kahn et al. (2016) find that, among immigrants, this net adverse selection into self-employment occurs only in non-science occupations. Science entrepreneurs are positively selected on paid earnings potential. On the exit margin, Sarada (2015) shows that those who persist in self-employment have increasing consumption and savings, consistent with selection of the best entrepreneurs, while Hamilton (2000), Evans and Leighton (1989), and Bruce and Schuetze (2004) provide some mixed evidence on the return to self-employment experience in paid work.

3 Data and Sample Construction

3.1 The Sample

The data are from the 1976-2011 waves of the Panel Study of Income Dynamics (PSID). The PSID began interviewing a sample of households in 1968 and has continued to interview those households and their descendants annually until 1997 and bi-annually from 1999 onward. The long panel structure of the PSID allows us to observe workers before, during, and sometimes after spells of entrepreneurship. Updates to the survey sample, interview non-response, and lifecycle transitions as older workers retire and young workers enter the labor force combine to create an unbalanced panel.

Because we need to keep track of lifetime accumulated work experience in each sector, we include individuals in the sample only when we can follow their work experience starting at age 25 or earlier. We also restrict the sample to men. Women in the PSID are far more likely to persist in self-employment despite low relative earnings, suggesting a larger role for non-pecuniary considerations or a fuzzy divide between self- and non-employment. It seems clear that men and women should not be pooled to estimate a single model. Women’s entrepreneurial decisions are explored in Lim (2015) and remain a fruitful area for future research. We discuss sample construction and variable definitions in more detail in the Data Appendix. After restrictions, the sample includes just under 7,000 men. On average, we observe 10 years of earnings for each individual in the sample, with a maximum of 29 observations spanning 36 years.

3.2 Identifying Entrepreneurs

We define an entrepreneur as someone who is self-employed in their main job. We choose this definition, rather than one based primarily on business ownership, because entrepreneurship here is modeled as a labor supply choice that may also represent a financial investment, rather than primarily as an investment choice.⁹

Table 1 describes workers who ever spend time in entrepreneurship and those who only ever

⁹Light and Munk (2015) investigate the difference between self-employment and business ownership in more detail.

work in the paid sector. Observing entrepreneurship is quite common. Of those with experience in entrepreneurship, 82% own a business at some point during the sample. Workers who report being self-employed without owning a business are generally contractors or work in low-physical capital occupations like consulting.¹⁰ Entrepreneurs are more likely to be white and slightly better educated on average than the paid-only group, but generally look similar to other workers. Moves in and out of entrepreneurship are frequent. The average member of the sometime entrepreneur sample is observed for 12 years in paid work and only 9 years in entrepreneurial work.

Figure 1 plots the cumulative transitions back to the paid sector after entering self-employment. Overall, 40% of workers who experiment with entrepreneurship return to the paid sector within five years, 25% return after only one year. Experienced entrepreneurs are substantially more likely to own a business and that business is more likely to be incorporated, which Levine and Rubinstein (2013) find is a good proxy for business sophistication and profits. The prevalence of business ownership and incorporation increases with experience in self-employment, both because workers who enter self-employment as business owners are more likely to persist and because some workers who begin without a business establish one several years into self-employment. This pattern of "upgrading" into business ownership is suggestive of another dimension of experimentation; workers sample self-employment without paying the costs of incorporation and adjust their status after learning more about their success.

3.3 Earnings in Paid Work and Entrepreneurship

In the PSID, paid workers are asked about their total labor earnings from wages, salaries, bonuses, commissions, and tips. Owners of incorporated businesses are also asked about total labor earnings, but owners of unincorporated businesses are instead asked about their net profit from that business. We measure annual earned income as reported total labor income when possible and fill in with net business profit when necessary. We discuss the potential concerns with this mixed measure in the Data Appendix.

¹⁰21% of workers who always work in the paid sector also own businesses in at least one year, but these businesses are run on the side while working for someone else on the main job.

Workers' total take-home earnings are the relevant measure for the choice between paid work and self-employment. Taxes create a wedge between total reported earnings and the money that workers take home.¹¹ Accounting for this wedge is of particular importance because U.S. payroll taxes differ for paid and self-employed workers and because progressive income taxes dampen the upside risk of experimenting with entrepreneurship. We work with pre-tax estimates when modeling the dynamics of earnings, but adjust by a stylized tax schedule when modeling how workers make their labor sector choices. We construct these expected post-tax earnings in each sector using an approximation of the average U.S. tax schedule over the sample period. We describe the details of this procedure in Appendix B.

Both pre- and post-tax earnings in entrepreneurship are more variable than earnings in the paid sector. Figure 2 plots the distribution of real weekly pre-tax earnings, Panel A, and post-tax earnings, Panel B, in each sector. The distribution of earnings for workers currently in entrepreneurship is flatter than for paid workers, with more weight on the lowest values and a thicker long right tail. Using the Survey of Income and Program Participation, Hamilton (2000) finds that the mean of pre-tax earnings in entrepreneurship is somewhat higher than the mean in paid work, but the median is lower. We can replicate his finding using our PSID sample, but find that both mean and median earnings are lower in self-employment after accounting for taxes.

4 Learning and Selection Out of Entrepreneurship

Figure 1 shows that moves from entrepreneurship back to the paid sector are quite common. Figure 3 illustrates that these transitions back to the paid sector are not random. The top line of this figure plots the median post-tax earnings in each year of self-employment for workers who remain self-employed for 6 or more years. This profile is strictly above the median earnings by years of

¹¹Under-reporting by the self-employed motivates much of the literature on taxation and entrepreneurship. Hurst et al. (2014) evaluate this under-reporting in household surveys. According to their findings, the PSID survey data is likely to be unbiased for paid earnings, taxes on paid earnings, and total self-employment taxes that were remitted to the government (as a consequence of using reported earnings in self-employment). Actual self-employment earnings may be higher than reported. We do not correct for under-reporting, in keeping with much of the prior literature on entrepreneurial earnings. Instead, our estimates highlight the difference between lifetime and reported cross-sectional earnings from self-employment.

experience for workers who leave within 5 years, which is in turn higher than median earnings in self-employment for workers who leave after one year. Workers who persist longer in entrepreneurship earn more in that sector from their first year. The same pattern holds in Figure A1, which plots the same series of observed earnings in entrepreneurship, now relative to workers' last observed earnings in paid work.¹² In all groups, the median entrepreneur earns less in his first year of entrepreneurship than he did in his last year of paid work. Those who remain in entrepreneurship for at least 6 years eventually earn more than their lagged paid earnings, while individuals who return to the paid sector within 5 years continue to earn less at the median.

This negative correlation between entrepreneurial earnings and the probability of leaving entrepreneurship is consistent with a learning model, but also with some alternative models. For example, workers might know their abilities with certainty, but choose their sector based on both their known potential earnings and unobserved and variable preferences for self-employment. Workers who are drawn into self-employment primarily through a transitory positive preference shock would be both less likely to have high entrepreneurial earnings and more likely to exit quickly. Because both stories are consistent with the same patterns of earnings and transition data, they are difficult to disentangle within a reasonably tractable structural model. Instead, we now present some reduced form evidence in favor of a learning model, which will motivate our modeling choices in the next section.

First, as shown in Figure A1, the majority of workers earn less in their first year of self-employment than they did in their last year of paid work. This fall in earnings suggests imperfect information about self-employment earnings or a willingness to take on low earnings in exchange for potential future earnings growth. Table 2 reinforces that entrepreneurial earnings are difficult to predict in advance. For precision and ease of interpretation, we use estimates of individual earnings fixed effects in each sector (roughly, average residual earnings after netting out experience profiles) rather than raw average earnings. We describe the estimation of these fixed effects, using methods that do not rely on any particular model of selection between sectors, in Section 6. The first column

¹²This plot of relative earnings is noisier for two reasons: it presents the ratio of earnings in two years, which compounds measurement error, and it uses a smaller sample because some workers are self-employed from the beginning of the sample (17% of the sample used in Figure 3) and others are missing earnings shortly before transitioning to self-employment, either because of time out of work or a missed survey (a further 11%).

of Table 2 presents the results of a regression of the log individual fixed effect in self-employment, estimated using all years of observed entrepreneurial earnings, on each worker’s log earnings fixed effect in paid work. The coefficient on the paid fixed effect is large and precisely estimated. Alone, it explains 44% of the variation in the self-employment fixed effect.

The next four columns add additional information that might be expected to predict earnings in entrepreneurship. Evans and Leighton (1989) show that some workers may become self-employed because they are unable to find good work as employees.¹³ As expected, being “pushed in” to self-employment by a sharp fall in paid earnings or a long spell of unemployment is associated with lower earnings in self-employment, 20% lower on average. If workers know their abilities in advance, then workers with higher expected earnings should be willing to invest more when starting their business. The fourth column of Table 2 shows that forming a business within a year of entering self-employment, particularly an incorporated business, is associated with higher average earnings in all years of self-employment.¹⁴ However, the amount invested in that business over the early years is not a strong predictor of entrepreneurial earnings. Work experience in the same industry prior to entering self-employment has a similarly weak relationship with entrepreneurial earnings. The full set of covariates explains only 7% more of the variation in entrepreneurial earnings than the paid earnings fixed effect alone. Overall, it appears that individuals can use their past earnings in paid work to predict how much they will earn in self-employment relative to other self-employed workers, but it is difficult to predict entrepreneurial earnings relative to one’s own paid earnings, which is the relevant comparison when deciding whether to enter.

We also consider some alternative explanations for the patterns of who leaves self-employment. Table 3 presents estimates from a series of Cox proportional hazards models on the probability of returning to paid work in each year for workers who were self-employed in the previous year. The first column estimates the effect of expected self-employment earnings for the coming year on the probability of exiting that year.¹⁵ Consistent with Figure 3, workers who expect to earn more in

¹³We also find evidence of this kind of behavior in the PSID; workers who experience a persistent fall in average weekly earnings of at least 30% are more than twice as likely to enter self-employment the following year as those who experience earnings growth or smaller losses (3.8% chance versus 1.6%).

¹⁴As shown in Figure 1, some workers report being self-employed without owning a business.

¹⁵These expected earnings, which are based on the history of earnings in self-employment, are constructed using the estimates described in Section 6.

self-employment are less likely to exit. Table 3 presents hazard ratios, so the third row of the first column indicates that a worker in the highest quartile of expected self-employment earnings is about a third as likely to exit self-employment as a worker in the lowest quartile of expected earnings, the reference group. The second column shows that the ratio of individual expected earnings in self-employment to expected earnings in paid work is an even stronger predictor of exit.

This relationship between earnings and exits could simply reflect the path by which workers entered self-employment. Workers who experienced large negative shocks to paid earnings just before becoming self-employed earn less in self-employment and may be more likely to exit if they became self-employed by necessity rather than choice. However, the third column of Table 3 shows that workers who were initially pushed into self-employment are only 10% more likely to exit, and this relationship is not statistically different from zero. Moreover, once these pushed-in workers enter self-employment, their exit decisions, conditional on earnings, look similar to those of other entrepreneurs. Workers who enter the sample self-employed, which implies that they have been self-employed since at least age 25, are substantially less likely to leave self-employment, which could reflect a strong preference for self-employment or perhaps an inherited family business.

Finally, workers with low earnings may exit self-employment not because these low earnings give them new information about future earnings, but because they do not have enough capital to weather a negative shock and are forced to close their business. In this case, workers in more capital-intensive industries should be more likely to respond to low earnings by leaving self-employment. The last column of Table 3 shows no strong pattern between capital intensity and the probability of exit.¹⁶ The hypothesis that exits are driven by learning cannot be rejected.¹⁷

Once workers have entered entrepreneurship and observed their earnings in that sector, those with low relative entrepreneurial earnings are far more likely to return to the paid sector. This pattern is constant for many self-employed workers, regardless of industry or reason for entering self-

¹⁶The cutoff for a high-capital-intensity industry is one where at least 25% of workers invest at least \$25,000 when starting their business, using industry-specific averages from the 2007 Survey of Business Owners.

¹⁷In addition, unobserved capital constraints do not appear to be the primary cause of exits from entrepreneurship as reported in surveys of business owners. Data from the 2007 Survey of Business Owners (SBO) Public Use Microdata sample suggests that the minority of exits from entrepreneurship are due to capital constraints. Of those who report that their past business no longer operates, 7.9% cite lack of business credit and 4% cite lack of personal credit as the cause of lack of operations. In contrast, 29% of these respondents cite insufficient sales or cash flow, suggesting the business failed to materialize.

employment. We see much weaker evidence that workers who have not yet worked as entrepreneurs have any advance knowledge of their potential earnings in that sector. These patterns motivate our model.

5 Model

This section describes the semi-structural model that captures utility over earnings and preferences for entrepreneurship. We use the model to account for mobility patterns over the lifecycle when simulating the distribution of earnings and entrepreneurial choices.

In each period $t = 0, 1, \dots, T_i$, starting after the last year of schooling and continuing to retirement, risk-neutral individual i chooses between supplying labor in the paid sector ($d_{it} = 0$) or the entrepreneurial sector ($d_{it} = 1$). Once the sector is chosen, earnings shocks are realized and flow payoffs are received. This assumption means that the individual makes sectoral choices without knowing earnings in advance, and adverse shocks cannot be escaped before their realization. This is consistent with a world in which there is some earnings risk over the course of a period, perhaps from job loss or changes, or where earnings, even in the paid sector, depend partially on bonuses or profit sharing.

The assumption of risk neutrality is reasonable if workers have some ability to smooth consumption. The option to return to paid work limits individual downside risk from entering entrepreneurship to a few periods. With smoothing, low earnings in entrepreneurship today affect total consumption only through the lifetime budget constraint. As a result, and consistent with the data presented in Hurst et al. (2014), flow consumption changes are likely to be much smaller than the change in flow earnings for those who have low earnings in entrepreneurship. Model complexity and data issues make adding risk aversion difficult for this paper.¹⁸ The approach taken here can capture some aspects of risk aversion, among other factors, through heterogeneous non-pecuniary benefits from entrepreneurship and from the utility costs of entering self-employment.¹⁹

¹⁸Adding endogenous savings and consumption decisions greatly complicates the model solution. The PSID includes very limited data on assets, which limits our ability to discipline these choices in the model.

¹⁹Setting a relatively high real interest rate (10%) also means that there is an adjustment for the risk of future earnings.

The expected flow utility from choosing the paid sector, $d_{it} = 0$, is

$$u(d_{it} = 0, S_{it}, \varepsilon_{it}^0) = \beta_1 E[(1 - \tau_W) W_{it} | S_{it}] + \varepsilon_{it}^0, \quad (1)$$

where S_{it} summarizes the individual's employment history, sectoral experience, age, and beliefs about entrepreneurial ability at time t , W_{it} is wage earnings in the paid sector, $1 - \tau_W$ is a tax adjustment function for paid earnings described in Appendix B.2, and ε_{it}^0 is a transitory taste shock for choosing paid work that is unobserved to the econometrician. The parameter β_1 translates earnings into units of utility, scaled relative to the variance of the taste shock. As a normalization, the expected utility of working in the paid sector for a wage of zero is set to $E(\varepsilon_{it}^0)$.²⁰

The expected flow utility from choosing entrepreneurship, $d_{it} = 1$, is

$$u(d_{it} = 1, S_{it}, \beta_{0i}, \varepsilon_{it}^1) = \beta_{0i} + \beta_1 E[(1 - \tau_R) R_{it} | S_{it}] + \beta_2 (d_{it-1} = 0) \alpha_i + \beta_3 (x_{Rit} = 0) \alpha_i + \varepsilon_{it}^1. \quad (2)$$

The components of utility in entrepreneurship include an unobserved taste or opportunity shock, ε_{it}^1 , utility from expected entrepreneurial earnings, R_{it} , adjusted for taxes on self-employment earnings through the function $1 - \tau_R$, flow utility from working in entrepreneurship, β_{0i} , and costs to entering entrepreneurship that are proportional to baseline wages in paid work, α_i . β_{0i} is a random parameter that varies across individuals, but is fixed for each individual over time. Individuals with $\beta_{0i} > 0$ are, all else equal, willing to give up some earnings to be self-employed; individuals with $\beta_{0i} < 0$ will only work for themselves if they expect to earn more in self-employment than their projected paid earnings.

²⁰To keep the model simple, we abstract from savings, eliminating intertemporal reasons that wage earnings do not equal consumption. This is innocuous under risk neutrality. When comparing paid and entrepreneurial earnings, we also abstract from the distinction between business and personal taxes except for the difference in the payroll tax. Measuring effective taxes for business owners is difficult in the PSID; however, this issue is unlikely to have a major effect on the estimates. The parameter β_1 is identified based on within-individual differences in expected earnings across sectors, up to scale. To the extent that differences in the tax treatment of earnings across sectors do not vary non-linearly with earnings levels, interpreting the scale of β_1 becomes difficult but there are no spillovers for the other parameters. Finally, note that the effective tax rate for most small business owners is very similar to individual income taxes; many small business are taxed as pass-through entities, equalizing tax schedules across paid work and self-employment. Businesses set up as C corporations face a different tax schedule on any earnings that are not distributed as salary; these earnings are subject to the corporate tax rate, and distributions are subject to shareholder taxes.

The parameter β_2 captures the utility cost of entering entrepreneurship. Similar to Evans and Jovanovich (1989), who model constraints on borrowing as a function of wealth, we assume that entry costs are proportional to baseline paid earnings ability, which is a proxy for expected lifetime wealth. These reduced-form entry costs may capture many constraints to entering entrepreneurship, including capital investments, effort costs of creating a new venture, the lost earnings associated with changing jobs, or the ex-ante psychological costs of uncertainty about a new enterprise. Some of these costs may be greater for workers entering self-employment for the first time. β_3 captures any additional costs for entrants with no prior entrepreneurial experience, $x_{Rit} = 0$.

We specify a flexible parametric model for earnings in each sector. If the agent is employed in the paid sector, his earnings depend on his accumulated work experience in the paid sector, x_{Wit} , and in entrepreneurship, x_{Rit} , along with fixed individual earnings ability in the paid sector, α_i , a transitory shock M_t , and a log AR(1) persistent shock, $P_t = P_{t-1}^\phi \zeta_t$. Paid earnings are given by

$$W_{it} = \exp(\alpha_i + G_W(x_{Wit}, x_{Rit})) P_{it} M_{it}. \quad (3)$$

The shocks ζ_{it} and M_{it} are distributed log-normally, with $\ln \zeta_{it} \sim N(0, \sigma_\zeta^2)$ and $\ln M_{it} \sim N(0, \sigma_M^2)$. Individual log earnings ability is also drawn from a normal distribution, $\alpha_i \sim N(\mu_\alpha, \sigma_\alpha^2)$, and we assume that individuals know this ability with certainty at the time they enter the model. While this assumption implies an asymmetry in information between sectors, it can be reframed as an assumption that workers have had sufficient informal work experience prior to entering the paid sector to deduce their ability with very little uncertainty.²¹

The tax adjustment, where a polynomial is used to compute expected taxes over different earnings levels, is described in the appendix. Many readers, however, will be familiar with the formula for the expected pre-tax earnings in each sector. For the paid sector in period t , pre-tax earnings are

$$E[W_{it}|S_{it}] = \exp \left[\alpha_i + G_W(x_{Wit}, x_{Rit}) + \phi \log(P_{it-1}) + \frac{\sigma_\zeta^2 + \sigma_M^2}{2} \right] \quad (4)$$

where $\frac{\sigma_\zeta^2 + \sigma_M^2}{2}$ is the convexity adjustment from the first moment of the log normal distribution.

²¹Workers may also estimate paid earnings ability quite accurately upon receiving job offers.

The state variables necessary to calculate this expectation are lagged years of experience in the paid sector and entrepreneurship, α_i , and the lagged value of the persistent shock.

We assume that P_{it} continues to depreciate during periods when agents work in self-employment, but does not experience any new innovations. The persistent shock in the paid sector therefore influences moves in and out of entrepreneurship as agents with low P_{it} will find self-employment temporarily more attractive.

Earnings for entrepreneurs are described by

$$R_{it} = \exp [\eta_i + G_R(x_{Wit}, x_{Rit})] \xi_{it}, \quad (5)$$

where ξ_{it} is a log-normally distributed transitory shock, $\ln(\xi_{it}) \sim N(0, \sigma_\xi^2)$ and $\eta_i \sim N(\mu_\eta, \sigma_\eta^2)$ is log entrepreneurial ability. Expected earnings in self-employment differ from earnings in paid work in several respects. Most importantly, while agents are assumed to know α_i with certainty, they know only the distribution of η_i . We allow α_i and η_i to be correlated, in which case workers form their initial beliefs about η_i based on their known α_i and then refine those beliefs through experience in self-employment. In addition, earnings in entrepreneurship depend on only a transitory shock, with no persistent stochastic element. This assumption is partially practical; a persistent shock would substantially complicate the process by which agents use realized entrepreneurial earnings to update their beliefs about η_i . We also interpret the persistent shock in the paid sector as partially capturing employment shocks, which are less relevant in self-employment. Finally, accumulated experience affects earnings in each sector differently, through the function $G_R(\cdot)$ rather than $G_W(\cdot)$.

To reiterate, expected earnings in entrepreneurship depend on the beliefs about entrepreneurial ability. For individuals with no entrepreneurial experience, their prior belief is based on α_i . To pin down this relationship, α_i and η_i are assumed to have a bivariate normal distribution with correlation ρ ; this form fits the data well and provides a tractable way to calculate the conditional

distribution of η_i given α_i .²² This conditional prior is thus normally distributed with mean

$$\hat{\eta}_{i0} = \mu_\eta + \frac{\sigma_\eta}{\sigma_\alpha} \rho (\alpha_i - \mu_\alpha) \quad (6)$$

and variance $\sigma_{\eta 0}^2 = \sigma_\eta^2 (1 - \rho^2)$. For individuals with x_{Rit} years of entrepreneurial experience, the mean belief is denoted $\hat{\eta}_{ix}$ and is updated according to Bayes' rule. This yields

$$\hat{\eta}_{ix} = \frac{\sigma_\xi^2 \hat{\eta}_{i0} + x_{Rit} \sigma_{\eta 0}^2 \overline{\log(\tilde{R}_{it-1})}}{x_{Rit} \sigma_{\eta 0}^2 + \sigma_\xi^2} \quad (7)$$

where $\overline{\log(\tilde{R}_{it-1})}$ is the mean of the residual log earnings history in entrepreneurship from experience levels 0 through x_{Rit-1} , net of the experience profile in entrepreneurship, $G_R(\cdot)$. The variance of the prior distribution is updated in a deterministic fashion in each period. For $x_{Rit} > 0$, the variance of the prior is $\sigma_{\hat{\eta}_{ix}}^2 = \frac{\sigma_{\eta 0}^2 \times \sigma_\xi^2}{x_{Rit} \sigma_{\eta 0}^2 + \sigma_\xi^2}$, which declines with experience in entrepreneurship.

The risk-neutral agents in this problem care about the level of earnings. Expected pre-tax earnings in entrepreneurship are given by

$$E(R_{it}) = \exp \left[\hat{\eta}_{ix} + G_R(x_{Wit}, x_{Rit}) + \frac{\sigma_{\hat{\eta}_{ix}}^2 + \sigma_\xi^2}{2} \right]. \quad (8)$$

The process for taking post-tax earnings is provided in the appendix. As is clear from Equation 8, the expected flow value of entrepreneurial earnings is increasing in $\sigma_{\hat{\eta}_{ix}}^2$. Holding fixed $\hat{\eta}$, expected entrepreneurial earnings change in two ways as workers accumulate entrepreneurial experience. Experienced entrepreneurs are more certain about their skills, which lowers $\sigma_{\hat{\eta}_{ix}}^2$ and therefore expected earnings. However, more experience raises expected earnings through $G_R(x_{Wit}, x_{Rit})$.

Formally, person i chooses to work in the sector that maximizes the present value of expected utility,

$$V_{it}(S_{it}, \beta_{0i}, \varepsilon_{it}) = \max_{d_i} E \left[\sum_{\tau=t}^{T_i} \delta^{\tau-t} u(d_{i\tau}, S_{i\tau}, \beta_{0i}, \varepsilon_{i\tau}) | S_{it} \right], \quad (9)$$

where δ is the discount rate and d_i is the state-contingent sequence of choices made by the individual

²²Formal tests of this assumption and sensitivity are discussed in Appendix B.

to maximize the present value of utility.

6 Model Solution and Estimation

The previous section describes a dynamic discrete choice model with unobserved heterogeneity in the taste parameter β_{0i} . Non-parametric point identification of β_{0i} (up to the other parametric assumptions imposed on the model) is not possible for the majority of workers who never take up entrepreneurship. We can bound the preferences of these workers from above by identifying the largest β_{0i} consistent with choosing the paid sector in all observed periods, but any preference below this threshold is equally consistent with the data. Because of that, we impose the distributional assumption that $\beta_{0i} \sim N(\mu_{\beta_0}, \sigma_{\beta_0}^2)$. This added discipline rules out thick tails in the preference distribution but allows flexibility through $\sigma_{\beta_0}^2$.

To estimate the model we must first define the expected value to workers of choosing to work in each sector. Because choices today affect future payoffs, the alternative-specific, or conditional, value functions include a flow-utility term and a continuation value. These conditional value functions describe the present value of the agents' problem at time t conditional on choosing sector d_t and then following an optimal strategy in the future. Omitting i subscripts to conserve notation except to make clear heterogeneity in parameters, the lifetime maximization problem in equation (9) can be rewritten as a sequence of single-period decisions using the Bellman equation,

$$V(S_t, \beta_{0i}, \varepsilon_t) = \max_{d_t \in \{0,1\}} \{u(d_t, S_t, \beta_{0i}, \varepsilon_t) + \delta E[V(S_{t+1}, \beta_{0i}, \varepsilon_{t+1}) | d_t, S_t, \beta_{0i}]\}. \quad (10)$$

The value of arriving at time t with state variables S_t , preferences β_{0i} , and shocks ε_t is the maximum of the conditional value functions. Recall that ε_t is a vector of sector-specific taste shocks that are iid across sectors and over time from distribution $g(\varepsilon_t)$. Define $f(S_{t+1}|S_t, d_t; \theta)$ as the transition density function describing the evolution of the observed state variables, parameterized by θ . If β_{0i}

is known, the conditional value function is therefore

$$\begin{aligned} v(d_t, S_t, \beta_{0i}, \varepsilon_t; \beta) &= u(d_t, S_t, \beta_{0i}, \varepsilon_t; \beta) + \delta E[V(S_{t+1}, \beta_{0i}, \varepsilon_{t+1}) | d_t, S_t, \beta_{0i}] \\ &= u(d_t, S_t, \beta_{0i}, \varepsilon_t; \beta) + \delta \iint V(S_{t+1}, \beta_{0i}, \varepsilon_{t+1}) dg(\varepsilon_{t+1}) df(S_{t+1} | S_t, d_t; \theta). \end{aligned} \quad (11)$$

The flow utilities from choosing each sector, described in equations (1) and (2), depend on the common parameters β and the heterogeneous parameter β_{0i} .

The probability of observing $d_t = j$, where $j \in \{0, 1\}$ indexes labor sector choice, is given by integrating over the unobserved taste shocks ε_t

$$p_t(d_t = j | S_t, \beta_{0i}; \beta) = \int I \left\{ \arg \max_{d_t \in \{0, 1\}} v(d_t, S_t, \beta_{0i}, \varepsilon_t; \beta) = d_j \right\} dg(\varepsilon_t). \quad (12)$$

The individual's likelihood contribution at time t is therefore

$$\mathcal{L}_t(d_t, S_{t+1} | S_t, \beta_{0i}; \theta, \beta) = p_t(d_t | S_t, \beta_{0i}; \beta) f(S_{t+1} | d_t, S_t; \theta). \quad (13)$$

However, since the individual's fixed preferences are also unobserved we instead use the marginal likelihood of workers' choices by integrating over the distribution of β_{0i} . Using this approach to recover parameters, we solve

$$\left(\hat{\theta}, \hat{\beta}, \hat{\mu}_{\beta_0}, \hat{\sigma}_{\beta_0}^2 \right) = \arg \max_{\theta, \beta, \mu_{\beta_0}, \sigma_{\beta_0}^2} \sum_{i=1}^N \log \left[\int \prod_{t=1}^{T_i} \mathcal{L}_t(d_{it}, S_{it+1} | S_{it}, \beta_{i0}; \beta, \theta) d\phi \left(\frac{\beta_{0i} - \mu_{\beta_0}}{\sigma_{\beta_0}} \right) \right]. \quad (14)$$

The parameters to be estimated are $\hat{\theta}$, which describe the determinants of earnings and the transitions of the observed state variables, $\hat{\beta}$, the parameters of the flow utility functions, and $\mu_{\beta_0}, \sigma_{\beta_0}^2$, the parameters of the distribution of non-pecuniary benefits in entrepreneurship. From equation (11), the likelihood for individual i at time t depends on the continuation value of each choice, so each calculation of the likelihood involves solving the full lifecycle model. Estimating some parameters in a first step eases the computational burden by minimizing the number of parameters that must be estimated within the full solution.

Further detail about the first and second stage of estimation is contained in Appendix B.

7 Estimates

7.1 Determinants of Earnings

The estimates from the first stage are presented in Table 4. Later, we consider sensitivity of the results to these estimates. These estimates describe the evolution of pre-tax earnings. Earnings in the paid sector rise by 5% on average with the first year of paid experience and 42% with ten years of experience. Accumulated paid experience is also associated with higher earnings in self-employment. A worker who accumulated ten years of paid experience before entering self-employment earns 20% more on average in self-employment than a worker who entered with no paid experience. Entrepreneurial earnings rise with entrepreneurial experience, but not as rapidly as paid earnings rise with paid experience. Entrepreneurs with ten years of self-employment experience earn 23% more than new entrepreneurs on average.

The next panel of Table 4 describes the distributions of individual fixed effects for log earnings in each sector. The mean of earnings in entrepreneurship is slightly lower than mean paid earnings. A worker at the mean of the paid ability distribution who has no experience would expect to earn 1.8% less in entrepreneurship than in paid work pre-tax, a difference of about \$600 if he worked all year. However, this difference increases to about \$3,000 after accounting for the differential payroll treatment between self-employment and paid work. Workers who have worked in the paid sector will face larger average drops in pre-tax earnings upon entering entrepreneurship because of the smaller returns to paid experience.

Ability in entrepreneurship has more than twice the variance of ability in the paid sector, 0.63 compared to 0.19. Abilities in the two sectors are strongly correlated, with $\rho = 0.7$, so the variance of a worker's belief about his own entrepreneurial ability is smaller than the population variance. With no experience in entrepreneurship, the variance of an individual's belief of his entrepreneurial ability is $\sigma_\eta^2(1 - \rho^2) = 0.32$. A standard-deviation in the initial forecast of entrepreneurial ability is equal to a change in pre-tax earnings of about 56% at the mean.

The relative variance of entrepreneurial ability and the transitory shock to entrepreneurial earnings, the last panel of Table 4, imply that this uncertainty dissipates quickly as workers gain entrepreneurial experience. The variance of a worker’s belief about his ability falls to 0.07 after observing one year of entrepreneurial earnings and 0.02 after five years. This correlation and the remaining variation indicate that workers may predict their average entrepreneurial earnings based on paid earnings, but there is significant variation around that prediction, creating the basis for learning about entrepreneurial earnings.

The bottom two panels of Table 4 present the stochastic parameter estimates from the first stage. We estimate that 83% of the persistent shock in the paid sector remains after a year. The transitory shock to earnings in entrepreneurship is larger than the combined shocks in the paid sector, but the main reason workers face more uncertainty about entrepreneurial earnings is the variance of permanent ability, η .

7.2 Determinants of Utility

This section discusses the flow payoff parameters estimated in the second stage, which makes use of the parameters described in Section 7.1. Recall that these parameter estimates come from maximizing the probability of choices that individuals make. The choice to enter entrepreneurship from paid work is somewhat rare, about 2% of paid workers transition to self-employment each year,²³ yet the average difference in flow pay is only a small fraction of lifetime income. If experimenting with entrepreneurship is a financially attractive choice, then the model will fit the low observed entry rate into entrepreneurship through entry costs into entrepreneurship, the distribution of non-pecuniary benefits, or the scaling of money relative to the taste shock.

Table 5 presents the flow payoff parameters estimated in the second stage of estimation. The first column presents estimates from our baseline model with heterogeneous tastes for entrepreneurship, while the second presents the estimates that result from shutting down any preference heterogeneity. The scale of utility is normalized by the Type-1 extreme value taste shocks, which have a standard deviation of $\frac{\pi}{\sqrt{6}}$. The parameter estimates in the top panel are therefore not easily interpretable.

²³We report these observed transition probabilities, along with the predictions of our model, in Table 6.

The second panel provides transformations of these parameters into 2010 dollars equivalents, using the coefficient on after-tax earnings in the flow utility function, β_1 .

The model allows the cost of entering self-employment to differ for repeat entrepreneurs and workers entering self-employment for the first time. In our baseline model, returning entrepreneurs face a utility cost of entering self-employment equivalent to a one-time payment of about \$325,000. First-time entrepreneurs pay an additional \$21,000, for a total entry cost of just over \$346,000.²⁴ These entry costs are larger than the average monetary investments reported in the PSID. The estimated utility entry costs must therefore capture far more than the direct financial investments in new businesses. Entering self-employment involves additional financial costs in terms of foregone earnings during the transition. These utility costs likely also capture the stress and pressure of developing a new business, a low arrival rate of business ideas, and, since we do not explicitly model risk-aversion, a distaste for the large initial uncertainty about entrepreneurial earnings. Because we do not model any costs of returning to paid work, these estimated entry costs will also capture the risk of incurring a second set of transition costs if self-employment proves unfruitful.²⁵

We estimate wide variation in tastes for entrepreneurship. The mean worker would accept \$76,377 less in annual paid earnings to avoid working as an entrepreneur, slightly more than the mean annual earnings for a 30 year old. This estimate, along with high entry costs, is driven by the large share of workers who never enter self-employment despite potential gains to experimentation. However, the standard deviation of preferences is equally large; many workers have a far smaller distaste for entrepreneurship and 15% of workers derive a positive flow utility from working for themselves. The transitory taste shocks are also fairly large, with a standard deviation of \$89,877. These shocks perturb the expected lifetime value of choosing to work in each sector, including the discounted lifetime earnings stream. This standard deviation represents 10% of the average expected value of choosing paid work for 30 year olds.

²⁴If we ignore taxes, the estimated value of money falls relative to the non-pecuniary components of utility, inflating the monetary equivalent of the entry costs and flow preferences for self-employment. Without progressive taxes, workers appear to be leaving more money on the table by not trying self-employment to see if they land in the right tail of earnings.

²⁵We rely on exits from entrepreneurship to identify the distribution of tastes for entrepreneurship. We cannot identify these tastes and a separate exit cost. Alternative models could enforce equal transition costs for entering and exiting self-employment or omit exit costs altogether. Both do a poor job of fitting the data.

When we shut down this preference heterogeneity, the most conspicuous change is a far different mean taste for entrepreneurship, \$624 per year rather than -\$76,377. Most individuals who persist in self-employment earn modestly less (after taxes) as entrepreneurs than they would in paid work. When all workers are constrained to have the same preferences, this pattern pushes the disutility from working in entrepreneurship upward to account for those who persist. The other changes in the second column stem from this key difference. When all workers are roughly indifferent between paid work and self-employment, smaller taste shocks are sufficient to prompt a return to paid work, so the standard deviation of the taste shock necessary to match observed transitions falls slightly. Since fixed tastes for entrepreneurship will no longer constrain any workers from entering entrepreneurship, the estimated entry costs rise. As shown in section 7.3, this simpler model cannot predict movements between sectors as well as the full model with heterogeneity.

The last panel of Table 5 presents the average means and standard deviations of the Bayesian posterior distribution of β_{0i} for each person in the sample given their likelihood of choices. Appendix B provides details about this calculation. There is a sharp divide between the posterior distributions for those who enter entrepreneurship and those who do not. For workers who enter entrepreneurship, the average mean of the posterior distributions is positive: \$1,600 per year. Our estimates are consistent with earlier studies, surveyed in Åstebro et al. (2014), that estimate positive tastes for entrepreneurship among workers who choose to work in that sector. In contrast, much of the rest of the population appears to have a strong distaste for entrepreneurship. Those who never enter self-employment have posterior preference distributions centered around a mean of -\$95,085.

The posterior preference distributions for all workers have smaller standard deviations than the population distribution, reflecting the added information from observed choices. This contraction is far more pronounced for workers observed in entrepreneurship, who have an average posterior standard deviation of \$31,002, than for the never entrepreneurs, with an average of \$62,273. Entrepreneurs' tastes for self-employment are identified off both their initial decision to enter self-employment and their annual decisions to persist in self-employment or return to paid work, allowing for narrow posterior distributions. In contrast, preferences for never entrepreneurs can only be bounded above: they must dislike entrepreneurship at least enough to have never entered.

7.3 Model Fit

Unobserved heterogeneity and taste shocks play important roles in determining sector choices at the individual level, particularly moves into entrepreneurship. We now assess the ability of the full model to predict which workers will select into each sector and compare that performance with two simpler models. Table 6 presents the average model-predicted probability of choosing entrepreneurship, separately for workers who worked in each sector last year and who we observe choosing each sector this year.

Within each cell, the first row shows the prediction of a dynamic model where workers are forward-looking income maximizers who have no preferences for either sector and face no entry costs. Expected lifetime earnings are projected imposing sector choice this year and then assuming workers make the lifetime income-maximizing choice in all future periods.²⁶ Income maximization predicts some of the observed variation in sector choice; among workers who were in the paid sector last year, 58.7% of workers who choose entrepreneurship this year have higher projected lifetime earnings from that choice, compared to 54.5% of workers who remain in paid work. For workers who were entrepreneurs last year, 72.2% of those who continue in self-employment have higher projected earnings from that choice, compared to 48.6% of workers who return to the paid sector. While income maximization correctly predicts that workers who choose entrepreneurship are more likely to do so, it predicts far too much entrepreneurship in total. Overall, the income-maximizing model predicts that 55% of paid workers will choose to become entrepreneurs next year when only 2% do.

The next row in each cell presents the choice probabilities predicted by our utility-maximizing model with no heterogeneity in preferences. We construct these predictions using the parameters reported in the second column of Table 5. Under this model, high entry costs push down the share of paid workers who are predicted to move into entrepreneurship, but they do so for everyone. Workers who choose to move from the paid sector to self-employment this period have only a 3.7% predicted probability of doing so, not much higher than the 2.4% average predicted probability for

²⁶With no entry costs, workers have an incentive to enter self-employment to wait out any negative persistent shocks to paid earnings. To prevent frequent moves between sectors, in this exercise we impose that workers who leave self-employment for paid work cannot return for a second entrepreneurial spell.

those workers who remain in the paid sector.

The final row of each cell presents the predicted likelihood of selecting entrepreneurship using the estimates of the full model and the individual posterior preference distributions described in the previous section. These estimates incorporate differences in observed characteristics across workers and also differences in unobserved tastes, as revealed by workers' history of choices. We use each individual's full sequence of sector choices to estimate their posterior preference distributions, so the predicted likelihood of choosing entrepreneurship in year t is partially determined by the observed future choices of each worker. This row should therefore be interpreted not as a test of ex ante predictive power, but rather as a measure of how well the key features of our fairly simple model with time-invariant individual tastes for entrepreneurship and time-varying earnings shocks and beliefs about entrepreneurial ability can match observed choices. Incorporating these predicted individual tastes improves the model's predictive power considerably. We estimate that workers who move into entrepreneurship are more than five times more likely to make that choice than workers who remain in paid jobs. Workers who exit back into paid work are three times more likely to make that choice than workers who remain in self-employment (34% probability rather than 11.5%).

Overall, these comparisons yield four key insights. First, entry costs are essential to quantitatively match the flow of workers into entrepreneurship. Second, preference heterogeneity is essential to understand which workers enter self-employment. The first two rows, which use only observable worker characteristics, do a relatively poor job of separating those who select entrepreneurship from those who do not. This finding is consistent with Levine and Rubinstein (2013), who find that entrepreneurs look similar to paid workers in terms of education, aptitude, and family background.²⁷ Third, these utility components are far less important for understanding who leaves self-employment. The income-maximizing model does nearly as good a job as our full model in predicting which entrepreneurs will return to paid work each year. Finally, even our best model does a poor job predicting the exact timing of individual transitions. We predict that workers who move from paid work to entrepreneurship had a 12% chance of doing so that year, far higher than

²⁷They find that non-cognitive measures have some additional predictive power to identify future entrepreneurs. Those traits will be part of what we capture in our individual tastes for entrepreneurship.

the 2% chance assigned to workers who remain, but well below 100%. Workers who return to paid work had a 34% predicted probability of doing so. These remaining prediction errors suggest an important residual role for time-varying shocks. We characterize these iid shocks as taste shocks in our model, but they may also capture other time-varying unobserved states. For example, they may represent the arrival of an innovative idea or access to capital on the entry margin, or bankruptcy or other business destruction on the exit margin.

The model does a particularly good job of matching the aggregate rate of transitions in and out of self-employment over the lifecycle. The pace of learning about one’s entrepreneurial earnings potential is set by the relative variances of the fixed effect and the transitory earnings shock in self-employment. These parameters are estimated in the first stage, which does not make use of observed transition patterns between sectors. The parameters estimated in the second stage, time- and age-invariant preferences along with iid shocks, can do little to directly generate changes in choice probabilities by age and experience. The next two figures therefore reflect fit along dimensions not directly targeted by the model.

Figure 4 plots the probability of moving from the paid sector into self-employment by age. The data series plots the share of paid workers with no entrepreneurial experience who choose to enter self-employment at each age, while the model series plots the average predicted likelihood of selecting self-employment for these workers using the posterior taste distributions. The predicted likelihoods match the broad pattern of the data: entry rates are higher for younger workers, averaging over 2% for workers in their 20s, and fall as workers age, to an average of close to 1% for workers over 45. The fall in entry rates later in life reflects two factors. First, paid experience increases earnings in the paid sector more than it increases earnings in self-employment, so more experienced paid workers are less likely to improve their earnings by changing sectors. Second, the option value of learning one’s entrepreneurial ability is highest for younger workers, who have more remaining working years to make use of that information. In the data, the probability of becoming self-employed rises through the early 20s, peaking at 2.6% at age 27, while the median age of those observed in entrepreneurship is 32. The model predicts that the youngest workers are most likely to move to entrepreneurship.

Figure 5 plots the probability of returning to the paid sector by years in self-employment, again comparing predicted probabilities using posterior preference distributions to observed patterns. In our model, as in the data, the hazard rate of leaving entrepreneurship falls sharply with entrepreneurial experience. This decline reflects the sharp fall in uncertainty about entrepreneurial ability. 23% of workers leave entrepreneurship after 1 year, while only 14% of workers who remain after five years exit the following year. The exit rates predicted by our model closely follow the empirical hazard rates.

8 The Monetary Value of Entering Self-Employment

We now present estimates of the expected returns to entering self-employment, highlighting the bias in cross-sectional earnings comparisons from non-random selection into self employment and the importance of accounting for the possibility of experimentation. For much of this section, we focus on observed and projected earnings for 30 year olds.²⁸ We focus on one age to simplify the interpretation of averages, particularly when we consider forward-looking projections of earnings. Over our full sample, average earnings blend an unbalanced panel of workers with different levels of accumulated work experience, uncertainty about entrepreneurial earnings, and remaining years of working life. Later we will discuss estimates for all ages and how they vary over the lifecycle. In the PSID sample, 8.5% of 30 year olds were self-employed over the past year. The median self-employed 30 year old has 4.25 years of experience in that sector. 6% of 30 year olds who worked for pay during the last year have already had a spell of self-employment.

8.1 Selection into Entrepreneurship

The first panel of Table 7 describes observed earnings for 30 year olds who spent the past year working in each sector. Previous research finds that the mean self-employed worker earned more before taxes than the mean paid worker, but the median self-employed worker earned less. We

²⁸The comparison of mean earnings in the first panel of Table 7 is sensitive to the choice of age because in each year the sample of self-employed workers is small. We chose age 30 because it is close to the median age at which workers enter entrepreneurship, 32, and because the distribution of observed earnings is representative of the overall patterns for young workers. The remaining panels of Table 7 look very similar for any age between 28 and 35.

replicate this pattern in pre-tax earnings, but find that young self-employed workers earn less after tax at both the mean and the median. Mean self-employed earnings relative to paid increase when older workers are considered. This is because the set of workers in self-employment overweights workers with high relative entrepreneurial earnings; this bias becomes more severe for older workers who have had more years to learn their entrepreneurial earnings ability and sort accordingly. From age 34 onwards, mean after-tax earnings for workers who chose self-employment are higher than mean after-tax earnings for paid workers, although the median self-employed worker continues to lag the median paid worker.

The next panel of this table summarizes projected after-tax current-year earnings in each sector for all 30 year olds in the sample, based on the parameter estimates in Table 4. In the first panel, workers contribute to only one column each: entrepreneurs are included in the entrepreneurial average and paid workers in the paid average. In this second panel, all workers contribute to both averages. The difference between these first two panels quantifies the selection bias in cross-sectional earnings. The mean after-tax earnings for current entrepreneurs is \$39,093, \$508 more than the mean expected earnings in entrepreneurship across all workers. This gap indicates some positive selection in the entrepreneur sample. The selection bias is far larger and of the opposite sign at the median. The small but important oversample of revealed super-star entrepreneurs distorts the mean upwards in the first panel, but initial entry into self-employment is negatively selected, bringing down the median.²⁹ The median entrepreneur earns \$8,920 less than the median paid worker, but if all 30 year olds were self-employed median earnings would be only \$5,141 lower than if all 30 year olds worked for others.

8.2 The Option Value of Experimentation

These first two panels describe a single year of earnings. We now consider the importance of incorporating projected earnings in future years. The next two panels of Table 7 describe expected discounted lifetime earnings, conditional on choosing each sector today. To construct expected

²⁹Initial entry into self-employment is U-shaped in paid earnings ability; very low earners are the most likely to enter, followed by very high earners, with middling earners entering at the lowest rate.

lifetime earnings we simulate 250 possible earnings streams for each observation in the sample, starting with the current year and continuing to retirement at age 65. Expected earnings evolve following the estimates in Table 4; the simulated worker chooses where to work based on these earnings, the preference parameters in Table 5, and his individual posterior expected taste for entrepreneurship. We calculate the total discounted earnings stream and take the mean of these totals for each observation. To make these numbers comparable to one-year earnings, we convert the discounted sum of lifetime earnings to annual equivalents: the constant annual earnings that would generate the same discounted lifetime earnings as our estimates.

The third panel describes expected lifetime earnings in a static model where workers must remain in one sector for all future periods. Because earnings in both sectors rise with experience, these annualized projected lifetime earnings are all higher than the single-year projections in the second panel. The fourth panel of Table 7 describes expected lifetime earnings imposing sector choice today and then allowing workers to choose the sector that maximizes expected utility in all future years. The difference between this panel and the one above it measures the option value of entering self-employment. At the mean, 30 year olds expect to earn \$3,125 more per year by entering entrepreneurship with the option to return to the paid sector than by working as an entrepreneur in all future years. With this option to return to paid work, the mean expected lifetime earnings for workers who choose self-employment next period is slightly higher than mean expected earnings for workers who choose paid work, while median expected earnings remains slightly lower if workers choose self-employment.³⁰ Comparing the top panel of Table 7 to the fourth panel, we see that for young workers, simple cross-sectional comparisons of average post-tax earnings across sectors substantially understate the expected returns to entering self-employment.

³⁰When comparing these estimates to prior papers, it is important to note that, to our knowledge, no other work on the lifetime returns to self-employment accounts for taxation. Like earlier studies, we find that the median man can expect higher pre-tax earnings over his life if he enters self-employment than if he spends all his working years in paid jobs.

8.3 Heterogeneity in the Option Value

The averages presented in Table 7 mask considerable heterogeneity in expected earnings in each sector, expected gains from entering self-employment, and the expected value of the option to return to paid work. The largest determinant of individual expected gains is past entrepreneurial experience. Workers who have never been entrepreneurs have considerable uncertainty about their entrepreneurial ability. Their expected gains from entering self-employment are positive and are tightly distributed around the expected value of the option to experiment. Once workers have entered self-employment, that uncertainty resolves and the anticipated gains from entrepreneurship become more dispersed. Workers now know with additional confidence whether they can earn substantially more in entrepreneurship, substantially less, or about the same.

Figure 6 illustrates these patterns by plotting the difference between the static and dynamic expected lifetime returns to entering self-employment, separately for 30 year old workers with and without entrepreneurial experience. Under the joint distribution of worker fixed effects estimated in Section 7.1, workers without entrepreneurial experience have expected entrepreneurial earnings centered close to their current earnings in paid work. The option value captures the benefit of being able to bound potential losses from entering self-employment, by exiting, while retaining the potential gains. For 30 year olds, the annualized option value is distributed symmetrically around a mean of \$3,265. On average, 30 year olds with no entrepreneurial experience expect to earn 11% more over their lifetimes from entering self-employment with the option to return to paid work than from remaining in self-employment in all future years.

Experienced entrepreneurs have already discovered whether they earn more or less in self-employment relative to paid work. Workers with high realized entrepreneurial earnings do not value the option to return to paid work because they have no intention of exercising it. In fact, their expected earnings in the full dynamic model may be lower than the static model, because taste shocks may push them back to paid work despite higher earnings in self-employment. These workers fill in the left tail of the distribution of the “option value” for experienced entrepreneurs. For workers who have discovered that they earn less in self-employment, the option to return to

paid work is no longer simply a contingency plan but the clear income-maximizing path; they expect to earn substantially more if they are allowed to return to paid work. The mean difference in annualized expected lifetime earnings from entering self-employment with and without the option to return to paid work is \$2,339 for 30 year olds with past entrepreneurial experience. The standard deviation of this gap is \$12,174 for experienced entrepreneurs, more than four times the standard deviation for workers with no entrepreneurial experience.

So far we have focused on the value of entrepreneurship for 30 year olds. Table A1 replicates Table 7 averaging over all observations of workers at all ages. The qualitative patterns we find at age 30 hold for all other ages, but there are some subtle differences over the lifecycle. The earlier workers learn about their entrepreneurial ability, the more working years they have to take advantage of their knowledge about their higher-earning sector. This pattern is evident in Figure 7, which shows that lifetime earnings gains from entrepreneurship decline with age. The steepest declines come as the individual approaches the end of the planning horizon.³¹ Balancing this option value is the expected loss in earnings from foregoing some of the returns to one's accumulated paid experience. Once workers enter self-employment they will begin climbing the experience profile in that sector, but older workers will have to work for many years before they can confidently expect to earn as much in self-employment as they did as experienced paid workers. Together, these concerns explain why, in both the data and our simulations, transitions to self-employment are more common for younger workers.

8.4 Sensitivity to Earnings Parameters

Table 8 reports sensitivity of the estimates in Table 7 to changes in the parameters governing the distribution of earnings. The estimates are constructed after perturbing the first-stage parameters that we expect have the strongest influence on the lifetime value of experimenting in entrepreneurship. We then re-estimate the second stage parameters incorporating these adjusted first-stage parameters, and finally re-simulate lifetime earnings. Column 1 is analogous to results reported in the third panel of Table 7 while column 2 is analogous to the fourth panel.

³¹When comparing this figure to past results, note that these numbers are not annualized.

Shifting the mean and variance of the distribution of entrepreneurial earnings move the value of entrepreneurship in the logical directions. We estimate the mean of this distribution by comparing the earnings fixed effects for workers who we observe in both sectors. Our estimates will be off if workers have some prior knowledge of their earnings in self-employment relative to their earnings in paid work and choose their sector accordingly. Under this story, we are most likely to be over-estimating mean entrepreneurial earnings. Changes in mean self-employment earnings translate directly into expected relative earnings in self-employment in the *static* model. In the *dynamic* model, expected gains from entering self-employment still increase with the population mean of earnings, but the changes are far more muted. The modest benefits of entering self-employment that we estimate are driven by strategic mobility and are not very sensitive to the estimated population mean of entrepreneurial earnings.

As we would expect, changes in the variance of entrepreneurial earnings have a smaller effect on static gains but a larger effect on the dynamic gains in self-employment. In the static model, increasing the variance of entrepreneurial earnings increases uncertainty but has little effect on mean expected earnings because workers will be forced to absorb both the added upside and downside risk. In the dynamic model, workers can return to paid work if their realized earnings are low, so increasing the variance of entrepreneurial earnings has a substantial net positive effect on expected earnings from entering self-employment. Conversely, if the variance of self-employment earnings was 5% smaller than the baseline estimate, the average dynamic gain from entering self-employment would be negative. If workers select into self-employment based on some (accurate) belief of their comparative advantage, then we will under-estimate the variance of earnings ability because we will mainly observe the upper part of the distribution. This concern suggests that the estimate where the variance is 5% larger is the more relevant figure and that we are more likely to be understating the dynamic gains to entrepreneurship.

The final panel considers changes in the correlation of abilities across paid work and entrepreneurship. We estimate the correlation between earnings abilities in these two sectors by comparing the estimated earnings fixed effects for workers who we observe in both sectors. We will mis-estimate this correlation if that sample is somehow selected on the correlation between

abilities. This correlation affects the expected returns to entering entrepreneurship in two ways. First, increasing the correlation of abilities increases ex ante certainty about entrepreneurial earnings and therefore lowers the value of experimenting with self-employment. Second, increasing the correlation in abilities raises the value of entering self-employment for workers with high earnings ability in the paid sector, because they are now more likely to also be high-earning entrepreneurs, and lowers the value for low earners in the paid sector. These effects largely cancel out at the mean, and both static and dynamic gains estimates exhibit little sensitivity to the estimated correlation.

9 The Effects of Policy on Self-Employment and Earnings

The previous sections show that workers face ex ante uncertainty about their potential earnings in self-employment and that this uncertainty creates a substantial value, in expectation, of entering self-employment to learn more about one's skills as an entrepreneur. This finding has several policy-relevant implications. First, policies that induce more people to enter self-employment might improve average earnings by facilitating better sorting by comparative advantage. To investigate, we consider the effects of a subsidy to enter self-employment. Second, the option value of self-employment is driven by the right tail of expected earnings. In consequence, the progressivity of income taxes may have important implications for the value of entering self-employment. We consider how flat taxes change sorting and earnings patterns. Some parameters in the semi-structural model have clear and policy-invariant interpretations, such as the utility value of after-tax income, while others, such as the entry costs, represent a mix of underlying factors and may be sensitive to the policy environment under which they are estimated. We argue that the two policies considered are unlikely to change the value of the other preference parameters, but also add some caveats to the interpretation of the results.

9.1 Subsidizing Entry into Self-Employment

Our first counterfactual policy offers workers with no prior entrepreneurial experience a subsidy worth \$15,000 if they enter self-employment. We hypothesize that our estimated entry costs capture

both direct financial investments and additional concerns, such as barriers in credit markets or expected (un-modeled) earnings losses if workers choose to return to paid work and must find a few job. In this context, this entry subsidy could represent a direct cash transfer, but could also proxy for a policy that resolves some of these concerns, such as the loosening of borrowing constraints studied in Jensen et al. (2014) or the added job security studied in Gottlieb et al. (2016).

For this exercise, we use the model, with the baseline parameter estimates, to project the probability that each worker in each year of the sample would select self-employment next year with the added incentive of the subsidy, and then project remaining lifetime earnings with workers behaving optimally in response to the subsidy. Figure 8 presents the probability of entry into self-employment under this subsidy and the flat tax policy experiments described in the next section. We plot the effects of these policies across the distribution of paid earnings ability to highlight distributional consequences.³² Throughout, we summarize the choices and earnings for workers with no prior entrepreneurial experience to focus on new entrants.

On average, this \$15,000 subsidy increases the probability that a paid worker will enter entrepreneurship next year by 0.2 percentage points, about a 10% increase on a base entry rate of 2%. The first-time entrepreneur subsidy changes entry the most, both in absolute and percentage terms relative to the baseline results, for those in the bottom 25% of the paid earnings distribution. The \$15,000 subsidy makes up a larger percentage of earnings for those who earn less in paid work and a larger share of the entry utility costs, which we estimate as increasing in paid earnings ability.

Figure 9 plots the change in expected lifetime earnings under the entry subsidy, relative to the baseline projections. These expected earnings reflect the average of 250 simulations, as in the results discussed in section 8.2. The top panel plots changes in pre-tax earnings, which captures the gross social return of these policies. Pre-tax earnings fall slightly in response to the subsidy for workers in the bottom 30% of the paid ability distribution. These workers are lured into self-employment

³²Humphries (2016) considers age-dependent subsidies in a model of self-employment with capital investment estimated on Swedish data. His paper does not feature learning, so there are no allocation issues across sectors. A recent paper by Hincapie (2017) does consider a model in which individuals are assumed to learn about earnings in both paid work and entrepreneurship. He considers a subsidy for young entrepreneurship and finds small long-term effects on entry and earnings. We do not find meaningful differences by age in the percent change in entry probability in response to the subsidy, except for the last few years of working life. Young workers demonstrate the largest responses in levels because, as shown in Figure 4, young workers are the most likely to enter self-employment.

by the subsidy but, because of the correlation in earnings between sectors and the larger variance of earnings in self-employment, earn even less on average as entrepreneurs than they did in paid work.³³ Pre-tax earnings are higher under the subsidy for workers with high paid earnings ability, though again the changes are small, reflecting improved sorting between sectors from a policy that induces more workers to learn about their entrepreneurial abilities. The second panel, which plots changes in after-tax earnings in response to the subsidy, illustrates how progressive taxes further mute these small changes in earnings.

This experiment yields two key insights. First, a policy to subsidize entry into entrepreneurship may be most effective when targeted at high-earning paid workers. These workers are less likely to enter self-employment, but also more likely to earn relatively more as entrepreneurs. Subsidies to low-earning workers may be counterproductive. Second, the effect of this policy on average earnings is likely to be small. The previous section showed positive but modest gains in expected lifetime earnings for workers who experiment with self-employment. Because we estimate strong non-pecuniary concerns about self-employment, we find that even a fairly generous \$15,000 subsidy has a limited effect on entrepreneurship rates. In the baseline lifetime earnings simulations, workers with no entrepreneurial experience have a mean probability of becoming self-employed in some future year of 24.0% (30.5% for workers who are currently 35 or younger). With the addition of the subsidy, this mean probability increases to 25.3% (32.0% for workers under 35).

If implemented as a cash transfer, this subsidy would be a net cost for the government. The average cost of providing the subsidy exceeds both the average increase in projected tax revenues and the average increase in total pre-tax earnings at all points in the paid earnings ability distribution.

9.2 Flat Taxes

We now turn to an examination of how changing the progressivity of taxes changes entry into and earnings in self-employment. We first consider a 15% flat tax that applies only to entrepreneurs while leaving progressive taxes on paid earnings unchanged. This policy mimics, for example, a

³³These projections do not include the value of the subsidy; expected after-tax income, including the subsidy for those who enter self-employment, is marginally higher under the subsidy for all workers.

flat tax on pass-through income. We also consider a flat 15% income tax for both entrepreneurs and paid workers. The latter is in line with proposals to simplify the tax code by eliminating the progressivity of income taxes. Statutory differences in the payroll tax are left fixed. Because we do not model risk aversion explicitly, the estimated utility costs of entering self-employment may partially capture a distaste for the additional uncertainty surrounding entrepreneurial earnings. Progressive income taxes provide partial insurance for this uncertainty by counteracting some of the costs of income losses and the benefits of income gains. If workers are risk averse, our simulations will overstate the effect of flat tax policies on self-employment rates by ignoring the cost of removing this partial insurance.

The value of experimenting with entrepreneurship is driven by the possibility that a worker will discover that he is a very successful entrepreneur. Flattening the tax schedule makes this possibility even more attractive by cutting tax rates for high earners. Moving to a flat tax scheme for all labor earnings increases the probability of entering self-employment for all workers. The largest increases, in both percentage and absolute terms, come from workers with high paid earnings ability because these workers have the highest chance of also being high-earning entrepreneurs. A uniform flat tax is more effective for persuading these high earners to enter self-employment than a direct subsidy for entrepreneurship.

A policy that flattens tax rates for only self-employment earnings still enhances the value of experimentation, but also gives workers the opportunity to choose between two tax schemes. High-earning paid workers, who face high marginal tax rates under a progressive scheme, now have a double incentive to enter self-employment. In contrast, low-earning paid workers will expect less favorable tax treatment as entrepreneurs. The net effect of this policy is a slight decline in the probability of entering entrepreneurship for workers in the bottom 20% of the paid earnings ability distribution.

The changes in pre-tax earnings, plotted in the top panel of Figure 9, are driven entirely by the effect of these policies on workers' labor supply choices. For high-ability workers, large increases in the probability of entering self-employment under the flat tax plans translate into meaningful increases in expected pre-tax lifetime earnings. Workers at the top of the paid earnings ability

distribution expect to earn 1% more before taxes under a uniform flat tax plan and 1.5% more if only self-employment earnings are taxed at a flat rate. Workers with lower paid earnings ability expect smaller gains in pre-tax earnings, but the net change is positive for all but the lowest-earning paid workers.

The bottom panel of Figure 9 combines these changes in total earnings with changes in tax treatment. Under a uniform flat tax, high earners experience large increases in after-tax income relative to the baseline policy, while low-ability workers experience large declines in after-tax income. If the flat tax is only applied to entrepreneurial earnings, high-ability workers experience somewhat smaller gains, both because non-pecuniary concerns still limit entry into self-employment and because not all high earners in the paid sector are also high-earning entrepreneurs. Under this second policy, workers with lower paid earnings ability can avoid after-tax earnings losses by not working for themselves.

These flat-tax policies are effective at increasing experimentation and pre-tax earnings, particularly for high-earning paid workers, but they are expensive to implement for the government. 15% is approximately equal to the average income tax burden under our progressive tax model, so average tax rates are roughly unchanged under the uniform flat tax scheme. However, the higher tax rates on low earnings cannot make up for the lower rates on high earnings.

10 Conclusion

We formulate a model where individuals have heterogeneous earnings abilities in self-employment and paid work. Workers are initially uncertain about potential earnings in entrepreneurship, but can learn about their earnings by entering self-employment and observing how well they fare. Upon discovering entrepreneurial earnings are below paid earnings, they can return to paid work, limiting the downside risk of their experimentation. We document that the patterns of movements in and out of self-employment in the PSID are consistent with this kind of initial uncertainty, experimentation, and strategic returns to paid work. We then estimate a lifecycle model of labor sector choice that incorporates learning about earnings abilities and dynamic re-optimization.

Other papers have recognized that the ex-ante value of experimenting with entrepreneurship differs from static expected earnings. We believe this is the first paper characterize how the value of resolving uncertainty about entrepreneurial earnings varies over the lifecycle and how policies to encourage entrepreneurship affect learning, earnings, and sorting into self-employment.

We estimate that workers face substantial ex ante uncertainty about their own potential earnings in self employment. Those who have never worked in self-employment have a standard deviation of expected pre-tax entrepreneurial earnings of 56%, or \$33,000 per year on average. This uncertainty makes experimenting with entrepreneurship attractive; workers may find they earn much more working for themselves, but are cushioned from the symmetric downside risk by the option to return to paid work. We estimate that, for young workers, the expected lifetime earnings from entering self-employment look substantially more favorable than the expected static earnings from the cross-sectional distribution. The option to return to paid work drives these returns, making the monetary returns from self-employment and paid work similar at the median.

With similar monetary returns between paid work and self-employment for the median worker, two different sources may potentially explain why less than 30% of the sample is ever observed in self-employment. First, we estimate costs to enter entrepreneurship that are larger than the direct monetary investments in new businesses observed in the data. These costs may pick up a range of additional factors: the arrival of a business idea, unobserved precautionary cash needs, and the uncertainty and effort costs that come with setting up a new venture. Second, earlier studies have hypothesized that workers experience non-pecuniary benefits from working for themselves.³⁴ We identify a more nuanced role for tastes for entrepreneurship. We allow for heterogeneous tastes for entrepreneurship and estimate wide dispersion. In fact, the average worker must dislike working for himself in order to rationalize low entry rates to self-employment. However, the variance around this mean is large enough to span positive and negative values. In particular, the average worker who we observe entering self-employment has an ex post positive estimated taste for entrepreneurship. These heterogeneous preferences play an important role in matching the behavior of both those who enter entrepreneurship and those who do not.

³⁴See Åstebro and Thompson (2011) for a review.

The role of preference heterogeneity is relevant for evaluating programs designed to encourage entrepreneurship. If workers are initially uncertain about their entrepreneurial earnings ability, then programs that facilitate resolving this uncertainty should improve the allocation of workers across sectors. We examine two such policies: subsidies for entry into entrepreneurship and flat taxes. Both are successful at increasing the share of workers who experiment with self-employment and modestly increasing average lifetime earnings, with the caveat that preference differences dilute sorting on relative earnings potential. However, neither policy has a net positive effect on government revenues. Other programs that lower the cost of learning may have more success. For example, Lerner and Malmendier (2011) find that Harvard Business School graduates who interacted with more former entrepreneurs during school were less likely to become entrepreneurs themselves, but more likely to succeed if they did so. Gottlieb et al. (2016) find that a Canadian reform that lowered the cost of returning to paid work after a temporary absence encouraged more workers to enter self-employment.

References

- Altonji, Joseph G and Nicolas Williams (2005) “Do wages rise with job seniority? A reassessment,” *Industrial & Labor Relations Review*, Vol. 58, pp. 370–397.
- Åstebro, Thomas, Holger Herz, Ramana Nanda, and Roberto A Weber (2014) “Seeking the roots of entrepreneurship: insights from behavioral economics,” *The Journal of Economic Perspectives*, Vol. 28, pp. 49–69.
- Åstebro, Thomas and Peter Thompson (2011) “Entrepreneurs, Jacks of all trades or Hobos?” *Research Policy*, Vol. 40, pp. 637–649.
- Blau, David M (1987) “A time-series analysis of self-employment in the United States,” *Journal of political economy*, Vol. 95, pp. 445–467.
- Bruce, Donald (2000) “Effects of the United States tax system on transitions into self-employment,” *Labour economics*, Vol. 7, pp. 545–574.
- Bruce, Donald and Herbert J Schuetze (2004) “The labor market consequences of experience in self-employment,” *Labour Economics*, Vol. 11, pp. 575–598.
- Caliendo, M and A Uhlendorff (2008) “Self-employment dynamics, state dependence and crossmobility patterns. IZA Discussion Papers 3900,” *Institute for the Study of Labor (IZA), Bonn*.
- Caliendo, Marco, Frank Fossen, and Alexander Kritikos (2010) “The impact of risk attitudes on entrepreneurial survival,” *Journal of Economic Behavior & Organization*, Vol. 76, pp. 45–63.
- Carroll, Christopher D and Andrew A Samwick (1997) “The nature of precautionary wealth,” *Journal of monetary Economics*, Vol. 40, pp. 41–71.
- Daly, Moira (2015) “The long term returns of attempting self-employment with regular employment as a fall back option,” *Labour Economics*, Vol. 35, pp. 26–52.
- Evans, David S and Boyan Jovanovic (1989) “An estimated model of entrepreneurial choice under liquidity constraints,” *The Journal of Political Economy*, pp. 808–827.

- Evans, David S and Linda S Leighton (1989) "Some empirical aspects of entrepreneurship," *The American Economic Review*, pp. 519–535.
- Gibbons, Robert, Lawrence F. Katz, Thomas Lemieux, and Daniel Parent (2005) "Comparative Advantage, Learning, and Sectoral Wage Determination," *Journal of Labor Economics*, Vol. 23, pp. 681–724.
- Gottlieb, Joshua D., Richard R. Townsend, and Ting Xu (2016) "Experimenting with Entrepreneurship: The Effect of Job-Protected Leave," working paper.
- Hall, Robert E and Susan E Woodward (2010) "The Burden of the Nondiversifiable Risk of Entrepreneurship," *American Economic Review*, Vol. 100, pp. 1163–1194.
- Hamilton, Barton H (2000) "Does entrepreneurship pay? An empirical analysis of the returns to self-employment," *Journal of Political Economy*, Vol. 108, pp. 604–631.
- Hincapie, Andres (2017) "Where are the Young Entrepreneurs? A Study of Entrepreneurship over the Life Cycle."
- Humphries, John Eric (2016) "The Causes and Consequences of Self-Employment over the Life Cycle."
- Hurst, Erik, Geng Li, and Benjamin Pugsley (2014) "Are household surveys like tax forms? Evidence from income underreporting of the self-employed," *Review of economics and statistics*, Vol. 96, pp. 19–33.
- Hurst, Erik and Benjamin Wild Pugsley (2011) "What do small businesses do?" Technical report, National Bureau of Economic Research.
- Jensen, Thais Lærkholm, Søren Leth-Petersen, and Ramana Nanda (2014) "Housing collateral, credit constraints and entrepreneurship-evidence from a mortgage reform," Technical report, National Bureau of Economic Research.
- Jones, Benjamin F (2010) "Age and great invention," *The Review of Economics and Statistics*, Vol. 92, pp. 1–14.

- Jovanovic, Boyan (1979) “Job matching and the theory of turnover,” *Journal of Political Economy*, pp. 972–990.
- Kahn, Shulamit, Giulia La Mattina, and Megan MacGarvie (2016) “Misfits, Stars and Immigrant Entrepreneurship,” working paper.
- Kerr, William R, Ramana Nanda, and Matthew Rhodes-Kropf (2014) “Entrepreneurship as experimentation,” *The Journal of Economic Perspectives*, Vol. 28, pp. 25–48.
- Lazear, Edward P, Kathryn L Shaw, and Christopher T Stanton (2015) “The value of bosses,” *Journal of Labor Economics*, Vol. 33, pp. 823–861.
- Lerner, Josh and Ulrike Malmendier (2011) “With a Little Help from My (Random) Friends: Success and Failure in Post-Business School Entrepreneurship,” nber working paper no. 16918.
- Levine, Ross and Yona Rubinstein (2013) “Smart and Illicit: Who Becomes an Entrepreneur and Does It Pay?” working paper.
- Liang, James, Hui Wang, and Edward P Lazear (2014) “Demographics and Entrepreneurship,” Technical report, National Bureau of Economic Research.
- Light, Audrey and Robert Munk (2015) “Business ownership vs. self-employment,” working paper.
- Lim, Katherine (2015) “Self-Employment, Workplace Flexibility, and Maternal Labor Supply: A Life-Cycle Model,” working paper.
- Manoso, Gustavo (2016) “Experimentation and the Returns to Entrepreneurship,” *Review of Financial Studies*, p. hh019.
- Miller, Robert A (1984) “Job matching and occupational choice,” *The Journal of Political Economy*, pp. 1086–1120.
- Moscarini, Giuseppe (2005) “Job matching and the wage distribution,” *Econometrica*, Vol. 73, pp. 481–516.

- Moskowitz, Tobias J and Annette Vissing-Jorgensen (2002) “The returns of entrepreneurial investment: A private equity premium puzzle?” *The American Economic Review*, Vol. 92, p. 745.
- Papageorgiou, Theodore (2014) “Learning your comparative advantages,” *The Review of Economic Studies*, Vol. 81, pp. 1263–1295.
- Rust, John (1987) “Optimal replacement of GMC bus engines: An empirical model of Harold Zurcher,” *Econometrica: Journal of the Econometric Society*, pp. 999–1033.
- Sarada (2015) “The Unobserved Returns from Entrepreneurship,” *University of Wisconsin-Madison manuscript*.
- Train, Kenneth E (2009) *Discrete choice methods with simulation*: Cambridge university press.
- Vereshchagina, Galina and Hugo A Hopenhayn (2009) “Risk taking by entrepreneurs,” *The American Economic Review*, pp. 1808–1830.

Table 1: Summary Statistics for Men in the PSID

	Never Entrepreneurs	Sometime Entrepreneurs
Number of individuals	5,287	1,550
Years spent in paid work	17.3	11.0
Years spent in entrepreneurship	0.0	8.8
Ever own a business	0.21	0.82
White	0.79	0.88
Black	0.12	0.05
Hispanic	0.07	0.05
Other race	0.02	0.01
Less than HS diploma	0.08	0.08
High school diploma	0.33	0.28
Some college	0.31	0.32
College graduate	0.19	0.20
Graduate degree	0.08	0.11

Source: PSID 1976-2011. Averages use sampling weights. Years spent in each sector is as of last/most recent survey wave.

Table 2: Predicting Earnings in Self-Employment

	(1)	(2)	(3)	(4)	(5)
Intercept	-1.36*	-1.01*	-0.63	-0.36	0.70
	(0.35)	(0.39)	(0.40)	(0.40)	(0.50)
Log paid sector fixed effect	1.21*	1.16*	1.11*	1.04*	0.98*
	(0.05)	(0.06)	(0.06)	(0.06)	(0.06)
Pushed In			-0.23*	-0.22*	-0.21*
			(0.05)	(0.05)	(0.05)
Formed unincorp. bus. at entry				0.14*	0.12
				(0.07)	(0.07)
Formed incorp. bus. at entry				0.29*	0.27*
				(0.08)	(0.07)
Zero initial investment				0.20*	0.17
				(0.09)	(0.09)
Start-up investment \$5,001-25,000				0.06	0.08
				(0.08)	(0.08)
Start-up investment \$25,001-65,000				0.01	0.02
				(0.11)	(0.11)
Start-up investment > \$65,000				0.25	0.26
				(0.17)	(0.16)
Past experience in industry					0.04
					(0.05)
Race and Education	No	Yes	Yes	Yes	Yes
Enter sample in self-emp.	No	No	Yes	Yes	Yes
Missing initial investment	No	No	No	Yes	Yes
Industry	No	No	No	No	Yes
Observations	1,438	1,438	1,438	1,438	1,438
R-square	0.44	0.45	0.46	0.48	0.51
Adj. R-Square	0.44	0.45	0.46	0.48	0.49

* Indicate significance at 5%. The dependent variable is the log of the individual earnings fixed effect in self-employment. The estimation of these individual effects from observed earnings, and for the fixed effect in the paid sector, is described in Section 7.1. Workers are considered “pushed in” to self-employment if they experience at least a 30% drop in weekly paid earnings or work at least 14 weeks less than usual in the year prior to entering entrepreneurship. We include an indicator for entering the sample in self-employment and set the pushed-in indicator to zero for these workers. We observe money invested when workers began a spell of self-employment, including reported zeros, for only 46% of our self-employed sample. The question was asked only every five years for most of the sample, and only asked to workers who reported owning a business at the time of the survey.

Table 3: Exit Rates from Entrepreneurship

	(1)	(2)	(3)	(4)
Ent. earn q2	0.72*	0.78*		
	(0.00)	(0.03)		
Ent. earn q3	0.55*	0.78		
	(0.00)	(0.11)		
Ent. earn q4	0.35*	0.90		
	(0.00)	(0.62)		
Earn diff. q2		0.99	0.97	0.96
		(0.94)	(0.79)	(0.82)
Earn diff. q3		0.65*	0.58*	0.61*
		(0.01)	(0.00)	(0.02)
Earn diff. q4		0.28*	0.26*	0.23*
		(0.00)	(0.00)	(0.00)
Pushed in			1.10	1.11
			(0.56)	(0.54)
Pushed in* Earn diff. q2			1.08	1.05
			(0.75)	(0.83)
Pushed in* Earn diff. q3			1.17	1.17
			(0.60)	(0.60)
Pushed in* Earn diff. q4			1.83	1.89
			(0.16)	(0.14)
No prior paid work			0.55*	0.54*
			(0.00)	(0.00)
In high-capital ind.				1.18
				(0.24)
High cap.* Earn diff. q2				1.04
				(0.86)
High cap.* Earn diff. q3				0.92
				(0.74)
High cap.* Earn diff. q4				1.22
				(0.59)
Missing industry				1.30
				(0.23)
Observations	6,191	6,191	6,191	6,191
Log likelihood	-3,332	-3,318	-3,294	-3,291

Table reports hazard ratios with **p-values from z-tests in parentheses**. Ent. earn is projected entrepreneurial earnings for the coming year, as described in Section 7.1. Earn diff. is the difference between projected entrepreneurial earnings for the coming year and projected paid earnings. “Pushed in” is as defined in the notes to Table 2. No prior paid work indicates that workers have been self-employed since they entered the PSID sample. High-capital industries are those with an above-median share of workers who invest at least \$25,000 when entering entrepreneurship. This measure is taken from the 2007 Survey of Business Owners (SBO) Public Use Microdata.

Table 4: Determinants of Log Earnings

Experience profiles in paid sector			
1 year paid experience			0.051 (0.005)
10 years paid experience			0.421 (0.385)
Experience profiles in entrepreneurship			
1 year paid experience			0.028 (0.036)
10 years paid experience			0.202 (0.440)
1 year entrepreneurial experience			0.031 (0.012)
10 years entrepreneurial experience			0.229 (0.147)
Distribution of abilities			
Mean log ability in paid sector	μ_α		6.450 (0.017)
Variance of log ability in paid sector	σ_α^2		0.193 (0.020)
Mean log ability in entrepreneurship	μ_η		6.432 (0.089)
Variance of log ability in entrepreneurship	σ_η^2		0.631 (0.055)
Correlation of abilities across sectors	ρ		0.702 (0.063)
Paid sector earnings shocks			
Variance of AR(1) innovation	σ_ζ^2		0.024 (0.009)
Annual persistence of AR(1)	ϕ		0.831 (0.079)
Variance of transitory shock	σ_m^2		0.022 (0.009)
Entrepreneurship earnings shock			
Variance of transitory shock	σ_e^2		0.096 (0.008)

PSID 1976-2011. Estimated on weekly earnings as described in the text. Bootstrapped standard errors from 200 draws in parentheses. Paid earnings depend on a cubic in paid experience. Entrepreneurial earnings depend on a quadratic in paid experience and a cubic in entrepreneurial experience. The order of these polynomials were determined by a series of specification F-tests. We discuss omitting entrepreneurial experience from the paid earnings equation in the text. In the simulations annual earnings are projected from these estimates assuming 50 weeks of work per year.

Table 5: Utility Parameters

	Full model	No heterogeneous tastes
Parameter estimates		
β_1 : Utils per \$10,000	0.143 (0.005)	0.150 (0.017)
β_3 : Startup cost, as share of paid fixed effect	-1.433 (0.019)	-1.873 (0.065)
β_2 : Additional cost of first entry	-0.094 (0.030)	-0.081 (0.081)
μ_{β_0} : Mean non-pecuniary benefit	-1.090 (0.041)	0.009 (0.002)
σ_{β_0} : Std. dev. of non-pecuniary benefit	1.043 (0.028)	
Dollar equivalent interpretations		
Mean cost to enter self-emp.	-325,289	-405,622
Mean add'l cost of first entry	-21,243	-17,494
Std. dev. of transitory preference shock	89,877	85,784
Mean non-pecuniary benefit	-76,377	624
Std. dev. of non-pecuniary benefit	73,083	0
% of workers with $\beta_{0i} > 0$	15	0
Posterior Distributions of Non-pecuniary Benefit		
Mean $\mu_{\beta_{0i}}$, never entrepreneurs	-95,085	
Mean $\sigma_{\beta_{0i}}^2$, never entrepreneurs	62,273	
Mean $\mu_{\beta_{0i}}$, sometime entrepreneurs	1,600	
Mean $\sigma_{\beta_{0i}}^2$, sometime entrepreneurs	31,002	

The utility parameter per \$10,000 has no independent interpretation except to scale the variance of the extreme-value taste shock. Dollar equivalents are in 2010 USD. Entry costs are scaled to individual fixed earnings ability in paid work, $50 * \exp(\alpha_i)$, so the table reports mean entry costs across individuals. See text for details of estimation.

Table 6: Estimated Probability of Choosing Self-Employment

Sector Last Year:	Observed Choice this Year	
	Paid Sector	Self-Employment
Paid sector		
Income maximizing	54.5%	58.7%
Full model, homogeneous tastes	2.4%	3.7%
Full model, heterogeneous tastes	2.1%	12.3%
Share of prior paid workers	98.1%	1.9%
Self-Employment		
Income maximizing	48.6%	72.2%
Full model, homogeneous tastes	76.1%	83.8%
Full model, heterogeneous tastes	66.0%	88.5%
Share of prior entrepreneurs	7.8%	92.2%

Each column presents the model-predicted probability of choosing entrepreneurship, separately for workers who were in each sector last year and are observed to choose each sector this year. If the model predicted choices perfectly, we would assign workers observed choosing paid work this year (the first column) a 0% chance of choosing entrepreneurship and workers observed choosing entrepreneurship (the second column) a 100% chance of doing so. The first row in each cell gives the prediction of a model where workers are income maximizers with no entry costs or non-pecuniary benefits of entrepreneurship. The next row presents predictions from a model where agents are utility maximizers with our full preference specification, but without any heterogeneity of preferences. The final row presents predictions of the full model, using individual posterior distributions of preferences.

Table 7: The Value of Entering Self-Employment at Age 30

	Value of Paid	Value of Entrep.	Difference	Pct. Difference
Annual observed earnings in chosen sector				
Mean	41,119	39,093	-2,026	-4.9%
Median	39,743	30,823	-8,920	-22.4%
N	2,742	254		
Projected annual earnings for all workers				
Mean	41,636	38,585	-3,051	-7.3%
Median	40,176	35,035	-5,141	-12.8%
N	2,996	2,996		
Projected lifetime earnings, static model				
Mean	47,551	44,412	-3,139	-6.6%
Median	45,282	39,817	-5,464	-12.1%
N	2,996	2,996		
Projected lifetime earnings, dynamic model				
Mean	47,146	47,537	391	0.8%
Median	44,447	43,567	-880	-2.0%
N	2,996	2,996		

This table presents after-tax earnings estimates for 30 year olds. For discussion of these estimates at other ages and before taxes, see section 8. The first panel of this table summarizes observed annual earnings for 30 year old paid workers and entrepreneurs. In this panel the workers contributing to the paid averages are different from the workers contributing to the entrepreneurial average. In the remaining panels, earnings are projected for all 30 year old workers conditional on choosing each sector, so the same sample of workers contributes to the averages for each sector. For ease of comparison, the projected lifetime earnings are converted to constant-annual-income equivalents, $\bar{C}$ such that $\sum_{s=t}^T \left(\frac{1}{1+r}\right)^s Y_s = \bar{C} \sum_{s=t}^T \left(\frac{1}{1+r}\right)^s$. The third panel projects discounted lifetime earnings assuming that workers choose each sector this year and remain there for all future years. The fourth panel projects discounted lifetime earnings assuming workers choose each sector this year and then behave optimally according to the full model in each subsequent year.

Table 8: Sensitivity of Key Averages for 30 Year Olds

	Pct. Earnings Gain in Self-Emp., Static	Pct. Earnings Gain in Self-Emp., Dynamic
Baseline	-6.6%	0.8%
Mean self-employed earnings ability, μ_η		
5% larger	-2.9%	0.9%
5% smaller	-10.1%	0.6%
Var. self-employed earnings ability, σ_η^2		
5% larger	-5.5%	1.5%
5% smaller	-7.7%	-0.4%
Correlation of earnings ability, ρ		
5% larger	-6.4%	1.0%
5% smaller	-6.8%	0.6%


This table summarizes the sensitivity of key outcomes to several parameters estimated in the first stage. The table presents the mean for all 30 year olds of the percent difference between expected lifetime earnings conditional on entering self-employment next year and expected lifetime earnings conditional on choosing paid work. The first column corresponds to static-model values reported in the third panel of Table 7. The second column corresponds to dynamic-model values reported in the fourth panel of Table 7. The parameters are adjusted by +/- 0.05 for μ_η , +/- 0.03153 for σ_η^2 , and +/- 0.035099 for ρ .

Figure 1: Composition of Entrepreneurs Over Time




Source: PSID 1976-2011. Note: Because of bi-annual interviews after 1997 the sample is not consistent over years since entering entrepreneurship. For example, we may see some workers 2 and 4 years post-entry and others 3 and 5 years post-entry.

Figure 2: Distribution of Earnings in Paid Work and Entrepreneurship




Source: PSID 1976-2011. Distribution of real weekly earnings in 2010 dollars. Truncated at \$4,000 per week, which excludes the top 2% of earnings. Weighted as described in the text.

Figure 3: Earnings Profiles by Persistence in Entrepreneurship




Source: PSID 1976-2011. Profiles are average real annual earnings for entrepreneurs in 2010 dollars, less estimated taxes. Estimated taxes are described in Appendix B.2. Weighted as described in the text. The gap between each of the two lower profiles and the top profile are statistically significant with 99% confidence.

Figure 4: First Entry into Entrepreneurship by Age



Source: PSID 1976-2011 and predicted likelihood of choosing entrepreneurship from the model. Both series describe the average probability of choosing to work in entrepreneurship for individuals who worked in the past sector in the previous year and have no past entrepreneurial experience. Data are weighted as described in the text.

Figure 5: Exits from Entrepreneurship by Time in Entrepreneurship



Source: PSID 1976-2011 and predicted likelihood of choosing paid work from the model. Both series describe the probability of selecting paid work for individuals who worked as entrepreneurs in the previous year. Data are weighted as described in the text.

Figure 6: Distribution of the Annualized Option Value for 30 Year Olds




Figure plots the annualized option value of entrepreneurship, calculated by taking the expected lifetime earnings conditional on choosing entrepreneurship this year and behaving optimally in future years less the annualized expected lifetime earnings conditional on choosing entrepreneurship in all future years. The average values of these projections are presented in the 4th and 3rd panels of Table 7, respectively. The solid line plots the option value for workers who have no entrepreneurial experience by age 30. The dashed line plots the option value for workers with some entrepreneurial experience, both current and former entrepreneurs.

Figure 7: Expected Post-tax Lifetime Earnings Gains from Entering Self-Employment by Age






Figure plots the expected lifetime earnings gains conditional on choosing entrepreneurship this year and behaving optimally in future years for those without prior entrepreneurial experience. Note that this figure is the present value of future earnings, not the annualized present value that is contained in other figures and calculations. These calculations are also not comparable to those in Table 7 because only individuals without entrepreneurial experience are included. To adjust for composition differences over the lifecycle, the distribution of α , the paid earnings ability that forecasts entrepreneurial earnings, is held constant.

Figure 8: Entry into Self-Employment Under Counterfactual Policies

A. Entry Probability by Paid Earnings Ability




B. Change in Entry Probability Relative to Baseline



Figures plot the average predicted probability of selecting self-employment each period, conditional on having no prior entrepreneurial experience, in our baseline model and in the models implementing policy counterfactuals. Probabilities are calculated using each individual's posterior distribution of taste for entrepreneurship. The second panel plots the average percent difference in probabilities, not the percent difference of the averages.

Figure 9: Effect of Counterfactual Policies on Expected Lifetime Earnings



Figures plot changes in simulated lifetime earnings, relative to the simulations under the baseline policy structure. Both sets of simulations assume that workers' utility depends on after-tax earnings. The top panel plots changes in projected total, pre-tax earnings, while the second plots changes in projected post-tax earnings.

A Data Appendix

A.1 Sample Construction

The original PSID sample includes a representative group of American households in 1968 and an oversample of low-income households. The PSID has continued to interview members of these households, their offspring, and individuals who marry into these families. Over time, some original families have been dropped and the sample has been augmented with several samples of recent immigrants to better reflect the current mix of US households. The baseline sample includes all households interviewed in each year; PSID-constructed weights are used to adjust for probability of inclusion in the survey. These changes in the sample, along with workers aging into and out of the workforce and occasional non-response, create an unbalanced panel.

Our PSID estimation sample is constrained by the need to keep track of accumulated work experience in each sector. Individuals are included in the sample only when we can follow their work experience starting at age 25 or earlier. Both survey responses and the panel dimension of the data are used to construct measures of accumulated experience in each sector. In other words, we keep all workers who enter the sample by age 25, plus any worker who enters the sample at an older age but reports being in their current job since before they were 25. We follow respondents starting with the 1976 survey, the year the PSID began asking about job tenure, but, where possible, we use reported paid or self-employed work from the 1968-1975 waves to help construct experience measures. In the first year a worker is observed, we initialize their sector experience using their report of how long they have been at their current job. For example, if a worker enters the sample working in the paid sector and has been working at that job for 8 years, we say he has 8 years of paid experience and no entrepreneurial experience. Experience in each sector is then updated with observed work going forward.

Because we estimate individual-specific earnings effects, we also restrict the sample to workers with at least five years of observed earnings. To abstract from early retirement behavior and schooling decisions, we include only men between the ages of 22 and 55.

A.2 Variable Definitions

We define an entrepreneur as someone who is self-employed in their main job. In most cases, this simply means the worker is self-employed. The PSID asks workers who hold more than one job simultaneously to identify their main job, which tends to be the job that accounts for the majority of hours worked and earnings. We classify individuals who work for someone else on their main job and for themselves in a side job as paid workers, and vice versa.

In the later part of the PSID sample, business owners are asked if they work in their business. In these later years we observe that some business owners who report working in their business and having only one job nonetheless report working for someone else. This combination of responses is somewhat puzzling, but on the whole we would like to classify these workers as entrepreneurs throughout the sample. To do so consistently, we also identify workers as entrepreneurs when their first report of owning a business coincides with a job change. This rule captures most of the respondents who report working in their businesses in the later half of the sample.

In the PSID, paid workers and incorporated business owners report total labor earnings, but unincorporated business owners are only asked about total net profits from their business. For these workers, we use net profits as our measure of labor income. One potential drawback is that using net profits will overstate labor earnings in cases where workers have also invested substantial financial capital in their businesses. Information about business assets are too incomplete in the PSID to allow us to adjust net profits for a reasonable return to capital. However, the data that are available suggest that these adjustments would be small. The median worker who opens a business upon entering entrepreneurship invests \$5,000, in 2010 dollars, over the first 1-5 years.³⁵ This number likely overstates the median entry investments for all self-employed workers, since workers who become self-employed without reporting ownership of a business are not asked about capital investments. Even at the 90th percentile, unincorporated business owners invest only \$59,000. At a

³⁵The exact question wording is "Since [date of last survey], did you (or anyone in your family) put money into a business or farm?" and, if yes, "How much did you put in altogether?" Until 1999, this question was asked only every five years, about investments for the past five years. We define the first answer to these questions after a worker first reports earning a business as the "initial investment," but we are only able to match 46% of the self-employed workers to a response to this question (it was only asked for workers who still reported owning a business at the time of the survey) and these numbers should be interpreted with caution.

5% rate of return, a \$59,000 investment would contribute only \$2,900 per year to business income, while the median annual income for business owners is \$60,000.

Because work hours are difficult to define for self-employed workers, earnings are measured at the weekly level using annual earnings divided by reported weeks worked. Both labor income and business profits are reported over the past calendar year, while job tenure is reported as of the survey date. Earnings for years when individuals spent part of their time in the paid sector and part in entrepreneurship are excluded from earnings regressions. To estimate the effect of sector experience on earnings we construct experience as of the beginning of the last calendar year, using reports from both the current and previous waves of the survey.

B Estimation Appendix

B.1 Determinants of Earnings and State Variable Transitions

The set of state variables, S_{it} , includes age, accumulated experience in each sector, x_{itR} and x_{itW} , paid sector ability, α_i , expected entrepreneurial skill, $\hat{\eta}_{ix}$, the lag of the persistent shock in the paid sector, P_{it} , and the lagged sector choice, d_{it-1} . Within the model, age and experience evolve deterministically based on sectoral choices, α_i is fixed, and $\hat{\eta}_{ix}$ and P_{it} , depend only on sector choice and exogenous shocks. Because none of these processes depend on β_{0i} , when we substitute equation (13) into (14), $f(S_{t+1}|d_t, S_t; \theta)$ moves outside of the integral over β_i and is additively separable in the log likelihood. We can therefore estimate $\hat{\theta}$ separately in a first stage, with appropriate controls for non-random selection into each sector.³⁶

To estimate the effects of sector experience on earnings we must account for workers' strategic transitions between sectors. Depending on the joint distribution of abilities in each sector, transitions from paid work to self-employment may be positively or negatively correlated with paid ability, α . We estimate a version of Equation 3, describing paid earnings, that includes individual fixed effects to purge permanent components from the estimates.

³⁶This approach rules out some potentially interesting extensions. For example, if x_{itR} and x_{itW} represent accumulated skills rather than years of experience, individuals could have heterogeneous learning abilities. However, with the relatively small sample in the PSID, it would be difficult to allow additional flexibility.

While this strategy yields reasonable estimates for the effect of paid experience on paid sector earnings, individual fixed effects are unlikely to remove entirely selection bias from the estimated effects of entrepreneurial experience on paid sector earnings. For example, if workers sometimes know the realization of shocks before choosing their sector, say because they are offered a particularly attractive new job, then workers who return to the paid sector will have disproportionately high post-self-employment paid earnings. Evans and Jovanovic (1989) and Hamilton (2000) both find positive effects of entrepreneurial experience on earnings, but neither study controls for selection. Like Bruce and Schuetze (2004), we find small and imprecise returns to entrepreneurial experience after various attempts to control for non-random selection between sectors. Because of the difficulty of fully accounting for this selection, in the current specification we impose that entrepreneurial experience has no effect on paid sector earnings. This zero relationship suggests that experimenting with entrepreneurship involves an opportunity cost of lost paid sector experience, but no additional costs or benefits. Even without separate indicators for entrepreneurial experience, our specification of the paid sector earnings process does a good job of fitting earnings for workers who are newly returned to the paid sector from spells of entrepreneurship.

After netting out the experience profile, residual paid earnings, w_t , consist of fixed paid ability, the persistent wage shock $\log(P_t) = \phi \log(P_{t-1}) + \zeta_t$, and the transitory shock $\log(M_t)$. Following Carroll and Samwick (1997), we distinguish the persistent and transitory shocks by comparing the variance and covariance of residuals over long and short intervals. We use two-, four-, and six-year intervals, which we can construct both before and after the PSID's move to bi-annual interviews in 1997. The variance-covariance moments are

$$\begin{aligned}
\text{var}(w_t) &= \sigma_\alpha^2 + \frac{\sigma_\zeta^2}{(1 - \phi^2)} + \sigma_m^2 \\
\text{cov}(w_t, w_{t-2}) &= \sigma_\alpha^2 + \frac{\phi^2 \sigma_\zeta^2}{(1 - \phi^2)} \\
&\vdots \\
\text{cov}(w_t, w_{t-6}) &= \sigma_\alpha^2 + \frac{\phi^6 \sigma_\zeta^2}{(1 - \phi^2)}.
\end{aligned} \tag{15}$$

These moments provide a consistent estimate of the population variance of paid ability, σ_α^2 . In the second stage, the conditional value functions depend on each individual's paid ability, α_i . Average residual earnings in the paid sector are an unbiased estimator of individual paid ability, but they are inconsistent in short panels, assuming fixed-T asymptotics. With only a few years of paid earnings, we cannot distinguish an individual with low paid ability from an individual with a low persistent shock. To address this difficulty, we first adjust for time-series dependence by re-weighting residuals within individuals to place more weight on years that are less correlated with other observations. We then follow Lazear et al. (2015) and shrink individual predictions towards the consistently estimated cell-specific means using an empirical Bayes approach. The resulting estimator is

$$\hat{\alpha}_i = \bar{w}_z + \sigma_\alpha^2 (\Sigma + \sigma_\alpha^2)^{-1} (w_{it} - \bar{w}_z), \quad (16)$$

where Σ is the variance-covariance matrix of the residuals, using the consistent estimates of σ_α^2 and σ_ζ^2 , and $\bar{w}_z$ is the within-cell average residual from the paid wage equation.³⁷ Finally, we use the average $\hat{\alpha}_i$ across individual workers, not individual observations, as the measure of mean paid sector ability, μ_α . This estimate of μ_α is unbiased if there is no selection into being observed at least once in the paid sector and is consistent with large N asymptotics.

Experience-earnings profiles in entrepreneurship are subject to the same selection concerns as profiles in the paid sector. In entrepreneurship, the bias is a direct consequence of the learning model: workers refine their beliefs about entrepreneurial ability through observing earnings in entrepreneurship. As beliefs become more precise, workers who are now confident they could earn more from paid work are more likely to return to that sector. This selection out of entrepreneurship will overstate the effect of entrepreneurial experience on entrepreneurial earnings. Instead, we account for selection out of entrepreneurship by instrumenting entrepreneurial experience with individually de-measured experience, as in Altonji and Williams (2005).

A second source of potential bias involves selection into entrepreneurship. Workers who wait to enter entrepreneurship, and therefore enter with more paid experience, are more likely to have

³⁷We take means within bins by race/ethnicity and years of completed schooling.

been pushed into self-employment by negative shocks in the paid sector. If this negative selection is correlated with entrepreneurial performance, then estimates of the effect of paid experience on entrepreneurial earnings will also be biased (probably downwards). Because most workers have only one spell in entrepreneurship, during which paid experience remains constant, we cannot estimate the effect of paid experience on entrepreneurial earnings within a fixed effect model. To control for this selection, we include the worker's persistent paid-sector earnings shock in the last period before entering entrepreneurship in the entrepreneurial earnings equation.³⁸ We also include a set of indicator variables for race and completed schooling. These covariates improve the precision of our experience profile estimates, but they all capture elements of the entrepreneurial fixed effect, η_i . To estimate the stochastic components of entrepreneurial earnings we construct residual earnings, r_t by netting out only the effects of the experience profile, leaving any heterogeneity captured by these other covariates as part of the residual. The stochastic components governing entrepreneurial earnings are identified by

$$\begin{aligned}\text{cov}(r_{t+1}, r_t) &= \sigma_\eta^2 \\ \text{cov}(r_{t+1} - r_t, r_t - r_{t-1}) &= -\sigma_\xi^2.\end{aligned}\tag{17}$$

In the final step of this stage, we estimate mean entrepreneurial ability, μ_η , and the correlation between abilities in each sector, ρ . These estimates describe the transferability of skills across sectors and allow us to recover $\hat{\eta}_{i0}$ and $\sigma_{\hat{\eta}_{ix}}^2$. For workers who we observed in entrepreneurship, we construct $\hat{\eta}_i$ in the same way as $\hat{\alpha}_i$. The set of entrepreneurs may not reflect the full distribution of entrepreneurial abilities. To account for this selection, we estimate $\widehat{\mu}_\eta$ as $\widehat{\mu}_\alpha$ plus the average difference between $\hat{\alpha}_i$ and $\hat{\eta}_i$ for individuals who we observe in both sectors. We use this same sample of workers who appear in both sectors to estimate ρ , the correlation between α and η . We construct $\hat{\eta}_{ix}$, the worker's belief about his η with x years of entrepreneurial experience, using equations (6) and (7). Finally, we construct $\hat{\alpha}_i$ for workers observed only in the entrepreneurial sector from $\hat{\eta}_i$ by inverting equation (6).

³⁸This shock is estimated in the paid earnings equation and does not depend on the experience-earnings profile in entrepreneurship. Omitting this control variable leads us to estimate negative effects of paid experience on entrepreneurial earnings.

For these estimates we assume that earnings abilities across the two sectors have a bivariate normal distribution. This assumption is formally rejected by a Doornik-Hansen test on the individuals who we observe in both sectors, but not for violations that would be particularly problematic for our analysis. One concern is that individuals may face substantially more or less conditional uncertainty about their earnings in entrepreneurship than is implied by bivariate normality. Appendix Figure A2 plots the observed and theoretical deviations of workers' realized entrepreneurial earnings ability from the η_i predicted from their paid earnings. The observed deviations are moderately more dispersed than bivariate normality would imply, consistent with measurement error in our estimates of the fixed effects, but they are centered on zero and display no skewness. A second concern is that violations of bivariate normality would lead these prediction errors to be systematically positive or negative along some parts of the paid earnings ability distribution. Appendix Figure A3 does not indicate any such patterns.

B.2 Post-Tax Earnings

Our estimates require total tax liabilities for various levels of head-of-household earnings. We account for two different sources of tax liabilities: payroll taxes and income taxes. For paid workers, half of the payroll tax is paid by the employer. This means that paid-sector earnings reported in survey data are likely to be effectively net of the employer portion of the payroll tax. Self-employed individuals are responsible for both the individual and employer portions of the payroll tax, meaning that reported earnings for the self-employed are unlikely to be net of the employer portion. This different treatment requires an adjustment to make take-home pay comparable.

The payroll tax includes a social security component and a medicare component. Each component has different rates and different maximum earnings amounts after which the marginal rate falls to zero; these maximum thresholds and rates change yearly. We take the timeseries average of the maximum earnings amount over the sample period after converting to 2010 dollars for the social security component of the tax. Between 1994 and 2011, the medicare portion was uncapped, so the medicare tax does not have a maximum. We then take the timeseries average payroll tax rates for the individual and individual plus employer portion for both the social security and

medicare portions of the tax. The individual portion is applied to earnings for paid workers, while the individual plus employer portion is applied to earnings for the self-employed.

We use the TAXSIM program from the NBER to compute income tax liabilities. This requires not just actual income taxes but expected income taxes conditional on decisions made. To compute this expectation, for every household in the sample observed in a given year, we compute that year’s total household tax liability for the household’s actual nominal earnings and for a variety of different earnings scenarios for the head, holding fixed spousal wages. These scenarios vary the head’s labor earnings at the following nominal equivalents of real 2010 dollars: $\{\$10,000, \$15,000, \dots, \$150,000\}, \$250,000, \$350,000, \$500,000, \$1,000,000\}$. When actual earnings data are used, business income is treated as labor earnings because we suspect that many households take business income as salary.³⁹ While generating the tax data, we assume that married households have 2 dependents while single households have none.

With statutory income taxes in hand for a variety of earnings scenarios, we then compute average household income tax rates as the sum of federal and state income taxes divided by total earnings. The average income tax rate is the relevant measure for discrete choices between sectors, as we do not model hours or effort choices. Using the exact tax rate for an individual household would require adding state variables to the dynamic model. Rather than add state variables for marital status, dependents, or the exact tax situation in a given year, we allow average tax rates to vary with age, earnings, and permanent paid earnings ability, α . Conditioning on permanent paid ability allows for fluctuations through demographic factors that covary with permanent earnings capacity. Conditioning on age allows for lifecycle fluctuations in the tax rate. To capture this variation, we regress the average income tax rate observed for a household in year t on a third-order polynomial in age, a third-order polynomial in α , and a fifth-order polynomial in the head’s earnings. In estimating the model and forward simulating earnings, we compute expected taxes using the fitted values of the average rate function.

While this approach ignores some variation across years and due to demographics, for our pur-

³⁹The effective tax rate for most small business owners is very similar to individual income taxes; many small business are taxed as pass-through entities, equalizing tax schedules across paid work and self-employment. Businesses set up as C corporations will face a different tax schedule on any earnings that are not distributed as salary; these earnings are subject to the corporate tax rate, and distributions are subject to shareholder taxes.

pose of simulating choices and computing the after-tax returns to entrepreneurship, the question is whether these details affect statutory after-tax earnings differentially for paid and self-employment. We suspect any differences are minor.

B.3 Estimates of Flow Payoff Parameters

In the second stage, we take $\hat{\theta}$ as given and maximize the likelihood function to estimate $\hat{\beta}$, $\hat{\mu}_{\beta_0}$, and $\hat{\sigma}_{\beta_0}^2$. In each interview, the PSID asks about employment and earnings in the last calendar year and about employment at the time of the interview, generally in late spring. This second stage maximizes the likelihood of being self-employed at the time of the interview, conditional on the state variables at the end of the past calendar year.⁴⁰

We follow Rust (1987) and assume that the shocks ε_{it}^0 and ε_{it}^1 are serially independent with Type-1 extreme value distributions. This gives a conditional logit form for the conditional value function

$$\begin{aligned} v(d_t, S_t, \beta_{0i}, \varepsilon_t; \beta) &= u(d_t, S_t, \beta_{0i}, \varepsilon_t) + \delta E_t [\max \{v(0, S_{t+1}, \beta_{0i}, \varepsilon_{t+1}; \beta), v(1, S_{t+1}, \beta_{0i}, \varepsilon_{t+1}; \beta)\}] \\ &= u(d_t, S_t, \beta_{0i}, \varepsilon_t) + \delta \int \log \left(\sum_j \exp [v(d_{t+1}, S_{t+1}, \beta_{0i}, \varepsilon_{t+1}; \beta)] \right) df(S_{t+1} | S_t, d_t; \theta) + \delta \gamma, \end{aligned} \quad (18)$$

where γ is Euler's constant. We use these conditional value functions, Equation (12), and estimates of $\hat{\theta}$ to maximize the likelihood in Equation (14) for $\hat{\beta}$, $\hat{\mu}_{\beta_0}$, and $\hat{\sigma}_{\beta_0}^2$. We set the discount rate at $\delta = \frac{1}{1.10}$. The integral is over next year's stochastic state variables: the persistent shock in the paid sector and the belief of entrepreneurial ability. We approximate both distributions using 5-point Gaussian quadrature.

The likelihood of the choices is then given by the conditional logit formula, $p_t(d_t = 1 | S_t, \beta_{0i}; \beta) = \frac{\exp(v(1, S_t, \beta_{0i}; \beta))}{\exp(v(1, S_t, \beta_{0i}; \beta)) + \exp(v(0, S_t, \beta_{0i}; \beta))}$. To get the marginal likelihood for sequences of choices, we inte-

⁴⁰This approach does not require linking interviews in two subsequent years and is therefore not affected by the PSID's transition to bi-annual interviews after 1997. We adjust sector choice if later interviews indicate that a worker spent most of the year in one sector but reported the other at the time of the interview.

grate over the distribution of β_{0i} ,

$$\mathcal{L}(d_i|\beta, \theta, \mu_{\beta_0}, \sigma_{\beta_0}^2) = \int \prod_{t=1}^T p_t(d_t = 1 | S_t, \beta_{0i}; \beta)^{d_t=1} p_t(d_t = 0 | S_t, \beta_{0i}; \beta)^{d_t=0} d\phi\left(\frac{\beta_{0i} - \mu_{\beta_0}}{\sigma_{\beta_0}}\right).$$

Again, we approximate the integral over β_{0i} using 5-point Gauss-Hermite quadrature. We then maximize $\sum_{i=1}^N \log(\mathcal{L}(d_i|\beta, \theta))$, the log likelihood, with respect to β, μ_{β_0} , and $\sigma_{\beta_0}^2$.

B.4 Posterior Preference Distributions

The parameters governing the distribution of β_{0i} describe the population distribution of preferences for entrepreneurship. It is also possible to estimate where in this distribution each individual in the PSID sample is likely to fall based on their estimated earnings potential and their observed sequence of sector choices. Define $h(\beta_{0i}|d_i, S_i; \mu_{\beta_0}, \sigma_{\beta_0}^2)$ as the posterior probability density of preferences for individual i , given their history of choices, d_i , and sequence of state variables, S_i . As derived in Train (2009) using Bayes' rule, this posterior distribution is determined by

$$h(\beta_{0i}|d_i, S_i; \mu_{\beta_0}, \sigma_{\beta_0}^2) = \frac{\mathcal{L}(d_i|\beta, \theta, \beta_{0i}) \phi\left(\frac{\beta_{0i} - \mu_{\beta_0}}{\sigma_{\beta_0}}\right)}{\mathcal{L}(d_i|\beta, \theta, \mu_{\beta_0}, \sigma_{\beta_0}^2)}, \quad (19)$$

where $\mathcal{L}(d_i|\beta, \theta, \beta_{0i}) = \prod_{t=1}^T \mathcal{L}_t(d_{it} | S_{it}, \beta_{0i}; \beta, \theta)$ is the probability of a sequence of choices conditional on having preference β_{0i} , $\phi\left(\frac{\beta_{0i} - \mu_{\beta_0}}{\sigma_{\beta_0}}\right)$ is the population probability distribution of that preference, and $\mathcal{L}(d_i|\beta, \theta, \mu_{\beta_0}, \sigma_{\beta_0}^2)$ is the unconditional probability of those choices as defined in Equation (19).

The mean and variance of the individual posterior distributions are given by

$$\mu_{\beta_{0i}} = \int \beta_{0i} h(\beta_{0i}|d_i, S_i; \mu_{\beta_0}, \sigma_{\beta_0}^2) d\beta_{0i} = \frac{\int \beta_{0i} \mathcal{L}(d_i|\beta, \theta, \beta_{0i}) \phi\left(\frac{\beta_{0i} - \mu_{\beta_0}}{\sigma_{\beta_0}}\right) d\beta_{0i}}{\mathcal{L}(d_i|\beta, \theta, \mu_{\beta_0}, \sigma_{\beta_0}^2)} \quad (20)$$

and


$$\sigma_{\beta_{0i}}^2 = \frac{\int (\beta_{0i} - \mu_{\beta_{0i}})^2 \mathcal{L}(d_i|\beta, \theta, \beta_{0i}) \phi\left(\frac{\beta_{0i} - \mu_{\beta_0}}{\sigma_{\beta_0}}\right) d\beta_{0i}}{\mathcal{L}(d_i|\beta, \theta, \mu_{\beta_0}, \sigma_{\beta_0}^2)}. \quad (21)$$

We can then calculate the individual posterior likelihoods using Equation (19) and substituting these individual posterior distributions in for the population distribution of preferences.

Table A1: The Value of Entering Self-Employment for all Ages				
	Value of Paid	Value of Entrep.	Difference	Pct. Difference
Annual observed earnings in chosen sector				
Mean	46,250	58,610	12,360	26.7%
Median	41,025	37,586	-3,439	-8.4%
N	54,266	6,326		
Projected annual earnings for all workers				
Mean	46,623	41,029	-5,594	-12.0%
Median	42,536	36,075	-6,460	-15.2%
N	60,593	60,593		
Projected lifetime earnings, static model				
Mean	49,113	46,499	-2,614	-5.3%
Median	45,818	40,573	-5,246	-11.4%
N	60,593	60,593		
Projected lifetime earnings, dynamic model				
Mean	49,795	50,328	533	1.1%
Median	45,642	44,838	-804	-1.8%
N	60,593	60,593		


The first panel of this table summarizes observed annual earnings paid workers and entrepreneurs. In this panel the workers contributing to the paid averages are different from the workers contributing to the entrepreneurial average. In the remaining panels, earnings are projected for all workers conditional on choosing each sector, so the same sample of workers contributes to the averages for each sector. For ease of comparison, the projected lifetime earnings are converted to constant-annual-income equivalents, $\bar{C}$ such that $\sum_{s=t}^T \left(\frac{1}{1+r}\right)^s Y_s = \bar{C} \sum_{s=t}^T \left(\frac{1}{1+r}\right)^s$. The third panel projects discounted lifetime earnings assuming that workers choose each sector this year and remain there for all future years. The fourth panel projects discounted lifetime earnings assuming workers choose each sector this year and then behave optimally according to the full model in each subsequent year.

Figure A1: Relative Earning Profiles by Persistence in Entrepreneurship




Source: PSID 1976-2011. Profiles are average observed annual earnings, less estimated taxes, for entrepreneurs relative to their observed annual earnings, less estimated taxes, in their last year of paid work. For this plot, we use only workers who we observe both in self-employment and in paid work no more than three years before they entered self-employment. Estimated taxes are described in Appendix B.2. Weighted as described in the text.

Figure A2: Prediction Error for Entrepreneurial Ability



Source: PSID 1976-2011. The first distribution in this figure plots the difference between the average observed residual entrepreneurial earnings and the residual predicted by average paid earnings for workers observed in both sectors. The second line plots the theoretical distribution of this prediction error, under the assumption of bivariate normality, using the estimated variances and covariance reported in Table 4.

Figure A3: Predicted Entrepreneurial Earnings Ability



Source: PSID 1976-2011. Plots average earnings residuals in paid work and self-employment for the subset of workers who we observe in both sectors.